

**MODELING DYNAMIC ALLOCATIONS
AND DEALLOCATIONS OF LOCAL
MEMORY FOR TRANSLATION
VALIDATION**

ABHISHEK ROSE



DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING
INDIAN INSTITUTE OF TECHNOLOGY DELHI
FEBRUARY 2025

MODELING DYNAMIC ALLOCATIONS AND DEALLOCATIONS OF LOCAL MEMORY FOR TRANSLATION VALIDATION

by

Abhishek Rose

Department of Computer Science and Engineering

Submitted

in fulfillment of the requirements of the degree of
Doctor of Philosophy

to the



Indian Institute of Technology Delhi
FEBRUARY 2025

Certificate

This is to certify that the thesis titled **Modeling Dynamic Allocations and Deallocations of Local Memory for Translation Validation** being submitted by **Mr. Abhishek Rose** for the award of **Doctor of Philosophy in Computer Science and Engineering** is a record of bona fide work carried out by him under my guidance and supervision at the Department of Computer Science and Engineering, Indian Institute of Technology Delhi. The work presented in this thesis has not been submitted elsewhere, either in part or full, for the award of any other degree or diploma.



Prof. Sorav Bansal

Microsoft Chair Professor

Department of Computer Science and Engineering

Indian Institute of Technology Delhi

New Delhi 110016

Acknowledgements

I would like to thank all who contributed to this dissertation.

My sincere gratitude goes to my supervisor, Prof. Prof. Sorav Bansal, for expert guidance, insightful feedback, and unwavering support.

I also appreciate the valuable input of my committee members, Prof. Sanjiva Prasad, Prof. Rahul Purandare, and Prof. Subodh Sharma, whose critiques greatly improved this work.

I am grateful to my fellow students, Shubhani, Indrajit, Vaibhav, and Pratik, for their collaboration and stimulating discussions.

I also acknowledge the staff of the Department of Computer Science and Engineering for their administrative and technical support.

Finally, I thank the reviewers for their thorough evaluations and constructive comments, which helped refine the final dissertation.

Abhishek Rose

Abstract

End-to-End Translation Validation is the problem of verifying the executable code generated by a compiler against the corresponding input source code for a single compilation. This becomes particularly hard in the presence of dynamically-allocated local memory where addresses of local memory may be observed by the program. In the context of validating the translation of a C procedure to executable code, a validator needs to tackle constant-length local arrays, address-taken local variables, address-taken formal parameters, variable-length local arrays, procedure-call arguments (including variadic arguments), and the `alloca()` operator.

We make the following contributions in our work:

1. A formalization of the execution semantics for an unoptimized intermediate representation (IR) of a C program and its compiled 32-bit x86 assembly in the presence of dynamically (de)allocated local memory. This includes modeling of the various dynamic allocation constructs in C, such as address-taken local variables, constant- and variable-length local arrays, address-taken formal parameters, procedure-call arguments (including variadic arguments), and the `alloca()` operator.
2. A notion of correct translation from the IR to the assembly through a refinement definition. The definition incorporates the concept of undefined behavior (UB) within the IR program, originally translated from C, where refinement is permitted to hold

trivially.

3. An algorithm that converts the correct translation check to first-order logic queries over bitvectors, arrays, and uninterpreted functions that can be discharged using off-the-shelf SMT solvers. The algorithm is capable of operating in both blackbox and whitebox modes, with the blackbox mode enabling its usage with third-party compilers that may not employ a specific allocation strategy, such as preallocation. In particular, we are perhaps the first to enable support for dynamic stack allocation strategy for procedure-call arguments used by almost all production compilers (e.g., GCC, Clang/LLVM).
4. A prototype implementation of the algorithm and its comprehensive evaluation on a set of diverse benchmarks, including both micro-benchmarks and a real-world `bzip2` program. Our prototype performs blackbox translation validation of C procedures with up to 100+ SLOC against their corresponding assembly implementations with up to 140+ instructions generated by an optimizing production compilers (such as GCC, Clang/LLVM, ICC) with complex loop and vectorizing transformations.

Contents

Acknowledgements	iii
Abstract	v
List of Figures	xviii
1 Introduction	1
1.1 Problem Statement and Motivating Example	4
1.1.1 An address-taken local example	5
1.1.2 Establishing Correct Translation	7
1.1.3 Subtleties	9
1.1.4 A sketch of proposed solution	12
1.2 Prior Work	13
1.2.1 IR-to-IR Translation Validation	13
1.2.2 IR-to-Assembly and Assembly-to-Assembly Translation Validation	14

1.2.3	Verified Compilation	15
1.3	Contributions	15
1.4	Outline	17
2	Execution Semantics and Notion of Correct Translation	19
2.1	Intermediate Source and Assembly Representations	19
2.1.1	Unoptimized IR	21
2.1.2	Assembly	25
2.1.3	Allocation and Deallocation	25
2.2	Transition Graph Representation	26
2.2.1	Address Set	27
2.2.2	Memory Regions	27
2.2.3	Ghost Variables	29
2.2.4	Error Codes	30
2.2.5	Outside world and observable trace	30
2.2.6	Expressions	31
2.2.7	Graph Instructions	31
2.3	Translations of C and A to their Graph Representations	33
2.3.1	Translation of C	34
2.3.2	Translation of A	40

2.4	Observable traces and Refinement Definition	44
2.5	Refinement Definition in the presence of local variables and procedure calls when all local variables are allocated on the stack in A	47
2.5.1	(De)Allocation indicating alloc_s and dealloc_s instructions	48
2.5.2	Annotated procedure-call instruction	50
2.5.3	Refinement Definition with only stack-allocated locals and procedure calls	51
2.5.4	Capabilities and Limitations of $C \dot{\supseteq} A$	53
2.6	Refinement in the presence of potentially register-allocated or eliminated local variables in A	57
2.6.1	Virtual (de)allocations through alloc_v and dealloc_v instructions	57
2.6.2	Revised semantics for assembly procedure instructions	60
2.6.3	Refinement Definition with both stack-allocated and register-allocated or eliminated locals	62
2.7	Towards A More General Refinement Definition and Execution Semantics	66
2.7.1	Comparison with $C \ddot{\supseteq} A$	69
3	Witnessing Refinement through a Determinized Cross-Product	75
3.1	Program Paths	76
3.2	Determinized Product Graph as a Transition Graph	77
3.3	Analysis of the determinized product graph	78

3.3.1	X requirements	79
3.3.2	Soundness of X requirements	82
3.3.3	Global Invariants in C, \ddot{A} , and X	88
3.4	Callers' Virtual Smallest Semantics	90
3.4.1	Soundness of Callers' Virtual Smallest semantics	91
3.5	Safety-Relaxed Semantics	93
3.5.1	Soundness of Safety-Relaxed Semantics	94
4	Automatic Construction of a Product-Program	101
4.1	The DYNAMO algorithm	101
4.1.1	Enumerating A paths	105
4.1.2	Correlating C paths	107
4.1.3	Identifying A annotation	114
4.1.4	Validating structure of identified paths	118
4.1.5	Incremental construction of (\ddot{A}, X)	120
4.1.6	Checking requirements on partial X	122
4.1.7	Correlating paths to error nodes due to annotated instructions	125
4.1.8	Soundness of DYNAMO algorithm	125
4.1.9	Counterexample Guided Best-First Search	125
4.2	Invariant Inference	126

4.2.1	Global Invariants	128
4.3	Running Example of the Algorithm	130
5	SMT Encoding	141
5.1	Preliminary Steps	141
5.2	Representing address sets using allocation state array	143
5.2.1	Encoding of address set updating instructions	144
5.2.2	Full-array encoding	145
5.3	Interval Encoding	146
5.3.1	Interval encoding for $r \in G \cup F \cup Y \cup Z_l \cup \{stk\}$	146
5.3.2	Interval encoding for $r \in \{hp, cl, cs\}$	146
5.3.3	Soundness of Interval Encoding	148
5.4	Semantics with Simpler SMT Encoding for <i>stk</i> Region of \ddot{A}	157
6	Evaluation	161
6.1	Implementation of DYNAMO	161
6.1.1	System components	161
6.1.2	Discharging Proof Obligations	163
6.1.3	Pseudo-register allocation in LLVM _d	164
6.1.4	Instrumentation of Clang/LLVM for generating annotation hints	164

6.2	Experiments	164
6.2.1	Evaluating efficacy of DYNAMO	165
6.2.2	Evaluating modeling cost of local allocations	168
6.2.3	Evaluating DYNAMO on a real-world program	170
6.2.4	Analysis of Failures	172
6.3	Other Applications	178
7	Conclusion	181
7.1	Summary	181
7.2	Limitations and Directions for Future Work	182
	Appendices	185
A	Soundness and Completeness Implications of isPush() Choice	187
A.1	\mathbb{K} needs to be at least 2^{d-1} in the presence of VLAs	188
A.2	$\mathbb{K} = 2^{d-1}$ can still lead to completeness problems	189
A.3	$\mathbb{K} = 2^{d-1}$ can also lead to soundness problems	189
A.4	Solution	190
B	More details of the experiments	191
B.1	Command-line used for compiling benchmarks in experiments	191
B.2	Full results for the <code>bzip2</code> experiment	195

CONTENTS	xiii
C Full source code of the benchmarks	199
List of Publications	205
Biography	207
Bibliography	218

List of Figures

1.1	Two ways of obtaining a certified executable from an input source code. The certified executable is guaranteed to have identical semantics as the input source code.	3
1.2	C program with an address-taken local and its IR and 32-bit x86 assembly lowerings. <code>@STR</code> denotes the address of the format string <code>"%d"</code> . <code>mem₄[<i>a</i>]</code> denotes a 4-byte memory access to address <i>a</i>	5
1.3	C program with variable-length array (VLA) and its assembly lowerings. <code>mem₄[<i>a</i>]</code> denotes a 4-byte memory access to address <i>a</i> . An execution where $m \geq n$ triggers undefined behavior in the C program and C semantics do not put restriction on behavior of the translated assembly program in such an execution.	10
1.4	Conceptual representation of the flow of information from high-level source program and the low-level assembly program to the compiler/-translation validator.	11
1.5	High-level components and flow of our translation validation approach.	12

2.1	C program with variable-length array (VLA) and its lowerings to unoptimized IR and assembly. Subscript ‘ s ’ denotes signed comparison. Red font (parts of) instructions in assembly are added by our algorithm.	20
2.2	Pseudo-code for translation of a C procedure-call expression to LLVM _d instructions. <code>alloc</code> , <code>dealloc</code> , <code>call</code> , <code>va_start_ptr</code> , etc. are LLVM _d instructions.	23
2.3	Translation of C’s variadic macros to LLVM _d instructions. <code>roundup₄(a)</code> returns the closest multiple of 4 greater than or equal to a	24
2.4	Translation rules for converting LLVM _d instructions to graph instructions. <code>op</code> represents an arithmetic, logical, or relational operator. c represents a constant.	36
2.5	Translation rules for converting LLVM _d instructions to graph instructions.	38
2.6	Translation rules for converting pseudo-assembly instructions to graph instructions. <code>op</code> represents an arithmetic, logical, or relational operator.	41
2.7	Translation rules for converting pseudo-assembly instructions to graph instructions.	43
2.8	Additional translation rules for converting pseudo-assembly instructions to graph instructions for procedures with only stack-allocated locals.	49
2.9	Example of transformation where relative order of (de)allocations and procedure calls is not preserved. The refinement definition will not admit the hypothetical assembly but will admit the compiler generated one.	54
2.10	Translation rules for converting the <code>alloc_v</code> and <code>dealloc_v</code> instructions to graph instructions.	58

2.11	Revised translation rules for converting pseudo-assembly instructions to graph instructions. The $\text{IF}\{z \in Z_l\}\{\dots\}\text{ELSE}\{\dots\}$ construct selects one of the translation depending on the result of syntactic predicate $z \in Z_l$.	60
2.12	Translation rules for the converting pseudo-assembly instructions to graph instructions for \ddot{A} . (ALLOCV') is derived from (ALLOCV) in fig. 2.10.	67
4.1	C source fragment and its abbreviated control-flow graph.	113
4.2	Predicate grammar for constructing candidate invariants. v represents a bitvector variable (registers, stack slots, and ghost variables), c represents a bitvector constant. $\odot \in \{\leq_{s,u}, <_{s,u}, >_{s,u}, \geq_{s,u}\}$. v^* represents a bitvector value drawn from a restricted grammar (explained in text).	127
4.3	Global invariants that hold at each non-entry, error-free node $n_x \in \mathcal{N}_x^{\text{DW}}$.	128
4.4	Reproduced C program and its unoptimized IR and assembly from fig. 2.1.	131
4.5	Abbreviated Transition Graphs for the unoptimized IR and assembly of the <code>fib</code> procedure from fig. 4.4.	132
5.1	Revised translation rules for the new stk^+ -based semantics for \ddot{A}	158
6.1	High-level components of DYNAMO implementation. Trusted Code Base (TCB) blocks have double border and a red background.	162
6.2	Comparison of running times of procedures in table 6.1 with full-interval (filled bars), partial-interval (thick black lines), and full-array (empty bars) encoding. Y-axis is logarithmically scaled.	167
6.3	Summary of refinement check results for the programs in table 6.1.	168

6.4	Procedure <code>s122</code> from ‘ <code>globals</code> ’ version of (modified) TSVC suite. . . .	169
6.5	Comparison of running times of TSVC benchmarks with exactly same code modulo allocation strategy. Y-axis is logarithmically scaled. . . .	170
6.6	Scatter plot of refinement time (in minutes) vs assembly lines of code (ALOC). Both axes are logarithmically scaled.	172
6.7	<code>vs11</code> procedure from table 6.1 (appears as <code>vs1N</code>) and the control-flow graph (CFG) of its GCC compiled assembly.	174
C.1	Benchmarks with VLAs.	200
C.2	Benchmarks with use of <code>alloca</code>	201
C.3	Benchmark <code>rod</code> with mixed use of VLA and address-taken variable. . .	202
C.4	Benchmarks <code>mp</code> and <code>ms</code> with variable argument list. <code>mp</code> is adapted from <code>minprintf</code> of K&R[24]	203
C.5	Structure of <code>bzip2</code> ’s functions	204

Chapter 1

Introduction

Safety and mission-critical software systems such as those found in medical equipment and nuclear reactors require strong correctness guarantees. Techniques in formal methods can provide mathematically-sound guarantees on the behavior of a software artifact. Absence of null-pointer dereference, buffer and integer overflow are examples of behavior-bounding guarantees ensured by formal techniques. Such techniques are typically applied on the source code, e.g., C source code, which is different from the machine executable code, e.g., x86 assembly code, that is executed on physical hardware. The translation of the source code into machine executable code is performed by a compiler in an act of compilation. A compilation (is expected to) preserves the *semantics* of the input source code in the output executable code. This *semantics preservation* or *correctness* property of compilation enables transfer of a behavior bounding guarantee, assured by formal methods techniques applied on the source code, from the source code to the executable code. A bug in the compilation process, however, can impede this transfer thereby giving a false sense of assurance. For example, a miscompilation may introduce an error that was absent in the original source program. Thus, application of formal methods on the source code is ineffective until the compiler or, more specifically, a desired instance of compilation itself is validated to be correct.

Besides upholding the semantic preservation or correctness property, compilers are also expected to perform optimizing transformations (optimizations) such that the compiled program exhibits certain performance characteristics, e.g., fast execution, power and memory efficiency, etc. An overwhelming majority of the complexity of modern optimizing compilers lies in this optimizing phase [16]. This additional responsibility of

optimization makes ensuring correctness harder. Indeed, correctness bugs in optimizing compilers are not uncommon [52, 46, 27, 45].

A way to validate a compilation, to ensure *semantics preservation*, is to formally verify the compiler itself. Such a *verified compiler* will only produce an output executable code, a *certified executable*, if it can ensure, through a formal argument, that the semantics of the input source code are preserved in the output executable code. However, writing a verified compiler is a formidable task. Modern non-verified optimizing compilers such as GCC [15] and Clang/LLVM[10] are multi-million lines of source code projects with thousands of changes being made with every release. Verifying the semantic preservation property on them or rewriting them to be easily verifiable while keeping with the pace of development seems impractical. Though there exist verified compilers such as CompCert [29] and CakeML [25], the optimization they perform and the source code language constructs they support are relatively limited in scope, prohibiting their large scale adoption. Another difficulty in development of verified compilers lies in the amount of expertise they require: writing formal proofs using proof assistants and/or using semi-automated verification tools arguably requires higher expertise than required of a typical compiler engineer.

Translation validation (TV for short) [37, 36] is an alternative approach where instead of verifying the entire compiler, an instance of compilation or translation is verified. A translation validation tool, i.e., a translation validator, takes two programs as input: (1) the *source* program passed as input to a translator; and (2) the *target* program emitted as output by the translator. The *target* program is compared against the input *source* program to verify the *semantics preservation* property. The validation, if successful, is accompanied by a proof of correctness that can be independently verified. The generation of an externally verifiable artifact allows the TV tool itself to potentially remain unverified while still being able to produce sound results. In the *end-to-end translation validation* setting, *source* represents the source code of a program written in a high-level language, such as C, and *target* represents the translated machine executable code, such as x86 machine code, obtained after compiling *source*. In contrast, translation validation can also be performed in a “pass-by-pass” manner, where the validator is invoked at the end of a transformation pass in the compiler. In this setting, *source* and *target* denote the programs before and after the transformation, respectively.

A verified compiler can also be obtained using a combination of a non-verified compiler and an *end-to-end* translation validator; fig. 1.1b shows a construction. In such a verified

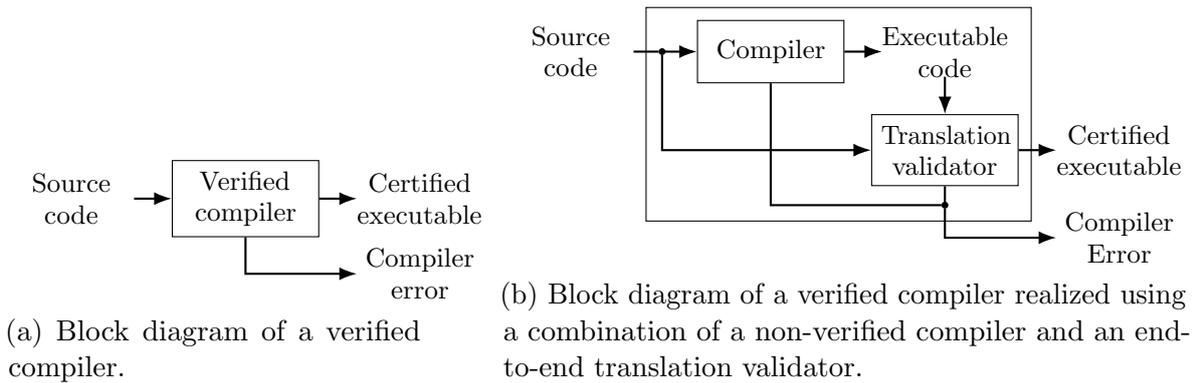


Figure 1.1: Two ways of obtaining a certified executable from an input source code. The certified executable is guaranteed to have identical semantics as the input source code.

compiler, the end-to-end translation validator is invoked at the end of each compilation by the non-verified compiler to validate the output compilation. If the translation validator is able to successfully validate the compilation, the combined tool produces a certified executable. This separation of compilation and its validation afforded by translation validation makes the combined construction an appealing practical alternative to the verified compiler approach.

A major challenge in the construction of an end-to-end translation validator lies in the difference between the levels of abstractions in the two input programs' representations. In the context of translation from an *unoptimized intermediate representation (IR)* of a C program to an assembly code representation of executable code, which is the focus of this work, this manifests as difference in the abstractions available in the C programming language (equivalently, in its unoptimized IR) and the assembly language.

Prior work on translation validation relates the two programs through construction of a product program that correlates the program points of the constituent programs. While the control-flow constructs could potentially be rather directly mapped to or modeled through a combination of conditional and unconditional jump operations available in the assembly language[57, 6, 54], the translation or modeling of concepts such as objects (their allocation and deallocation), types (that may affect aliasing), and pointer provenance [19] (for determining aliasing) is not immediately clear. For example, instead of a potentially infinite set of pseudo-registers or local variables, assembly has a finite set of machine registers and a stack region delineated by a stackpointer

register. Furthermore, C allows taking (and observing) address of a variable. Translation validation of programs with such constructs requires identification of relations between stack addresses in assembly and addresses of the variables in the C program. This identification and modeling of the relation between an address-taken variable in C and its corresponding implementation somewhere in the assembly stack is not straightforward. The difficulty of this problem has been acknowledged in previous work on end-to-end translation validation with [40] calling it their “largest limitation” and other works assuming that the input programs do not exhibit such programming patterns [8, 17]. We motivate and describe the problem clearly through an example.

1.1 Problem Statement and Motivating Example

The goal of this work is end-to-end validation of a compilation from a C program to an assembly program where the C program may contain dynamic allocation and deallocation of local memory through local variable declaration including C99’s variable-length arrays (VLAs), procedure-call arguments including variable-length arguments (variadic arguments), and the `alloca()` operator [2]. We transform the C program into an unoptimized intermediate representation (IR) where the allocation and deallocation of local memory is made explicit through IR instructions, and implementation-defined constructs such as types and their representation have been determinized. This unoptimized IR is obtained using a syntactic translation and does not involve any optimization. We use 32-bit x86 assembly with `cdec1` calling conventions, described in System V Application Binary Interface (ABI) for i386 systems [32], as a concrete assembly program representation. The translation from C to unoptimized IR uses the specifics of 32-bit x86 assembly ABI¹.

Thus, we are interested in showing that a 32-bit x86 assembly program \mathbb{A} is a correct translation of the unoptimized IR \mathbb{C} of a C program where \mathbb{C} may dynamically allocate and deallocate local memory and may contain constructs such as constant-length and variable-length local arrays, address-taken local variables and formal parameters, procedure-call arguments including variable-length (variadic) arguments, and the `alloca()` operator.

¹For example, the pointer type `void*` is determinized to a 32-bit bitvector representation matching the pointer representation in 32-bit x86.

```

C0:  int read_int() {
C1:    int x;
C2:    scanf("%d", &x);
C3:    return x;
C4:  }

```

(a) C program with an address-taken local variable `x`.

```

I0:  int read_int():
I1:    x = alloc int;
I2:    t1 = @STR; // address of "%d"
I3:    t2 = call scanf(t1, x);
I4:    r = *x;
I5:    dealloc I1;
I6:    ret r;

```

(b) (Abstracted) Unoptimized IR.

```

A0:  read_int:
      ; allocate 36 bytes on stack
A1:  esp -= 36;
A2:  eax = esp + 20;
      ; push procedure arguments on stack
A3:  push eax;
A4:  push @STR; ; address of "%d"
      ; (implicit) arguments to scanf:
      ; mem4[esp], mem4[esp+4]
A5:  call scanf;
A6:  eax = mem4[esp+28];
      ; deallocate 44 bytes from stack
A7:  esp += 44;
      ; return value in register eax
A8:  ret;

```

(c) (Abstracted) 32-bit x86 assembly of fig. 1.2a.

Figure 1.2: C program with an address-taken local and its IR and 32-bit x86 assembly lowerings. `@STR` denotes the address of the format string `"%d"`. `mem4[a]` denotes a 4-byte memory access to address *a*.

1.1.1 An address-taken local example

Consider the C program and its unoptimized IR and 32-bit x86 assembly shown in fig. 1.2. The `read_int()` procedure in fig. 1.2a defines a local integer `x`, passes the address of `x` to an external callee `scanf()`, and finally returns the value of `x`. Because the address of variable `x` is taken, through the *address-of* operator `&`, `x` is an *address-taken local*.

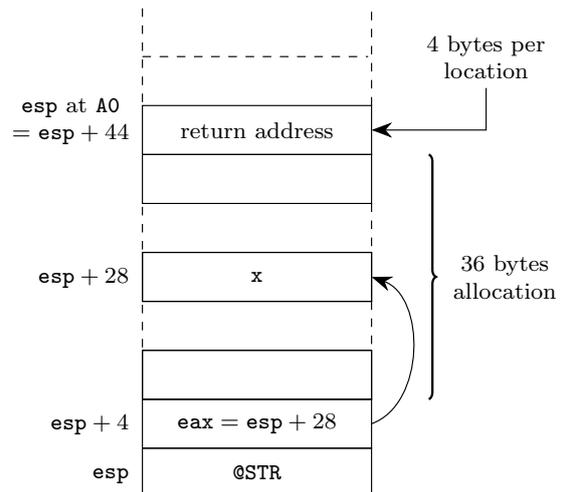
Unoptimized IR: The unoptimized IR of `read_int()`, shown in fig. 1.2b, is similar to LLVM IR [48] and has explicit instructions `alloc` and `dealloc` for local memory allocation and deallocation — the latter instruction is not available in LLVM IR. The `alloc` instruction in line I1 performs allocation of an `int` object and returns the start address of the allocated region in `x`. Both the allocated region and its initial memory contents are non-deterministic but subject to two *well-formedness conditions*:

- (1) the freshly allocated region is disjoint from other currently allocated regions;
- (2) the start address of the allocated region is aligned by the alignment requirement for the allocated type (e.g., `x` of type `int` should be aligned by 4).

An allocation due to execution of `alloc` is expected to satisfy these well-formedness conditions. Unlike C, the identifier `x` in IR represents the address of variable `x` so that the call to `scanf` in line I3 accepts `x` as argument (and not `&x` as in fig. 1.2a). The value of the variable `x` is obtained by dereferencing (`*x` in line I4). An allocation is uniquely identified by the location or PC of the `alloc` instruction. The `dealloc` instruction deallocates *all* allocated space due to execution(s) of an `alloc` instruction identified by its PC. In fig. 1.2b, `dealloc I1` deallocates the address region allocated by execution of instruction I1.

Assembly: Figure 1.2c shows the (abstracted) 32-bit x86 assembly of `read_int()` as generated by GCC. In the compiled assembly, the local `x` is allocated on the assembly stack which is a distinguished region in the machine address space. The lower, flexible end-point of the assembly stack region is identified by the stackpointer register `esp` so that decrementing `esp` *allocates* stack space and incrementing `esp` *deallocates* stack space. The stackpointer decrement in line A1 allocates 36 bytes on the stack — 4 bytes of this allocated space are reserved for the local `x` (an analysis of the next few instructions reveal that `x` is located at an offset of 20 bytes in this newly allocated space). Unlike its IR counterpart, the x86 `call` instruction does not accept arguments other than the callee label (or address); instead, the arguments are passed implicitly through the stack. The stack addresses used for passing the arguments are specified in the ABI’s calling conventions which enable identification of the passed arguments in the callee procedure.

The `push` instructions in A3 and A4 set up the two arguments to `scanf` on stack — notice that the arguments are pushed in reverse order so that A3 pushes the second argument, address of `x`, and A4 pushes the first argument `@STR`, address of the format string. The figure on the right shows the layout of stack just before the call instruction: at the bottom of stack, labeled by `esp`, is `@STR`, above it is `eax` which contains the address `esp+28` of `x` (assigned as `esp+20` in A2) and so on. The stackpointer increment in line A7 deallocates



the stack space allocated due to the two `push` instructions and instruction A1 in a single step (thereby deallocating both arguments to `scanf()` and the local `x`). The return

value of an assembly procedure is passed in register `eax` which is set to value of `x` in `A6` (just before deallocation).

1.1.2 Establishing Correct Translation

Let `C` denote the IR procedure `read_int()` in fig. 1.2b and `A` denote the assembly procedure `read_int()` in fig. 1.2c. In this section, we attempt to define a notion of correct translation and present an (loosely formal) argument that `A` is a correct translation of `C`. We will illustrate some key problems in validating a translation from `C` to `A` through this exercise.

According to the C standard [21], `A` will be considered a correct translation of `C` if for each execution of `A`, there exists an execution of `C` such that both executions produce identical *observables* — the observables produced by an execution of `C` constitute its *observable behavior* (para 6 in §5.1.2.3 of [21]). Notice that this definition implies a subset relationship for behaviors which makes sense because `C` is non-deterministic². Thus, correct translation is defined in terms of *refinement* of observable behaviors — the assembly program refines (for e.g., by determinization) the behavior present in the C program.

The observables produced by an execution of `C` are (in order):

- O1.** the occurrence of procedure-call to `scanf()`;
- O2.** the program state passed to or accessible to the external callee `scanf()`, which includes the arguments `@STR` and `x` (line I3 in fig. 1.2b);
- O3.** the return value `*x`.

Because the definition of the `scanf()` procedure is external to the translation unit, no assumptions can be made about its behavior. Therefore, we conservatively consider *both* the occurrence of the call to `scanf()` and the program state accessible to `scanf()` as observables.

To establish correct translation, each execution of `A` must produce observables such that an execution of `C` also produces identical observables (recall that `C` is non-deterministic). Consequently, an execution of `A` must produce the observables **O1**, **O2**, **O3** (in sequence). It can be observed that `A` produces **O1**: there is a call to `scanf()` at `A5` in fig. 1.2b

²Recall that the `alloc` instruction returns a region with non-deterministic start address and non-deterministic memory contents.

with no observable event³ before it. Thus, every execution of **A** at least produces **O1**. Next, we must witness production of observable **O2** in (an execution of) **A**.

Recall that **O2** observes the program state passed to `scanf()` which includes both the address and the value of variable **x**, both of which are non-deterministic in **C** — the value of **x** is included because passing the address makes the variable reachable from `scanf()` (through a dereference). To produce an identical observable, (an execution of) **A** must also *pass* identical program state to `scanf()`. Recall that in **A** the arguments are passed implicitly through stack so that the number of passed arguments is not identifiable solely from the `call` instruction in **A**; the calling conventions specify the addresses of the arguments but not their count. This problem is exacerbated by the fact that `scanf()` is a variadic procedure so it may accept a variable number of arguments and even the type signature of `scanf()` in **C** cannot be used for determining the number of arguments. Even if we assume that only two arguments (as set up in **A3** and **A4** of fig. 1.2c) are passed to `scanf()`, we need to relate the (stack) address `esp+20` (set up in **A3**) with the non-deterministic address **x** in **C**. While **x** is subject to the well-formedness conditions, there is no guarantee that `esp+20` respects the same conditions so that it is possible for `esp+20` and **x** to not agree on the same value — recall that we want to establish that an execution of **C** may always produce a value **x** that is equal to the value `esp+20` produced by an execution of **A**. Moreover, as the stack is shared by both the allocated locals and spilled pseudo-registers, discriminating different stack writes becomes a problem. For example, it needs to be ascertained that the stack writes at lines **A3** and **A4** of fig. 1.2c do not mutate the stack region corresponding to the local **x**⁴.

Further, because we are not able to ascertain the number of arguments passed to `scanf()` in **A**, we must conservatively assume that it may read an arbitrary number of arguments (by reading the stack addresses specified in the calling conventions). The external call to `scanf()` can be soundly (and over-approximately) modeled (in both **C** and **A**) as an arbitrary but deterministic mutation of the accessible program state subject to the input and accessible program state. With this modeling, the difference in number of arguments makes the observable **O2** in **A** different from the observable **O2** in **C**.

Lastly, the observable **O3** is also related to the modeling of procedure call to `scanf()`:

³Observable events include a call to an external procedure and a procedure return.

⁴While in this particular example it is easy to distinguish the writes due to use of the `push` instruction that writes to a freshly allocated region, doing so in presence of stack spills require tracking of addresses belonging to an allocated variable.

the return value of procedure **A** is obtained by reading machine memory at address `esp+28` immediately after the procedure call (line **A6** in fig. 1.2c); in **C** the return value is identified as the value `*x`. As both values are mutated arbitrarily by `scanf()`, due to difference in the number of arguments passed to `scanf()`, **O3** in **A** cannot be related to observable **O3** in **C**.

To summarize, we identified the following problems in relating the observables (**O2** and **O3**) in an execution of **A** with an execution **C**:

- In relating a stack allocation in **A** with an allocation in **C** so that a stack region in **A** can be distinctly identified and related to a non-deterministic but constrained allocated region in **C**.
- In identifying the behavior of a procedure call in **A** so that the program state accessible to an assembly callee can be precisely identified.

1.1.3 Subtleties

In this section, we take a closer look at the two programs in figs. 1.2a and 1.2c and consider some subtleties associated with a valid translation.

Assembly program may allocate more memory

Consider the stack allocation of 36 bytes at line **A1** in fig. 1.2c. It is possible for this stack allocation to overflow into other allocated space (e.g., heap) leading to abnormal termination of **A**'s execution. Recall that correct translation requires every observable behavior of **A**, including abnormal termination, to be an observable behavior of **C**. We observe that **C** in fig. 1.2b may similarly terminate due to allocation failure while executing `alloc`. Recall that an `alloc` instruction allocates a region that does not overlap with currently allocated space; if it is not possible to meet this requirement, the execution terminates abnormally. Notice, however, that **A** allocates more memory than **C**: 36 bytes in **A**, compared to 4 bytes in **C**⁵. Thus, it is possible for a given execution of an assembly program starting at `read_int()` to run out of memory but a similar execution of the unoptimized IR program to terminate normally. A correct translation definition, therefore, must take into consideration abnormal termination of an assembly program due to running out of memory as an admissible observable behavior.

⁵The extra allocation, in this particular example, is to align the stackpointer to a 16 byte boundary for the procedure-call at **A5**

<pre> C0: int foo(int n, int m) C1: { // VLA with n elements C2: char buf[n]; C3: buf[m] = 0; C4: return bar(buf, n); C5: } </pre>	<pre> A0: foo: A1: push ebp; ebp = esp; A2: esp -= 8; A3: edx = mem4[ebp+8]; ; argument 'n' A4: ecx = mem4[ebp+12]; ; argument 'm' ; variable-sized stack allocation A5: esp -= 0xFFFFFFFF & (edx+15); A6: mem1[esp+ecx] = 0; A7: eax = esp; A8: esp -= 8; A9: push edx; push eax; A10: call bar; A11: esp = ebp; pop ebp; A12: ret; </pre>
--	--

(a) C program with VLA

(b) (Abstracted) 32-bit x86 assembly

Figure 1.3: C program with variable-length array (VLA) and its assembly lowerings. $\text{mem}_4[a]$ denotes a 4-byte memory access to address a . An execution where $m \geq n$ triggers undefined behavior in the C program and C semantics do not put restriction on behavior of the translated assembly program in such an execution.

We wish to point out this peculiarity arises due to the finite machine state available to an assembly program: while an unoptimized IR program may utilize an infinite set of pseudo-registers, the assembly has a finite set of registers and has to spill additional state into the stack. A notable consequence of having such a definition is that it is easy to construct a *sound* but *vacuous* translation for every C program that always terminates abnormally due to out-of-memory.

C program may trigger Undefined Behavior (UB)

C has undefined behavior semantics [51]. Consider the C program and its compiled assembly in fig. 1.3. The `foo()` procedure in fig. 1.3a accepts two integers `n` and `m` as parameters, allocates a variable-length array (VLA) `buf` with `n` characters and calls an external procedure `bar`, passing `buf` as an argument. Figure 1.3b shows the compiled assembly of the `foo` procedure. The assembly procedure receives the two parameters `n` and `m` on stack and reads them into registers `edx` and `ecx` in A3 and A4 respectively. The VLA is allocated on stack via a stackpointer decrement in A5 — the extra arithmetic ensures that the resulting stackpointer value remains aligned by 16. The write to stack address `esp+ecx` in A6 corresponds to the zero initialization of m^{th} element of the VLA (C3 in fig. 1.3a) in the assembly program.

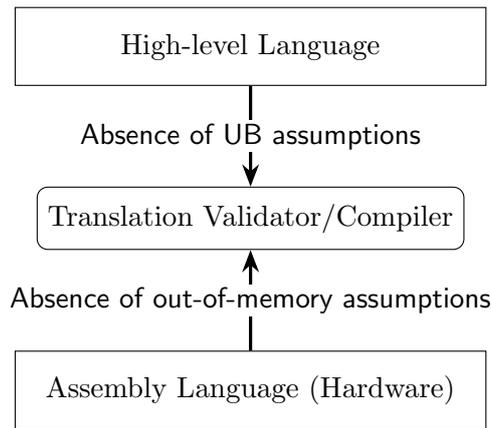


Figure 1.4: Conceptual representation of the flow of information from high-level source program and the low-level assembly program to the compiler/translation validator.

Let C and A denote the C and assembly `foo()` procedures in figs. 1.3a and 1.3b respectively. Consider an execution of A where $m \geq n$. In such an execution, the write at `A6` may potentially overstep stack bounds, accessing a memory location that is not mapped and thereby causing termination. Another possibility is that the write address remains within stack bounds, overwriting some address in the caller stack of the `foo()`, breaking some program invariant and eventually leading to some logic error. An execution of C under the same input ($m \geq n$) will trigger an out-of-bounds write at `C3`. An out-of-bounds access is undefined behavior (UB) in C and the C semantics do not put any restriction on valid translation for such executions. As a result, any behavior, including termination or a logic error, of the translated assembly program is considered a valid translation for a construct that causes undefined behavior in C. A correct translation definition must accommodate handling of undefined behavior in order to cater to the various possibilities available to the compiler.

Figure 1.4 summarizes the flow of information due to semantics of C and A . The undefined behavior semantics in C enable both the compiler and the translation validator to make certain “absence of UB” assumptions about the error-free execution of C — the information flowing downwards from “High-level Language” in fig. 1.4. Similarly, both the compiler and the translation validator may assume availability of adequate stack space; shown as “absence of out-of-memory assumptions” flowing from “Assembly Language (Hardware)” in fig. 1.4. A compiler performs a translation under these conditions and, therefore, the translation validator must assume them while validating the transformations produced by the compiler.

1.1.4 A sketch of proposed solution

Prior TV efforts identify a lockstep correlation between (potentially unrolled) iterations of loops in the two procedures to show equivalence [8]. These correlations can be represented through a *product program* that executes procedures **C** and **A** in lockstep, using a careful choice of program path correlations, to keep the machine states of both procedures related at the ends of correlated paths [54, 17].

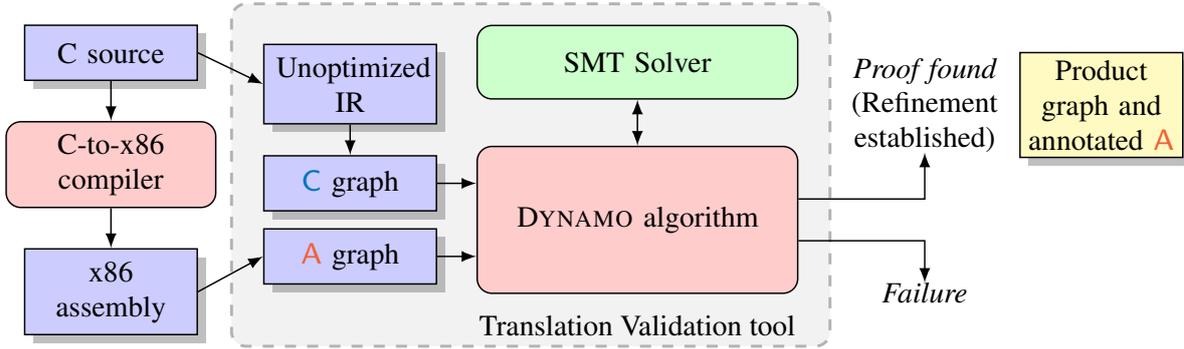


Figure 1.5: High-level components and flow of our translation validation approach.

Our TV algorithm, called DYNAMO, additionally attempts to identify a lockstep correlation between the *dynamic (de)allocation events* and *procedure-call events* performed in both procedures, i.e., we require the order and values of these execution events to be identical in both procedures. To identify a lockstep correlation, our algorithm annotates assembly procedure **A** with two kinds of annotations:

- (de)allocation instructions for identifying (de)allocation of a region in assembly;
- procedure-call annotations for identifying procedure-call arguments and memory regions accessible to an assembly callee.

Our key insight is to define a *refinement relation* between **C** and **A** through the existence of an annotated **A**. We carefully define the semantics for these annotations so that the refinement relation holds only if an annotation correctly captures the (de)allocation and procedure-call events in **A**. We also generalize the definition of a product program so it can be used to witness refinement in the presence of non-determinism due to addresses of dynamically-allocated local memory, undefined behavior (UB), and stack overflow in **A**. For example, our product program encodes the UB semantics that allow anything to happen in assembly (including out-of-bound stack access) if UB is triggered in **C**.

Our proof obligations are expressed in first-order SMT (Satisfiability Modulo Theories) logic over bitvectors, arrays, and uninterpreted functions that can be discharged by an off-the-shelf SMT solver. On successful execution, our algorithm produces two artifacts: (1) a product program encoding lockstep execution of unoptimized IR **C** and assembly procedure **A** in the form of a product graph and (2) an annotated version of **A** that satisfies the refinement relation with **C**. Figure 1.5 shows the high-level components of our translation validation approach.

1.2 Prior Work

Translation validation (TV) was proposed by Pnueli, Siegel, and Singerman [37] for validating translation of programs written in SIGNAL (a synchronous multi-clock data-flow language) to C code. This was followed shortly by Necula [36] who validated GCC optimization passes. Since then there has been much development in this area [54, 41, 49, 44, 56, 55, 35, 23]. TV has been successfully applied for finding bugs in real-world compilers [31, 34, 22], verifying complex loop unrolling and vectorizing transformations [8, 17], verification of individual passes inside a verified compiler [50], and, perhaps most ambitiously, in translation validation of a verified OS kernel [40]. However, no prior work tackled the problem of modeling dynamic local memory allocation and deallocation in the context of IR-to-assembly translation.

We give a brief survey of the prior TV approaches in rest of the section. We discuss their capability with respect to handling of dynamic local memory (de)allocations and highlight differences from an end-to-end IR-to-assembly TV such as ours.

1.2.1 IR-to-IR Translation Validation

An IR-to-IR translation validator (TV) validates transformation of a program represented in an intermediate language (IR)⁶ into a program represented using (usually) same or (rarely) different IR [36, 31, 49, 44, 56, 55, 34, 35, 23]. Unlike IR-to-assembly, modeling dynamic local memory allocations is significantly simpler for IR-to-IR TV. For example, (pseudo)register-allocation of local variables can be tackled by identifying relational invariants that equate the value contained in a local variable’s memory region

⁶In our discussion, we use intermediate language (IR) to denote a language suitable for use in target-independent phases of a compiler transformation pipeline. Typical characteristics of such a language include infinite pseudo-registers, types, and abstract mechanisms for representing allocation and deallocation of memory.

(in the original program) with the value in the corresponding pseudo-register (in the transformed program) [22, 23]. If the address of a local variable is observable by the C program (e.g., for an address-taken local variable), we need to additionally relate the variable addresses across both programs. These address correlations can be achieved by first correlating the corresponding allocation statements in both programs (e.g., through their names as identifiers) and then assuming that their return values are equal. Provenance-based syntactic pointer analyses, that show separation between distinct variables [3, 43], thus suffice for translation validation across IR-to-IR transformations.

Recent work on bounded TV⁷ by Lee et al. [28] presents an efficient SMT encoding for modeling memory where memory is segmented into *separate blocks* and each allocation is identified distinctly using a block-ID. As the number of allocations is bounded, a pointer is represented as a combination of a block-ID and an offset into a block. While this suffices for the bounded TV setting, our problem setting with potentially unbounded allocations (e.g., due to use of `alloca()` in a loop) requires a more general representation of a dynamically-allocated variable and a more general SMT encoding.

1.2.2 IR-to-Assembly and Assembly-to-Assembly Translation Validation

Prior work on IR-to-assembly and assembly-to-assembly TV [40, 41, 39, 13, 8, 17, 42] assumes that local variables are either absent or their addresses are not observed in the program and so they are removed through (pseudo)register-allocation. This assumption simplifies the validation effort as the required relations are between (pseudo)registers and spilled pseudo-registers on stack — the stack addresses of these spilled pseudo-registers are not relevant and do not need to be related.

A problem that arises uniquely in case of IR-to-assembly TV is disambiguation of a stack memory access from a *non-stack* (or heap) memory access. Sewell [39] employs a trusted static analysis over the assembly program to label each memory access as either aliasing with stack or not aliasing with stack. This distinction enables identification of the stack as separate from rest of the memory array such that any write to stack does not affect the rest of the memory state and vice versa. A similar approach of using two

⁷In bounded TV, the behavior of the input programs is bounded. For example, the programs are made acyclic by unrolling all loops by a finite unroll factor. This limits both execution and memory consumption of the programs. As a result, the validation guarantee is limited to executions that fall within the specified bounds.

memory arrays, one for stack and another for non-stack, is used in [42].

Other cases of address-taken parameters including the variadic parameter, where the correlation of the stack address representing the currently-iterated parameter in the variable-argument list is required, are also assumed to be absent in these prior works. Our work targets both constructs.

1.2.3 Verified Compilation

Prior work on verified compilation, embodied in CompCert [29], validates its own transformation passes from IR to assembly, and supports both address-taken local variables and variadic parameters. However, CompCert sidesteps the task of having to model dynamic allocations by ensuring that the generated assembly code *preallocates* the space for all local variables and procedure-call arguments at the beginning of a procedure’s body. Because preallocation is not possible if the size of an allocation is not known at compile time, CompCert does not support variable-sized local variables or `alloca()`. Moreover, preallocation is prone to stack space wastage. In contrast to a verified compiler, TV needs to validate the compilation of a third-party compiler, and thus needs to support an arbitrary allocation strategy. In particular, almost all production compilers use dynamic allocation strategy for procedure call arguments. An end-to-end TV must support dynamic (de)allocation to be a useful alternative.

1.3 Contributions

The central focus of this thesis is to investigate the applicability of translation validation as an alternative to verified compilation. Recent work [40, 17] has demonstrated that end-to-end translation validation is a feasible option for validating third-party compilations. However, as discussed in previous sections, the lack of support for handling (dynamic) local allocation restricts the applicability of the approach to the wide array of programs found in real-world. This work aims to rectify this limitation. We offer the following high-level contributions in this space.

- A formalization of the execution semantics for the unoptimized IR translation of a C program and its compiled 32-bit x86 assembly in the presence of dynamically (de)allocated local memory. This includes modeling of the dynamic local allocation constructs in C and assembly, such as constant and variable-length local arrays, and

`alloca()` operator in C and stack-allocated procedure-call arguments (including variadic arguments) in assembly. Our modeling is sound in the sense that we do not make any simplifying assumptions and model various constructs precisely. For example, we model C integers as bitvectors with wrap-around semantics (instead of mathematical integers) and the assembly stack is modeled as a part of the byte-addressable memory and not as a separate entity as done in prior work [40].

- A notion of correct translation from unoptimized IR to assembly through a refinement definition. The definition incorporates the concept of Undefined Behavior (UB) within the IR program (originally translated from C) where refinement is permitted to hold trivially when the IR program triggers UB. Similarly, the definition assumes the absence of stack overflow in assembly — the OS (or runtime environment) is expected to ensure the availability of sufficient stack space, and a translation validation is predicated on the availability of this space. A key limitation of this refinement definition is its failure to support interprocedural transformations (e.g., inlining, tail-call elimination, etc.).
- An algorithm that converts the correct translation check to first-order logic queries over bitvectors, arrays, and uninterpreted functions (AUFBV in SMT) through the automatic construction of a “product graph”. These first-order logic queries are discharged using off-the-shelf SMT solvers. The algorithm is capable of operating in both *blackbox* and *whitebox* modes, with the blackbox mode enabling its usage with third-party compilers with arbitrary allocation strategies (with some limitations described later). This is in contrast to the verified compiler CompCert [29] that only employs the preallocation strategy for local allocations. In particular, we are perhaps the first effort to enable support for dynamic stack allocation strategy for procedure-call arguments used by almost all production compilers (e.g., GCC, Clang/LLVM, ICC).
- A prototype implementation of the algorithm and its comprehensive evaluation on a set of diverse benchmarks, including both micro-benchmarks and a real-world `bzip2` program. Our prototype is capable of performing blackbox translation validation of C procedures with up to 100+ SLOC against their corresponding assembly implementations with up to 140+ instructions generated by an optimizing production compilers (such as GCC, Clang/LLVM, ICC) with complex loop and vectorizing transformations.

1.4 Outline

The remainder of this thesis is divided into chapters as follows. We also provide a summary in each case.

- In chapter 2, we formalize the execution semantics of the unoptimized IR \mathbb{C} obtained from translation of a C program and its compiled 32-bit x86 assembly \mathbb{A} , and present the definition of our refinement relation for encoding the notion of correct translation.

The execution semantics of a procedure \mathbb{C} of \mathbb{C} and corresponding assembly procedure \mathbb{A} are described in terms of translation of \mathbb{C} and \mathbb{A} to a labeled transition graph representation. The refinement relation is defined through existence of an annotation in \mathbb{A} such that its observable behaviors matches with \mathbb{C} .

- In chapter 3, we describe a cross-product or product program \mathbb{X} between the annotated assembly \mathbb{A} and unoptimized IR \mathbb{C} . We impose a set of requirements on \mathbb{X} so that the construction of \mathbb{X} implies the existence of a refinement relation between \mathbb{C} and \mathbb{A} .

We further describe *callers's virtual smallest* and *safety relaxed* semantics of \mathbb{A} and \mathbb{C} that help in realizing an efficient automatic construction algorithm for the product program between the modified \mathbb{A} and \mathbb{C} . We show that it is possible to construct a product program between the original \mathbb{A} and \mathbb{C} from a product program between \mathbb{A} and \mathbb{C} with modified semantics.

- In chapter 4, we describe our algorithm DYNAMO for the simultaneous automatic construction of the product program \mathbb{X} and annotations for \mathbb{A} .

DYNAMO builds both the annotated \mathbb{A} and product program \mathbb{X} incrementally. This incremental construction enables use of counterexample-guided heuristics described in prior works.

- In chapter 5, we describe the SMT encoding for the proof obligations generated by DYNAMO.

We describe an allocation state array based encoding, which uses quantifiers over SMT arrays, and a faster interval encoding, which uses SMT bitvectors.

- In chapter 6, we present an implementation of DYNAMO and its comprehensive evaluation over a set of varied benchmarks.

Our benchmarks include C procedures sourced from various sources including prior work on translation validation adapted suitably for our setting and the `bzip2` program

from SPEC2000 [20] with up to 100+ SLOC. The compiled 32-bit x86 assemblies for these C procedures include vectorized code with up to 140+ instructions.

- In chapter 7, we give concluding remarks and some directions for future work.

Chapter 2

Execution Semantics and Notion of Correct Translation

Our objective is to show that an x86 executable \mathbb{A} is a correct translation of a C program \mathbb{C} . Towards it, we first formalize the execution semantics of \mathbb{C} and \mathbb{A} in terms of translation to a transition graph representation. We then define correct translation in terms of refinement of behavior from \mathbb{C} to \mathbb{A} . Our novel contribution here is showing refinement through existence of an annotation of \mathbb{A} such that the annotated \mathbb{A} refines \mathbb{C} .

This chapter is organized as follows: in section 2.1, we present our unoptimized IR and assembly representations; in section 2.2, we introduce our transition graph representation and in section 2.3, we formalize the execution semantics of IR and assembly through translation to transition graph. We present and discuss our refinement definition in sections 2.4 to 2.7.

2.1 Intermediate Source and Assembly Representations

As our very first step, we translate the C program \mathbb{C} to an unoptimized intermediate representation (IR), which we also refer to as \mathbb{C} — this translation is syntactic and does not involve any analysis or optimization over the C program \mathbb{C} . Similarly, we disassemble the x86 executable to obtain assembly program \mathbb{A} . In the section, we

```

int printf(const char*, ...);

C0: int fib(int n, int m) {
C1:   int v[n+2];
C2:   v[0]=0; v[1]=1;
C3:   for(int i=2; i<=m; ++i)
C4:     v[i]=v[i-1]+v[i-2];
C5:   printf("fib(%d) = %d", m, v[m]);
C6:   return v[m];
C7: }

(a) C program with VLA.

I0: int fib(int* n, int* m):
I1:   i=alloc 1, int, 4;
I2:   v=alloc *n+2, int, 4;
I3:   v[0]=0; v[1]=1;
I4:   *i=2;
I5:   if(*i >_s *m) goto I9;
I6:     v[*i]=v[*i-1]+v[*i-2];
I7:     ++(*i);
I8:     goto I5;
I9:   pI9=alloc 1, char*, 4;
I10:  pI10=alloc 1, struct{int;int;}, 4;
I11:  *pI9=__S__; /* __S__ is the
                  address of the
                  format string */
I12:  *pI10=*m; *(pI10+4)=v[*m];
I13:  t=call int printf(pI9, pI10);
I14:  dealloc I10;
I15:  dealloc I9;
I16:  r=v[*m];
I17:  dealloc I2;
I18:  dealloc I1;
I19:  ret r;

A0: fib:
A1:   push ebp; ebp = esp;
A2:   push {edi, esi, ebx};
A3:   esp -= 12;
A31: vI1 = allocv 4, 4, I1;
A4:   eax = mem4[ebp+8]; ebx = mem4[ebp+12];
A5:   esp -= 0xFFFFFFFF0 & (4*(eax+2)+15));
A51: allocs esp, 4*(eax+2), 4, I2;
A6:   esi = ((esp+3)>>2)*4;
A7:   mem4[esi] = 0; mem4[esi+4] = 1;
A8:   if(ebx <_s 1) jmp A15;
A9:   edi = 0; edx = 1; eax = 2;
A10:  ecx = edx+edi;
A11:  edi = edx; edx = ecx;
A12:  mem4[esi+4*eax] = ecx;
A13:  ++eax;
A14:  if(eax <_s ebx) jmp A10;
A15:  edi = mem4[esi+4*ebx];
A16:  esp -= 4;
A17:  push {edi, ebx, __S__};
A171: allocs esp, 4, 4, I9;
A172: allocs esp+4, 8, 4, I10;
A18:  call int printf
      (<char*> esp,
       <struct{int; int;}> esp+4)
      {hp, cl, I9, I10};
A181: deallocs I10;
A182: deallocs I9;
A19:  eax = edi;
A191: deallocs I2;
A192: deallocv I1;
A20:  esp = ebp-12;
A21:  pop {ebx, esi, edi, ebp};
A22:  ret;

(c) (Abstracted) 32-bit x86 assembly code.

```

(b) (Abstracted) Unoptimized IR.

Figure 2.1: C program with variable-length array (VLA) and its lowerings to unoptimized IR and assembly. Subscript ‘_s’ denotes signed comparison. Red font (parts of) instructions in assembly are added by our algorithm.

describe the unoptimized intermediate source and assembly representations and the programming constructs we support through them. We discuss our logical model in the context of compilation to 32-bit x86 for the relative simplicity of the calling conventions in 32-bit mode.

We use the example shown in fig. 2.1 in our exposition. Figure 2.1 includes a C program, its unoptimized IR, and compiler generated assembly. The `fib` procedure in fig. 2.1a accepts two integers `n` and `m`, allocates a variable-length array (VLA) `v` of `n+2` elements, computes the first `m+1` Fibonacci numbers in `v`, calls `printf()`, and returns the m^{th} Fibonacci number. Notice that for a UB-free execution, both `n` and `m` must be non-negative and `m` must be less than `(n+2)`. The memory for local variables (`v` and `i`) and procedure-call arguments (for the call to `printf()`) is allocated dynamically through the `alloc` instruction in the IR program (fig. 2.1b). In the assembly program (fig. 2.1c), memory is allocated through instructions that manipulate the stackpointer register `esp`. We will continue to refer figs. 2.1b and 2.1c in remaining text.

2.1.1 Unoptimized IR

The unoptimized IR we use in our representation of the C program is mostly a subset of LLVM[48] — it supports all the primitive types (integer, float, code or PC labels) and the derived types (pointer, array, struct, procedure) of LLVM. Being unoptimized, our IR does not need to support LLVM’s fine-grained undefined behavior semantics enabled by `undef` and `poison` values, it instead treats all error conditions identically as undefined behavior (UB). Syntactic conversion of C to LLVM IR entails the usual conversion of types and operators. A global variable name `g` or a parameter name `y` appearing in a C procedure body is translated to the variable’s start address in IR, denoted `[lb.g]` and `[lb.y]` respectively¹. A local variable declaration or an invocation of the `alloca()` operator [2] is converted to LLVM’s `alloca[1]` instruction, and to distinguish the two, we henceforth refer to the latter as the “`alloc`” instruction. Unlike LLVM, our IR also supports a `dealloc` instruction that deallocates a variable at the end of its scope — we describe the semantics of both `alloc` and `dealloc` in section 2.1.3. Henceforth, we refer to our IR as LLVM_d, short for LLVM + `dealloc`.

We use a modified Clang[10], the de facto frontend for C in LLVM project, for our translation from C to LLVM_d. We use LLVM’s `stacksave` and `stackrestore` intrinsics generated by Clang to introduce an explicit `dealloc` instruction for each `alloc` instruction that corresponds to a variable allocation. These intrinsics are inserted by Clang roughly at scope boundaries and we use them as proxy for scope while inserting `dealloc` instructions. The allocations due to `alloca()` operator are deallocated at the

¹As we will also see later in section 2.2.3, `[lb.v]` denotes the *lower bound* of the memory addresses occupied by variable with name `v`.

end of the procedure and do not require scope tracking.

In fig. 2.1b, explicit `alloc` and `dealloc` instructions are inserted for locals `v` and `i` in fig. 2.1a. The `dealloc` instructions for these `alloc` instructions are inserted just before the end of procedure at I7 and I8.

Translation from C to LLVM_d for Procedure Definitions and Calls

Like LLVM, a procedure definition in LLVM_d can only return a scalar value — aggregate return value (of `struct` type) is passed in memory. Unlike LLVM, where a procedure takes parameters by value, LLVM_d takes all parameters through pointers of corresponding types, e.g., both `n` and `m` are passed through pointers of type `int` in fig. 2.1b. This makes the translation of a procedure call from C to LLVM_d slightly more verbose, as explicit instructions to (de)allocate memory for the arguments are required at the callsite. The LLVM_d `call` instruction takes the pointers returned by these allocations as operands.

Figure 2.2 shows the pseudo-code for translation of a C procedure-call expression to LLVM_d. The procedure-call expression is represented by $\rho(e_1, e_2, \dots, e_m)$ (shown at the top in fig. 2.2) where ρ is either the procedure name or a pointer to a procedure. The type of ρ is $(\tau_1, \dots, \tau_n) \rightarrow \gamma$ where τ_1, \dots, τ_n represents the parameters' types and γ represents the return type. A well-formed expression should have $m \geq n$. A code fragment with a shaded background represents a *templated* IR instruction with template slots (to be filled at runtime) marked by $\langle \! \! \langle \rangle \! \! \rangle$, e.g., `$p_r := \text{alloc } 1, \langle \gamma \rangle, \langle \text{ALIGNOF}(\gamma) \rangle$` ; represents an IR instruction where $\langle \gamma \rangle$ and $\langle \text{ALIGNOF}(\gamma) \rangle$, are instantiated for a concrete type γ — the instruction corresponds to allocation of an element of type $\langle \gamma \rangle$ and alignment $\langle \text{ALIGNOF}(\gamma) \rangle$.

LLVM_d instructions `alloc`, `store`, and `call` have similar syntactical structure as their LLVM counterparts: `alloc` takes the number of elements to be allocated, the LLVM_d type, and required alignment respectively as parameters (see section 2.1.3 for semantics); `store` accepts the type of store target, its alignment, the target value, and the target pointer respectively as its parameters; and the return type-parametric `call` instruction takes the callee label or address and the call arguments as parameters and returns a value if the return type is non-void.

$\text{GEN}_\tau(e)$ returns the LLVM_d variable holding value of expression e after casting it to type τ ; `promoted_type(e)` returns the promoted type of C expression e obtained after application of *default argument promotion* rules (see §6.5.2.2 of [21]); `mk_struct_x86_cc(...)`

```

                                 $\rho(e_1, e_2, \dots, e_m)$      $(\tau_1, \dots, \tau_n) \rightarrow \gamma$  is the type of  $\rho, m \geq n$ 

```

```

argsP := [];    // list of argument pointers
if is_aggregate_type( $\gamma$ ) {
    EMIT( $p_r := \text{alloc } 1, \langle \gamma \rangle, \langle \text{ALIGNOF}(\gamma) \rangle$ );
    argsP := argsP ·  $p_r$ ;    // add pointer to allocated region as first argument
}
for i in 1...n {    // non-variadic arguments
    EMIT( $p_i := \text{alloc } 1, \langle \tau_i \rangle, \langle \text{ALIGNOF}(\tau_i) \rangle$ );
    EMIT( $\text{store } \langle \tau_i \rangle, \langle \text{ALIGNOF}(\tau_i) \rangle, \langle \text{GEN}_{\tau_i}(e_i) \rangle, p_i$ );
    argsP := argsP ·  $p_i$ ;
}
if m > n {
     $\kappa_1, \dots, \kappa_i, \dots, \kappa_{m-n} := \text{promoted\_type}(e_{n+1}), \dots, \text{promoted\_type}(e_m)$ ;
     $\eta := \text{mk\_struct\_x86\_cc}(\kappa_1, \dots, \kappa_i, \dots, \kappa_{m-n})$ ;    // x86 calling conventions compatible type
    EMIT( $pvar := \text{alloc } 1, \langle \eta \rangle, \langle \text{ALIGNOF}(\eta) \rangle$ );
    argsP := argsP ·  $pvar$ ;
    EMIT( $p_v := pvar$ );
    for i in 1...(m - n) {
        EMIT( $\text{store } \langle \kappa_i \rangle, \langle \text{ALIGNOF}(\kappa_i) \rangle, \langle \text{GEN}_{\kappa_i}(e_i) \rangle, p_v$ );
        EMIT( $p_v := p_v + \langle \text{OFFSETOF}(\eta, i) \rangle$ );
    }
}
if  $\gamma = \text{void}$  {
    EMIT( $\text{call void } \langle \rho \rangle(\langle \text{argsP} \rangle)$ );
} else if is_aggregate_type( $\gamma$ ) {
    EMIT( $\text{call } \langle \gamma \rangle \langle \rho \rangle(\langle \text{argsP} \rangle)$ );
    EMIT( $\vec{\text{result}} := \langle \text{AGG2REG}(p_r) \rangle$ );    // distribute the populated aggregate into IR variables
} else {
    EMIT( $\text{result} := \text{call } \langle \gamma \rangle \langle \rho \rangle(\langle \text{argsP} \rangle)$ );
}
// deallocate the allocated arguments
for a in reverse(argsP) {
    EMIT( $\text{dealloc } \langle a \rangle$ );
}

```

Figure 2.2: Pseudo-code for translation of a C procedure-call expression to LLVM_d instructions. alloc, dealloc, call, va_start_ptr, etc. are LLVM_d instructions.

returns a C ‘struct’ type whose member fields’ alignment matches the calling conventions’ requirements of parameters for 32-bit x86; OFFSETOF(η, i) returns the offset (in bytes) of i^{th} member field in struct type η . AGG2REG(p) returns the list of values in aggregate pointed to by p .

For a procedure-call expression $\rho(e_1, e_2, \dots, e_m)$ with parameter types τ_1, \dots, τ_n , the translation proceeds as follows. For an aggregate return type ρ , an allocation for the return value is performed through `alloc` and the resulting pointer is saved as the first argument to the callee. For each non-variadic argument e_i ($1 \leq i \leq n$), the algorithm performs allocation according to type τ_i (of i^{th} parameter) through `alloc` and store e_i into the allocated region (through `store`). A single `struct` with layout respecting the calling conventions requirement is allocated for the variable-argument list and each variadic argument e_j ($n+1 \leq j \leq m$) is stored at appropriate offset inside the allocated region. The `call` instruction is passed the pointers to allocated regions as arguments. In the epilogue, each allocated region is deallocated through a `dealloc` instruction in reverse order of allocation.

An example of this translation is shown in fig. 2.1, where the call to `printf` at C5 in fig. 2.1a translates to instructions I9 to I15 in fig. 2.1b. The LLVM_d program performs two allocations in I9 and I10, one for the format string ("`fib(%d) = %d`") and another for the variable argument list (`m, v[m]`); the latter is represented as a single object of type "`struct {int;int;}`" containing two `ints`. The `call` instruction in I13 takes the pointers returned by the two allocations as arguments and stores the (unused) return value in pseudo-register `t`.

The memory allocation for procedure-call arguments follows the 32-bit x86 calling conventions [32], where stack space is used for passing arguments, making it a suitable semantic choice for this TV setting.

Translation of C's variadic macros and `va_list` type

$\text{va_start}(ap, last)$	$\text{va_end}(ap)$
$a := \text{va_start_ptr}$ $\text{store void}^*, 4, a, \langle ap \rangle$	$\text{store void}^*, 4, 0, \langle ap \rangle$
$\text{va_arg}(ap, \tau)$	$\text{va_copy}(aq, ap)$
$a := \text{load void}^*, 4, \langle ap \rangle$ $result := \text{load } \langle \tau \rangle, \langle \text{alignof}(\tau) \rangle, a$ $a' := a + \langle \text{roundup}_4(\text{sizeof}(\tau)) \rangle$ $\text{store void}^*, 4, a', \langle ap \rangle$	$a := \text{load void}^*, 4, \langle ap \rangle$ $\text{store void}^*, 4, a, \langle aq \rangle$

Figure 2.3: Translation of C's variadic macros to LLVM_d instructions. $\text{roundup}_4(a)$ returns the closest multiple of 4 greater than or equal to a .

C’s variadic macros `va_start`, `va_arg`, `va_end`, and `va_copy` are translated as shown in fig. 2.3. The translation is presented in the form of translation rules where templated IR instructions (as used in fig. 2.2) are shown below C syntax. Like LLVM (in case of 32-bit x86 `cdecl` calling conventions), we translate `va_list` to a pointer type whose object (denoted by ap in fig. 2.2) is initialized to the first address of the variable argument list, obtained through LLVM_d’s `va_start_ptr` instruction, in `va_start`. `va_arg` increments the address in passed `va_list` ap according to the passed type τ and `va_end` resets the address in ap to NULL or 0 address. `va_copy` simply copies the value in source ap into destination aq .

2.1.2 Assembly

Broadly, an assembly program \mathbb{A} consists of a code section with a sequence of assembly instructions, a data section with read-only and read-write global variables, and a symbol table that maps string symbols to memory addresses in code and data sections. Our translation validator checks that the address regions specified by the symbol table are well-aligned and non-overlapping, and uses it to relate a global variable (or procedure) in \mathbb{C} to its address (or implementation) in \mathbb{A} . For read-only symbols common in both \mathbb{C} and \mathbb{A} , the validator verifies that the memory contents are identical.

We assume that the OS guarantees the caller-side contract of the ABI calling conventions for the entry procedure, `main()`. For 32-bit x86, this means that at the start of program execution, the stackpointer is available in register `esp`, and the return address and input parameters (`argc, argv`) to `main()` are available in the stack region just above the stackpointer. For other procedure calls, the validator verifies the adherence to calling conventions at a callsite (in the caller) and assumes adherence at procedure entry (in the callee). Heap (de)allocation procedures like `malloc()` and `free()` are left uninterpreted, and so, the only compiler-visible way to allocate (and deallocate) memory in \mathbb{A} is through the decrement (and increment) of the stackpointer stored in register `esp`.

2.1.3 Allocation and Deallocation

Allocation and deallocation instructions appear only in \mathbb{C} and do not appear in \mathbb{A} . Let \mathbb{C} represent a procedure in program \mathbb{C} .

An LLVM_d instruction “ $p_{\mathbb{C}}^a: v := \text{alloc } n, \tau, \text{align}$ ” at PC $p_{\mathbb{C}}^a$ allocates a contigu-

ous region of local memory with space for n elements of type τ aligned by `align` and returns its start address in v . The PC p_C^a of an `alloc` instruction is also called an *allocation site*. We denote an allocation site by z where $z = p_C^a$. Let the set of allocation sites in C be Z such that $z \in Z$. During translation of the C program to $LLVM_d$, we distinguish between allocation sites due to the declaration of a local variable (or a procedure-call argument) and allocation sites due to `alloca()` — we use Z_l for the former and Z_a for the latter, so that $Z = Z_l \cup Z_a$.

The address of an allocated region returned by `alloc` is non-deterministic, but is subject to two *Well-Formedness (WF) constraints*:

1. The newly allocated memory region should be separate from all currently allocated memory regions, i.e., there should be *no overlap*.
2. The address of the newly allocated memory region should be aligned by `align`.

An error-free execution of `alloc` will satisfy these two well-formedness constraints.

An $LLVM_d$ instruction “ p_C^d : `dealloc z`” deallocates *all* local memory regions allocated due to the execution of (`alloc` instruction at) allocation site $z \in Z$. It is valid (i.e., not UB) to execute “`dealloc z`” when the `alloc` instruction at z was never executed — this can happen for an allocation site $z = za \in Z_a$ due to `alloca()`, where the `alloc` instruction is (conditionally) not executed but the `dealloc` instruction, inserted at the end of procedure C , is (unconditionally) executed.

In fig. 2.1b, the `alloc` instruction at I2, allocates space for `*n+2` integers (of type `int`) with alignment 4 and stores the region’s start address in v . The allocated region is identified by the allocation site I2 and deallocated at I17 using `dealloc I2`.

2.2 Transition Graph Representation

An $LLVM_d$ or assembly instruction may mutate the machine state, transfer control, perform I/O, or terminate the execution. We represent a C procedure, C in \mathbb{C} , as a transition graph, $C = (\mathcal{N}_C, \mathcal{E}_C)$, with a finite set of nodes $\mathcal{N}_C = \{n^s = n_1, n_2, \dots, n_m\}$, and a finite set of labeled directed edges \mathcal{E}_C . A unique node n^s represents the start node or entry point of C , and every other node n_j ($2 \leq j \leq m$) must be reachable from n^s . A node with no outgoing edges is a *terminating node*. A variable in C is identified by its scope-resolved unique name. The machine state σ_C of C consists of the set of

input parameters \vec{y}^2 , set of temporary variables \vec{t} , and an explicit array variable M_C denoting the current state of memory. We use i_N to denote a bitvector type of size $N > 0$. The type $T(M_C)$ of M_C is $i_{32} \rightarrow i_8$.

An assembly implementation of the C procedure C , identified through the symbol table in \mathbb{A} , is the assembly procedure A . Similarly to C , $A = (\mathcal{N}_A, \mathcal{E}_A)$ is also represented as a transition graph. The machine state σ_A of A consists of its hardware registers \overrightarrow{regs} and memory M_A and is disjoint from the machine state σ_C of C .

Let $P \in \{C, A\}$. In addition to the memory (data) state M_P , we also need to track the allocation state, i.e., the set of intervals of addresses that have been allocated by the procedure. We use α (potentially with a subscript) to denote a memory address of bitvector type i_{32} . Let $i = [\alpha_b, \alpha_e]$ represent an *address interval* starting at α_b and ending at α_e (both inclusive), such that $\alpha_b \leq_u \alpha_e$ (where \leq_u is unsigned comparison operator for bitvectors). Let $[\alpha]_w$ be a shorthand for the address interval $[\alpha, \alpha + w - 1_{i_{32}}]$, where $1_{i_{32}}$ ($n_{i_{32}}$) is the two's complement representation of integer 1 (n) using 32 bits.

2.2.1 Address Set

Let Σ (potentially with a sub- or superscript) represent a set of addresses, or an *address set*. An empty address set is represented by \emptyset , and an address set of contiguous addresses is an address interval i . Two address sets overlap, written $ov(\Sigma_1, \Sigma_2)$, iff $\Sigma_1 \cap \Sigma_2 \neq \emptyset$. Extended to $m > 2$ sets, $ov(\Sigma_1, \Sigma_2, \dots, \Sigma_m) \Leftrightarrow \exists_{1 \leq j_1 < j_2 \leq m} ov(\Sigma_{j_1}, \Sigma_{j_2})$. $|\Sigma|$ represents the number of distinct addresses in Σ . For a non-empty address set, $lb(\Sigma)$ and $ub(\Sigma)$ represent the smallest and largest address respectively in Σ such that $\Sigma \subseteq [lb(\Sigma), ub(\Sigma)]$. $comp(\Sigma)$ represents the *complement* of Σ , so that: $\forall_\alpha : (\alpha \in \Sigma) \Leftrightarrow (\alpha \notin comp(\Sigma))$.

2.2.2 Memory Regions

To support dynamic (de)allocation and memory related transformations (e.g., re-ordering, elimination, etc.), an execution model in a validator needs to individually track regions of memory belonging to each variable, heap, stack, etc. We next describe the memory regions tracked by our model.

1. Let G be the set of names of all global variables in \mathbb{C} . For each global variable $g \in G$, we track the memory region belonging to that variable. We use the name of a global

²We use the notation \vec{x} for representing a set.

- variable $g \in G$ as its *region identifier* to identify the region belonging to g in both \mathbb{C} and \mathbb{A} .
2. For a procedure \mathbb{C} , let Y be the set of names of formal parameters, including the variadic parameter, if present. We use the special name `vrdc` to identify the variadic parameter. The memory region belonging to a parameter $y \in Y$ is identified by y in both \mathbb{C} and \mathbb{A} .
 3. The memory region allocated by allocation site $z \in Z$ is identified by z in \mathbb{C} . In \mathbb{A} , our algorithm potentially annotates allocation instructions corresponding to an allocation site z in \mathbb{C} . Thus, the memory region allocated by these annotated instructions is also identified by z in \mathbb{A} .
 4. hp denotes the region belonging to the *program heap* (managed by the OS) in both \mathbb{C} and \mathbb{A} . Recall that we leave `malloc()` and `free()` uninterpreted so that hp does not grow or shrink as \mathbb{C} (\mathbb{A}) executes. This effectively models hp as a static region, even when `malloc()` and `free()` may implement dynamic growth and shrinking of its subregions that we do not track in our model.
 5. Local variables and actual arguments may be allocated in the *call chain* of a procedure (caller, caller's caller, and so on). The accessible subset (accessible to procedure \mathbb{C} ³) is coalesced into a single region denoted by cl or *callers' locals* in both \mathbb{C} and \mathbb{A} .
 6. In procedure \mathbb{A} , stack memory can be allocated and deallocated through stackpointer decrement and increment. The addresses belonging to the stack frame of \mathbb{A} (but not to a stack-allocated local variable $z \in Z$ or a parameter $y \in Y$) belong to the *stk* (stack) region in \mathbb{A} . The *stk* region is absent in \mathbb{C} .
 7. Separate from *stk*, we use cs (*callers' stack*) to identify the region that belongs to the stack space (but not to cl) of the call chain of procedure \mathbb{A} . cs is absent in \mathbb{C} .
 8. Program \mathbb{A} may use more global memory than \mathbb{C} , e.g., to store precomputed constants to implement vectorizing transformations. Let F be the set of names of all non-empty *assembly-only global variables* in \mathbb{A} . For each $f \in F$, its memory region in \mathbb{A} is identified by f .

³A local variable or actual argument v of procedure \mathbb{C}' in the call chain of procedure \mathbb{C} is accessible in procedure \mathbb{C} only if the address of v is accessible in \mathbb{C} , i.e., v is address-taken in \mathbb{C}' .

9. The region cv^4 denotes the inaccessible subset of local variables and actual arguments in the call chain of C . cv is present in both C and A , but inaccessible (i.e., cannot be read from or written to) in C and potentially accessible in A — we will elaborate on the accessibility aspect later when we discuss semantics of a memory access in A .
10. The region **free** denotes the free space, that does not belong to any of the aforementioned regions, in both C and A .

Let $R = G \cup F \cup Y \cup Z \cup \{hp, cl, cv, stk, cs, free\}$ represent all *region identifiers*; let $S = \{stk, cs\}$ denote the stack regions in A and $B = G \cup Y \cup Z \cup \{hp, cl\}$ ($B = R \setminus (F \cup S \cup \{cv, free\})$) denote the accessible regions in both C and A .

Let $G_r \subseteq G$ be the set of read-only global variables in C ; and, let $G_w = G \setminus G_r$ denote the set of read-write global variables. Let $F_r \subseteq F$ and $F_w = F \setminus F_r$ be defined analogously.

For each non-free region $r \in R \setminus \{free\}$, the machine state σ_P of a procedure P includes a unique variable Σ_P^r that tracks region r 's address set as P executes. If Σ_P^r is a contiguous non-empty interval, we also refer to it as i_P^r . For $r \in G \cup F \cup Y \cup \{hp, cl, cv, cs\}$ ($r \in R \setminus (Z \cup \{stk, free\})$), Σ_P^r remains constant throughout P 's execution. For $\vec{r} \subseteq R$, we define an expression $\Sigma_P^{\vec{r}} = \bigcup_{r \in \vec{r}} \Sigma_P^r$. Because C does not have a stack or an assembly-only global variable, $\Sigma_C^{F \cup S} = \emptyset$ holds throughout C 's execution. At any point in P 's execution, the free space can be computed as $\Sigma_P^{free} = \text{comp}(\Sigma_P^{B \cup F \cup S \cup \{cv\}})$. Notice that we do not use an explicit variable to track Σ_P^{free} .

For the IR and assembly procedures in fig. 2.1, $G = F = \emptyset$, $Y = \{n, m\}$, $Z = \{I1, I2, I9, I10\}$, and, therefore, $R = \{n, m, I1, I2, I9, I10, hp, cl, cv, stk, cs, free\}$.

2.2.3 Ghost Variables

Our validator introduces *ghost variables* in a procedure's execution semantics, i.e., variables that were not originally present in P . We use \boxed{x} to indicate that x is a ghost variable. For each region $r \in G \cup Y \cup Z$ (resp. $r \in F$), we introduce $\boxed{em.r}$, $\boxed{lb.r}$, and $\boxed{ub.r}$ in C (resp. A) to track the *emptiness* (whether the region is empty), *lower bound* (smallest address), and *upper bound* (largest address) of Σ_C^r (resp. Σ_A^r) respectively; for $r \in G \cup Y$ (resp. $r \in F$), $\boxed{sz.r}$ tracks the size of Σ_C^r (resp. Σ_A^r), and for $z \in Z$, $\boxed{lstSz.z}$ tracks the *size of last allocation* due to execution of allocation site z . Two ghost variables

⁴ cv stands for *callers' virtual*. The reason for tracking this region will become apparent when we discuss virtual allocation in section 2.6.

Σ_P^{rd} and Σ_P^{wr} track the set of addresses read and written by P respectively. Let \oplus be the set of all ghost variables.

2.2.4 Error Codes

Execution of \mathbf{C} or \mathbf{A} may terminate successfully, may never terminate, or may terminate with an error. We support two error codes to distinguish between two categories of errors: \mathcal{U} and \mathcal{W} .

- In \mathbf{C} : \mathcal{U} represents an occurrence of UB, and \mathcal{W} represents a violation of a well-formedness (WF) constraint that needs to be ensured either by the language or the OS (both external to the program itself).
- In \mathbf{A} : \mathcal{U} represents an occurrence of UB or a translation error, and \mathcal{W} represents occurrence of a condition that can be assumed to never occur, e.g., if the OS ensures that it never occurs.

In summary, for a procedure P , \mathcal{W} represents an error condition that P can *assume* to be absent (because the external environment ensures it), while \mathcal{U} represents an error that P must *ensure* to be absent. For \mathbf{C} , the programmer must ensure absence of UB; for \mathbf{A} , the compiler must ensure absence of translation error.

2.2.5 Outside world and observable trace

Let Ω_P be a state of the outside world (OS/hardware) for P that supplies external inputs whenever P reads from it and consumes external outputs generated by P . Ω_P is assumed to mutate arbitrarily but deterministically based on the values consumed or produced due to the I/O operations performed by P during execution.

Let T_P be a potentially infinite sequence of observable trace events generated by an execution of P . A trace event t , produced during execution of P , is concatenated to T_P , written $T_P := T_P \cdot t$. The trace events generated during an execution of P are as follows: (1) procedure-call invocation (`fcall`), (2) procedure return (`ret`), (3) procedure termination or halt (`exit`), (4) local allocation (`allocBegin` and `allocEnd`), (5) local deallocation (`dealloc`), and (6) a distinct silent trace event (\perp), indicating execution of an instruction that does not otherwise produce an observable event. A trace event may be associated with a set of values, e.g., a procedure-call invocation is associated with callee-observable values and a procedure return is associated with caller-observable

values — we will describe the exact values in each case later when we talk about graph instructions that produce these events (section 2.2.7). Procedure termination includes both error (indicated by an error code) and error-free termination (indicated by `exit`), each of which is identified distinctly.

2.2.6 Expressions

Let variable v and variables \vec{v} or \vec{x} be drawn from $\text{Vars} = \vec{r} \cup \overrightarrow{\text{regs}} \cup \{\oplus\} \cup \{M_P, \Sigma'_P\}$ for all $P \in \{\mathbf{C}, \mathbf{A}\}$ and for all $r \in R \setminus \text{free}$. Let $e(\vec{x})$ be an expression over \vec{x} , and $E(\vec{x})$ be a list of expressions over \vec{x} . An expression $e(\vec{x})$ is a well-formed combination of constants, variables \vec{x} , and arithmetic, logical, relational, memory access (read and write), and address set operators. For memory reads and writes, `select` (`sel` for short) and `store` (`st` for short) operations are used to access and modify M_P at a given address α . Further, the `sel` and `st` operators are associated with a `sz` parameter: `selsz(arr, α)` returns a little-endian concatenation of `sz` bytes starting at α in the array `arr`. Similarly, `stsz(arr, α , data)` returns a new array that has contents identical to `arr` except for the `sz` bytes starting at α that have been replaced by `data` in little-endian format. To encode reads/writes to a region of memory, we define projection and update operations.

Definition 2.2.1 ($\pi_\Sigma(M_P)$). $\pi_\Sigma(M_P)$ denotes the **projection** of M_P on addresses in Σ , i.e., if $M'_P = \pi_\Sigma(M_P)$, then $\forall \alpha \in \Sigma : \text{sel}_1(M'_P, \alpha) = \text{sel}_1(M_P, \alpha)$ and $\forall \alpha \notin \Sigma : \text{sel}_1(M'_P, \alpha) = 0$. The sentinel value 0 is used for the addresses outside Σ .

We use $M_{P_1} =_\Sigma M_{P_2}$ as shorthand for $(\pi_\Sigma(M_{P_1}) = \pi_\Sigma(M_{P_2}))$, for $P_1, P_2 \in \{\mathbf{C}, \mathbf{A}\}$.

Definition 2.2.2 ($\text{upd}_\Sigma(M_P, M)$). $\text{upd}_\Sigma(M_P, M)$ denotes the **update** of M_P on addresses in Σ using the values in M . If $M'_P = \text{upd}_\Sigma(M_P, M)$, then $M'_P =_\Sigma M$ and $M'_P =_{\text{comp}(\Sigma)} M_P$ hold.

2.2.7 Graph Instructions

Each labeled directed edge $e_P \in \mathcal{E}_P$ is labeled with one of the following *graph instructions*:

1. A *simultaneous assignment* of the form $\vec{v} := E(\vec{x})$: Because variables \vec{v} and \vec{x} may include M_P , an assignment suffices for encoding memory loads and stores. Similarly,

because the variables may be drawn from Σ_P^z (for an allocation site z), an assignment is also used to encode the allocation of an interval i_{new} through $\Sigma_P^z := \Sigma_P^z \cup i_{\text{new}}$ and the deallocation of all addresses allocated due to z through $\Sigma_P^z := \emptyset$. Stack allocation and deallocation in \mathbf{A} can be similarly represented as $\Sigma_{\mathbf{A}}^{\text{stk}} := \Sigma_{\mathbf{A}}^{\text{stk}} \cup i_{\text{new}}$ and $\Sigma_{\mathbf{A}}^{\text{stk}} := \Sigma_{\mathbf{A}}^{\text{stk}} \setminus i_{\text{new}}$ respectively.

2. A *guard* instruction of the form $e(\vec{x})?$: Instruction $e(\vec{x})?$ indicates that when execution reaches its head, the edge is taken iff its *edge condition* $e(\vec{x})$ evaluates to **true**. For every other instruction, the edge is always taken upon reaching its head, i.e., its edge condition is **true**.

For a *non-terminating* node $n_P \in \mathcal{N}_P$ (i.e., n_P has at least one outgoing edge ⁵), the guards of all edges departing from n_P must be mutually exclusive, and their disjunction must evaluate to *true*.

3. A type-parametric *choose* instruction $\theta(\vec{\tau})$: Instruction $\vec{v} := \theta(\vec{\tau})$ non-deterministically chooses values of types $\vec{\tau}$ and assigns them to variables \vec{v} , e.g., a memory with non-deterministic contents is obtained by using $\theta(i_{32} \rightarrow i_8)$ so that $M_P := \text{upd}_{\Sigma}(M_P, \theta(i_{32} \rightarrow i_8))$ updates M_P at addresses in Σ with non-deterministically chosen data values.
4. A *read* (**rd**) or *write* (**wr**) I/O instruction: A read instruction $\vec{v} := \text{rd}(\vec{\tau})$ reads values of types $\vec{\tau}$ from the outside world into variables \vec{v} , e.g., an address set is read using $\Sigma := \text{rd}(2^{i_{32}})$ where $2^{i_{32}}$ represents the type of address set Σ .

A write instruction $\text{wr}(V(E(\vec{x})))$ writes the value constructed by value constructor V using $E(\vec{x})$ to the outside world. A value constructor $V(\dots)$ is defined for each type of observable trace event.

- For a procedure-call event, $\text{fcall}(\rho, \vec{v}, \vec{\tau}, M)$ represents a value constructed for a procedure call to callee with name (or address) ρ , the actual arguments \vec{v} , callee-observable regions $\vec{\tau}$, and memory M .
- For a procedure-return event, $\text{ret}(E(\vec{x}))$ represents a value constructed during procedure return that captures observable values computed through $E(\vec{x})$.
- For local allocation and deallocation events, $\text{allocBegin}(z, w, a)$, $\text{allocEnd}(z, i, M)$, and $\text{dealloc}(z)$ represent the values constructed for allocation (allocBegin and allocEnd) and deallocation (dealloc) due to allocation site z with the associated observables, size of allocation w , alignment a , allocated interval i , and memory M .

⁵Recall that a terminating node has no outgoing edges.

A read or write instruction mutates outside world Ω_P arbitrarily based on the read and written values. Further, the data items read or written are appended to the observable trace T_P . Let $\text{read}_{\vec{r}}(\Omega_P)$ be an uninterpreted function that reads values of types \vec{r} from Ω_P ; and $\text{io}(\Omega_P, \text{rw}, E(\vec{x}))$ be an uninterpreted function that returns an updated state of Ω_P after an I/O operation of type $\text{rw} \in \{\text{r}, \text{w}\}$ (read or write) with values $E(\vec{x})$. Thus, in its explicit syntax, $\vec{v} := \text{rd}(\vec{r})$ translates to a sequence of instructions: $\vec{v} := \text{read}_{\vec{r}}(\Omega_P); \Omega_P := \text{io}(\Omega_P, \text{r}, \vec{v}); T_P := T_P \cdot \vec{v}$, where \cdot is the trace concatenation operator. Similarly, $\text{wr}(V(E(\vec{x})))$ translates to: $\Omega_P := \text{io}(\Omega_P, \text{w}, V(E(\vec{x}))); T_P := T_P \cdot V(E(\vec{x}))$. Henceforth, we only use the implicit syntax for brevity.

5. An error-free and error-indicating *halt* instruction that terminates execution. $\text{halt}(\emptyset)$ indicates termination without error and $\text{halt}(\varkappa)$ indicates termination with error code $\varkappa \in \{\mathcal{U}, \mathcal{W}\}$. Upon termination without error, a special *exit* event is appended to trace T_P ; upon termination with error, the error code is appended to T_P .

The destination of an edge with a *halt* instruction is a terminating node. We create a unique terminating node for an error-free exit. We also create a unique terminating node for each error code, also called an *error node*. An edge terminating at an error node is called an *error edge*. \mathcal{U}_P and \mathcal{W}_P represent error nodes in P for errors \mathcal{U} and \mathcal{W} respectively. Execution transfers to an error node upon encountering the corresponding error. Let $\mathcal{N}_P^{\text{OK}} = \mathcal{N}_P \setminus \{\mathcal{U}_P, \mathcal{W}_P\}$ be the set of error-free nodes in P .

In addition to the observable trace events generated by *rd*, *wr*, and *halt* instructions, the execution of every instruction in P also appends an observable *silent trace event*, denoted \perp , to T_P . Silent trace events count the number of executed instructions as a proxy for observing the passage of time.

2.3 Translations of C and A to their Graph Representations

Figures 2.4 to 2.7 (and figs. 2.8, 2.10 and 2.11 later) present the key translation rules from LLVM_d and (abstracted) assembly instructions to graph instructions. Each rule is composed of three parts separated by a horizontal line segment: on the left is the name of the rule, above the line segment is the LLVM_d /assembly instruction, and below the line segment is the graph instructions listing. For example, the top left corner of fig. 2.4 shows the parametric (OP) rule which gives the translation of an operation using

arithmetic/logical/relational operator op in LLVM_d to corresponding graph instructions.

We describe the operators and predicates used in the rules in table 2.1. We use C-like constructs in graph instructions as syntactic sugar for brevity, e.g. ‘;’ is used for sequencing, ‘?:’ is used for conditional assignment, and **if**, **else**, and **for** are used for control flow transfer. We highlight the read and write I/O instructions with a shaded background and use **bold**, **colored** fonts for error-indicating **halt** instructions. We use “macros” **IF** and **ELSE** to choose translations based on a boolean condition on the input syntax.

2.3.1 Translation of C

Figures 2.4 and 2.5 shows the rules for translating LLVM_d instructions to graph instructions. We discuss each in the following paragraphs.

The parametric rule (OP) gives the translation for application of an arithmetic/logical/relational operator op over arguments \vec{x} (fig. 2.4). An application of op may trigger undefined behavior (UB) for certain inputs, as abstracted through the $\text{UB}_C(\text{op}, \vec{x})$ operation. While there are many undefined behaviors in the C standard, we model only the following that we have seen getting exploited for optimization by the compiler:

1. *Logical or arithmetic shift operation*: The second operand should be bounded by a limit which is determined by the bit width of the first operand. This is required when a shift operation in C is translated to an x86 shift opcode in A.
2. *Address computation (getelementptr inbounds opcode in LLVM IR)*: No overflow and underflow in the intermediate and final computations. An optimizing compiler may assume this for conversion of inequality relations to disequality relations.
3. *Division operation*: The denominator operand should be non-zero. This is required for showing an exception-free execution of the corresponding division operation in A.

We will describe the $\beta(v) := \dots$ part of the translation shortly.

The (LOAD_C) and (STORE_C) rules (fig. 2.4) show the translations for **load** and **store** instructions respectively. A UB-free execution of **load** and **store** requires the dereferenced pointer p to satisfy memory access safety constraints specified through $\text{accessIsSafeC}()$ predicate (defined in table 2.1) — a *safe* memory access is recorded in ghost variable Σ_C^{rd} for **load** and Σ_C^{wr} for **store**. An access through pointer p is safe iff p is non-NULL

Table 2.1: Definitions of operators and predicates used in translations in figs. 2.4 to 2.8 and 2.10 to 2.12

Operator	Definition
$\text{sz}(\tau)$	Returns the size (in bytes) of type τ . For example, $\text{sz}(\text{i}_{32}) = 4$ and $\text{sz}(\text{i}_8^*) = 4$.
$\text{T}(a)$	Returns the type τ of a where a may be a global variable, a parameter, or a register. For example, $\text{T}(\text{eax}) = \text{i}_{32}$.
$\Delta_\tau(\text{eax}, \text{edx})$	A type-parametric operator which derives the return value of an assembly procedure with return type τ from input registers eax and edx using the calling conventions, e.g., $\Delta_{\text{i}_8}(\text{eax}, \text{edx}) = \text{extract}_{7,0}(\text{eax})$, $\Delta_{\text{i}_{32}}(\text{eax}, \text{edx}) = \text{eax}$, $\Delta_{\text{i}_{64}}(\text{eax}, \text{edx}) = \text{concat}(\text{edx}, \text{eax})$, where $\text{extract}_{h,l}(a)$ extracts bits h down to l from a and $\text{concat}(a, b)$ returns the bitvector concatenation of a and b where b takes the less significant position.
$\nabla_\tau(v)$	Inverse of $\Delta_\tau(\text{eax}, \text{edx})$. Distributes the packed bitvector v of type τ into two bitvectors of 32 bit-width each, setting the bits not covered by v to some non-deterministic value.
$\text{ROM}_P^r(i)$	Returns a memory array containing the contents of read-only global variable named r in P . The contents are mapped at the addresses in the provided interval i .
$\text{addrSets}_F()$	Returns the address sets of the assembly-only global variables F using the symbol table in the executable \mathbb{A} .

Predicate	Definition
$\text{aligned}_n(a)$	Bitvector a is n bytes aligned. Equivalent to: $a \% n = 0$, where $\%$ is remainder operator.
$\text{isAlignedIntrvl}_a(p, w)$	A w -sized sequence of addresses starting at p is aligned by a and does not wraparound. Equivalent to: $\text{aligned}_a(p) \wedge (p \leq_u p + w - 1_{\text{i}_{32}})$.
$\text{accessIsSafeC}_{\tau,a}(p, \Sigma)$	Equivalent to: $\text{isAlignedIntrvl}_a(p, \text{sz}(\tau)) \wedge ([p]_{\text{sz}(\tau)} \subseteq \Sigma)$.
$\text{addrSetsAreWF}(\Sigma_P^{hp}, \Sigma_P^{cl}, \Sigma_P^{cv}, \dots, i_P^g, \dots, \Sigma_P^f, \dots, i_P^y, \dots, \Sigma_P^{\text{vrdc}}, \dots, i_P^y, \dots, \Sigma_P^{\text{vrdc}})$	The address sets passed as parameter are well-formed with respect to C semantics. Equivalent to: $(0_{\text{i}_{32}} \notin \Sigma_P^{GUFUYU\{hp,cl,cv\}}) \wedge \neg \text{ov}(\Sigma_P^{hp}, \Sigma_P^{cl}, \dots, i_P^g, \dots, \Sigma_P^f, \dots, i_P^y, \dots, \Sigma_P^{\text{vrdc}}) \wedge \neg \text{ov}(\Sigma_P^{GUYU\{hp,cl\}}, \Sigma_P^{cv}) \wedge (\Sigma_P^{\text{vrdc}} \neq \emptyset \Rightarrow \text{isInterval}(\Sigma_P^{\text{vrdc}})) \wedge \forall r \in GU(Y \setminus \{\text{vrdc}\})_{UF} : (i_P^r = \text{sz}(\text{T}(r)) \wedge \text{aligned}_{\text{alignmnt}(r)}(\text{lb}(i_P^r)))$, where $\text{isInterval}(\Sigma_P^{\text{vrdc}})$ holds iff the address set Σ_P^{vrdc} is an interval, $\text{alignmnt}(r)$ returns the alignment of variable r .
$\text{intrvlInSet}(\alpha_b, \alpha_e, \Sigma)$	The pair (α_b, α_e) forms a valid interval inside the address set Σ . Equivalent to: $(\alpha_b \neq 0_{\text{i}_{32}}) \wedge (\alpha_b \leq_u \alpha_e) \wedge ([\alpha_b, \alpha_e] \subseteq \Sigma)$
$\text{intrvlInSet}_a(\alpha_b, \alpha_e, \Sigma)$	Equivalent to: $\text{aligned}_a(\alpha_b) \wedge \text{intrvlInSet}(\alpha_b, \alpha_e, \Sigma)$
$\text{obeyCC}(e_{\text{esp}}, \vec{\tau}, \vec{x})$	Pointers \vec{x} match the expected addresses of arguments for a procedure call in assembly. Based on the calling conventions, obeyCC uses the value of the current stackpointer (e_{esp}) and parameter types ($\vec{\tau}$) to obtain the expected addresses of the arguments. For example, $\text{obeyCC}(\text{esp}, (\text{i}_8, \text{i}_{32}), (\text{esp}, \text{esp} + 4_{\text{i}_{32}}))$ holds.
$\text{overflow}_{mul}(a, b)$	Signed multiplication of bitvectors a, b overflows. E.g., $\text{overflow}_{mul}(2147483647_{\text{i}_{32}}, 2_{\text{i}_{32}})$ holds.
$\text{stkIsWF}(\text{esp}, \text{stk}_e, \text{cs}_e, \vec{\tau}, \Sigma_A^{hp}, \Sigma_A^{cl}, \Sigma_A^{GUF}, \dots, i_A^y, \dots, \Sigma_A^{\text{vrdc}})$	The pairs $(\text{esp}, \text{stk}_e)$, $(\text{stk}_e, \text{cs}_e)$ represent well-formed intervals for initial stk region and initial cs region with respect to parameter types $\vec{\tau}$ and other (input) address sets in \mathbb{A} . Equivalent to: $\text{aligned}_{16}(\text{esp} + 4_{\text{i}_{32}}) \wedge (\text{esp} \leq_u \text{esp} + 4_{\text{i}_{32}}) \wedge \neg \text{ov}([\text{esp}]_{4_{\text{i}_{32}}}, \Sigma_A^{GUFUYU\{hp,cl\}}) \wedge \text{obeyCC}(\text{esp} + 4_{\text{i}_{32}}, \vec{\tau}, \dots, \text{lb}(i_A^y), \dots) \wedge (\text{stk}_e <_u \text{cs}_e) \wedge \neg \text{ov}([\text{stk}_e + 1_{\text{i}_{32}}, \text{cs}_e], \Sigma_A^{GUFU\{hp\}}) \wedge \Sigma_A^{cl} \subseteq [\text{stk}_e + 1_{\text{i}_{32}}, \text{cs}_e]$
$\text{UB}_P(\text{op}, \vec{x})$	Application of operation op of procedure P on arguments \vec{x} triggers UB. E.g., $\text{UB}_C(\text{udiv}, (1_{\text{i}_{32}}, 0_{\text{i}_{32}}))$ holds.

$$\begin{array}{c}
\text{(OP)} \frac{p_C^j : v := \text{op}(\vec{x})}{\text{if } (\text{UB}_C(\text{op}, \vec{x})) \text{ halt}(\mathcal{U}); \\ v := \text{op}(\vec{x}); \\ \dots, x, \dots := \vec{x}; \quad \beta(v) := \beta^{\text{op}}(\dots, \beta(x), \dots);} \quad \text{(ASSIGNCONST)} \frac{p_C^j : v := c}{v := c; \\ \beta(v) := \emptyset;} \\
\\
\text{(LOAD}_C) \frac{p_C^j : v := \text{load } \tau, a, p}{\text{if } (\neg \text{accessIsSafe}_{C, \tau, a}(p, \Sigma_C^{\beta(p)})) \\ \text{halt}(\mathcal{U}); \\ v := \text{sel}_{\text{sz}(\tau)}(M_C, p); \\ \beta(v) := \beta_M(\beta(p)); \\ \Sigma_C^{\text{rd}} := \Sigma_C^{\text{rd}} \cup [p]_{\text{sz}(\tau)};} \quad \text{(STORE}_C) \frac{p_C^j : \text{store } \tau, a, v, p}{\text{if } (\neg \text{accessIsSafe}_{C, \tau, a}(p, \Sigma_C^{\beta(p) \setminus G_r})) \\ \text{halt}(\mathcal{U}); \\ M_C := \text{st}_{\text{sz}(\tau)}(M_C, p, v); \\ \beta_M(\beta(p)) := \beta_M(\beta(p)) \cup \beta(v); \\ \Sigma_C^{\text{wr}} := \Sigma_C^{\text{wr}} \cup [p]_{\text{sz}(\tau)};} \\
\\
\text{(VASTARTPTR)} \frac{p_C^j : p := \text{va_start_ptr}}{\text{if } (\Sigma_C^{\text{vrdc}} = \emptyset) \{ \\ p := 0_{\text{i32}}; \quad \beta(p) := \emptyset; \\ \} \text{ else } \{ \\ p := \text{lb.vrdc}; \quad \beta(p) := \{\text{vrdc}\}; \\ \}}
\end{array}$$

Figure 2.4: Translation rules for converting LLVM_d instructions to graph instructions. *op* represents an arithmetic, logical, or relational operator. *c* represents a constant.

($\neq 0_{\text{i32}}$ in our modeling⁶), aligned by the required alignment *a*, and have its access interval belong to the regions which *p* may *point to* or *p* may be *based on* (§6.5.6p8 of the C11 standard [21]).

To identify the regions a pointer *p* may be based on, we define two maps:

- (1) $\beta : \text{Vars} \rightarrow 2^R$ that tracks the set of regions a variable (e.g., a pointer) may be based on, so that for a variable $x \in \text{Vars}$, $\beta(x)$ returns the set of regions *x* may point to.
- (2) $\beta_M : R \rightarrow 2^R$ that tracks the set of regions that pointers in a memory region may be based on, so that for a region $r \in R$, $\beta_M(r)$ returns the set of regions that some (pointer) value stored in $\pi_{\Sigma_C^r}(M_C)$ may point to.

For convenience, we extend β and β_M to also take as input a set of variables and a set of regions respectively so that $\beta(\vec{x})$ is equivalent to $\bigcup_{x \in \vec{x}} \beta(x)$, and $\beta_M(\vec{r})$ is equivalent to $\bigcup_{r \in \vec{r}} \beta_M(r)$. Similar extension is used in assignment to β_M so that $\beta_M(\vec{r}_1) := \vec{r}_2$ is equivalent to ‘for r_1 in \vec{r}_1 { $\beta_M(r_1) := \vec{r}_2$; }’.

⁶The $\text{accessIsSafe}_{C, \tau, a}(p, \Sigma)$ definition in table 2.1 does not include the $\neq 0_{\text{i32}}$ clause because it assumes that $0_{\text{i32}} \notin \Sigma$ so that a $[p]_{\text{sz}(\tau)} \subseteq \Sigma$ check implies $p \neq 0_{\text{i32}}$. $0_{\text{i32}} \notin \Sigma_C^r$ for an allocated region *r* is an invariant in *C*.

The initialization and update of β and β_M due to each LLVM_d instruction can be seen in figs. 2.4 and 2.5. In (OP), for an operation op , $\beta^{\text{op}} : (2^R \times 2^R \dots \times 2^R) \rightarrow 2^R$ represents the over-approximate abstract transfer function for $v := \text{op}(\vec{x})$, that takes as input $\beta(x_1), \beta(x_2), \dots, \beta(x_m)$ for $\vec{x} = x_1, x_2, \dots, x_m$ and returns $\beta(v)$. We use:

- $\beta^{\text{op}}(\vec{r}) = \vec{r}$, if op is identity, bitwise complement and unary negation.
- $\beta^{\text{op}}(\vec{r}_1, \dots, \vec{r}_m) = \bigcup_{1 \leq j \leq m} \vec{r}_j$, if op is bitvector addition, subtraction, shift, bitwise-
{and,or}, extraction, or concatenation.
- $\beta^{\text{op}}(\vec{r}_1, \dots, \vec{r}_m) = \emptyset$, if op is bitvector multiplication, division, logical, relational, or any other remaining operator.

The rule (ASSIGNCONST) for a constant assignment to variable v sets $\beta(v)$ to empty set making it impossible to fabricate pointers from integer literals.

(VASTARTPTR) (fig. 2.4) gives the translation rule for the `va_start_ptr` instruction of LLVM_d — recall that the `va_start_ptr` instruction is used during translation of the C variadic macro `va_start()` to LLVM_d (fig. 2.3). The rule sets the assigned variable p to the first address of the variadic parameter region (obtained through `lb.vrdc`) if the address set $\Sigma_{\mathbb{C}}^{\text{vrdc}}$ of the variadic parameter region is non-empty, otherwise $0_{i_{32}}$ (NULL in our representation) is used. Because `lb.vrdc` is a pointer inside `vrdc` region, $\beta(p)$ is set to singleton `{vrdc}`.

The (ENTRY $_{\mathbb{C}}$) rule in fig. 2.5 presents the initialization performed at the entry of procedure \mathbb{C} . The allocation state, address set $\Sigma'_{\mathbb{C}}$ of each region $r \in R \setminus \{\text{free}\}$, and memory state, $M_{\mathbb{C}}$, of \mathbb{C} are initialized using reads from the outside world $\Omega_{\mathbb{C}}$ — the contents of read-only global regions (G_r) are initialized separately using their predefined values (through $\text{ROM}_{\mathbb{C}}^g(i_{\mathbb{C}}^g)$ defined in table 2.1). The read address sets are checked for well-formedness with respect to C semantics through `addrSetsAreWF()` (defined in table 2.1), or else error \mathcal{W} is triggered; well-formedness in this context requires that the address sets do not contain the NULL pointer ($0_{i_{32}}$), a global variable/argument address set is an interval, and the address sets do not overlap. The rule concludes with initialization of the ghost variables associated with the regions and the β, β_M maps used for tracking may based on information. Notice that the region `cv` is not included in the set of reachable regions through β and β_M making it unreachable throughout \mathbb{C} 's execution (this is ensured during a fresh allocation through an `alloc` instruction (ALLOC) as well).

(RETV) and (RET $_{\mathbb{C}}$) (fig. 2.5) present translations for procedure-return instructions

$$\begin{array}{c}
\text{(ENTRY}_C) \frac{p_C^j : \text{def } C(\vec{\tau})}{\begin{array}{l}
\Sigma_C^{hp}, \Sigma_C^{cl}, \Sigma_C^{cv}, \dots, i_C^g, \dots, i_C^y, \dots, \Sigma_C^{vrdc} := \text{rd}(2^{i_{32}}, 2^{i_{32}}, \dots, 2^{i_{32}}); \\
\Sigma_C^{stk}, \Sigma_C^{cs}, \dots, \Sigma_C^f, \dots, \Sigma_C^z, \dots, \Sigma_C^{rd}, \Sigma_C^{wr} := \emptyset, \emptyset, \dots, \emptyset; \\
\text{if } (\neg \text{addrSetsAreWF}(\Sigma_C^{hp}, \Sigma_C^{cl}, \Sigma_C^{cv}, \dots, i_C^g, \dots, \Sigma_C^f, \dots, i_C^y, \dots, \Sigma_C^{vrdc})) \\
\quad \text{halt}(\mathcal{W}); \\
M_C := \theta(i_{32} \rightarrow i_8); \quad M_C := \text{upd}_{\Sigma_C^{B \setminus G_r}}(M_C, \text{rd}(i_{32} \rightarrow i_8)); \\
\text{for } g \text{ in } G_r \{ M_C := \text{upd}_{i_C^g}(M_C, \text{ROM}_C^g(i_C^g)); \} \\
\text{for } r \text{ in } G \cup Y \{ \\
\quad \text{sz}.r, \text{em}.r := |\Sigma_C^r|, (|\Sigma_C^r| = 0_{i_{32}}); \\
\quad \text{if } (\neg \text{em}.r) \{ \text{lb}.r, \text{ub}.r := \text{lb}(\Sigma_C^r), \text{ub}(\Sigma_C^r); \} \\
\quad \beta(\text{lb}.r) := \{r\}; \\
\quad \} \\
\text{for } r \text{ in } G \cup Y \cup \{hp, cl\} \{ \beta_M(r) := G \cup \{hp, cl\}; \} \\
\text{for } z \text{ in } Z \{ \text{em}.z := \text{true}; \quad \beta_M(z) := \emptyset; \}
\end{array}}
\end{array}$$

$$\begin{array}{c}
\text{(RETV)} \frac{p_C^j : \text{ret void}}{\begin{array}{l}
\text{wr}(\text{ret}(\pi_{\Sigma_C^B}(M_C))); \\
\text{halt}(\emptyset);
\end{array}}
\end{array}$$

$$\begin{array}{c}
\text{(RETC)} \frac{p_C^j : \text{ret } v}{\begin{array}{l}
\text{wr}(\text{ret}(v, \pi_{\Sigma_C^B}(M_C))); \\
\text{halt}(\emptyset);
\end{array}}
\end{array}$$

$$\begin{array}{c}
\text{(ALLOC)} \frac{z : v := \text{alloc } n, \tau, a}{\begin{array}{l}
\text{IF}\{z \in Z_l\} \{ \\
\quad \text{if } (n \leq_s 0_{i_{32}} \vee \text{overflow}_{mul}(n, \text{sz}(\tau))) \\
\quad \quad \text{halt}(\mathcal{W}); \\
\quad \} \text{ ELSE } \{ \\
\quad \quad \text{if } (n = 0_{i_{32}}) \text{ halt}(\mathcal{W}); \\
\quad \quad \} \\
\text{wr}(\text{allocBegin}(z, n * \text{sz}(\tau), a)); \\
\alpha_b := \theta(i_{32}); \quad \alpha_e := \alpha_b + n * \text{sz}(\tau) - 1_{i_{32}}; \\
\text{if } (\neg \text{intrvlInSet}_d(\alpha_b, \alpha_e, \Sigma_C^{\text{free}})) \\
\quad \text{halt}(\mathcal{W}); \\
\Sigma_C^z, \quad \Sigma_C^z \cup [\alpha_b, \alpha_e], \\
M_C, \quad \text{upd}_{[\alpha_b, \alpha_e]}(M_C, \theta(i_{32} \rightarrow i_8)), \\
\text{lb}.z, \quad \text{em}.z? \alpha_b : \min(\text{lb}.z, \alpha_b), \\
\text{ub}.z, \quad \text{em}.z? \alpha_e : \max(\text{ub}.z, \alpha_e), \\
\text{em}.z, \quad \text{false}, \\
\text{lstSz}.z, \quad n * \text{sz}(\tau); \\
v := \alpha_b; \quad \beta(v) := \{z\}; \\
\text{wr}(\text{allocEnd}(z, [\alpha_b, \alpha_e], \pi_{[\alpha_b, \alpha_e]}(M_C)));
\end{array}}
\end{array}$$

$$\begin{array}{c}
\text{(DEALLOC)} \frac{p_C^j : \text{dealloc } z}{\begin{array}{l}
\Sigma_C^z := \emptyset, \\
\text{em}.z := \text{true}; \\
\text{wr}(\text{dealloc}(z));
\end{array}}
\end{array}$$

$$\begin{array}{c}
\text{(CALLV)} \frac{p_C^j : \text{call void } \rho(\vec{\tau} \vec{x})}{\begin{array}{l}
\beta^* := \beta_M^* \left(\bigcup_{x \in \vec{x}} \beta(x) \cup G \cup \{hp\} \right); \\
\text{wr}(\text{fcall}(\rho, \vec{x}, \beta^*, \pi_{\Sigma_C^{\beta^*}}(M_C))); \\
M_C := \text{upd}_{\Sigma_C^{\beta^* \setminus G_r}}(M_C, \text{rd}(i_{32} \rightarrow i_8)); \\
\beta_M(\beta^* \setminus G_r) := \beta^*;
\end{array}}
\end{array}$$

$$\begin{array}{c}
\text{(CALLC)} \frac{p_C^j : v := \text{call } \gamma \rho(\vec{\tau} \vec{x}) \quad \gamma \neq \text{void}}{\begin{array}{l}
\beta^* := \beta_M^* \left(\bigcup_{x \in \vec{x}} \beta(x) \cup G \cup \{hp\} \right); \\
\text{wr}(\text{fcall}(\rho, \vec{x}, \beta^*, \pi_{\Sigma_C^{\beta^*}}(M_C))); \\
M_C := \text{upd}_{\Sigma_C^{\beta^* \setminus G_r}}(M_C, \text{rd}(i_{32} \rightarrow i_8)); \\
v := \text{rd}(\gamma); \\
\beta(v), \beta_M(\beta^* \setminus G_r) := \beta^*, \beta^*;
\end{array}}
\end{array}$$

Figure 2.5: Translation rules for converting LLVM_d instructions to graph instructions.

‘`ret void`’ and ‘`ret v`’ respectively. Return from a procedure produces a non-silent observable event with the return value (in case of `ret v`) and the memory state of accessible regions in \mathbb{C} written to outside world through `wr(ret(...))`.

The (ALLOC) and (DEALLOC) rules (fig. 2.5) give the semantics for the allocation and deallocation of local memory, identified by an allocation site z , through `alloc` and `dealloc` respectively. For an allocation, if $z \in Z_l$ ⁷, the computation of allocation size, obtained by multiplying the number of elements allocated (n) with size of each element ($\text{sz}(\tau)$), has a *no overflow* constraint for a UB-free execution (shown through translation-selecting $\text{IF}\{z \in Z_l\}$). The translation uses the choose instruction ($\theta(\text{i}_{32})$) for identifying the non-deterministic start address of the freshly allocated interval $[\alpha_b, \alpha_e]$. A freshly allocated interval must satisfy the two well-formedness (WF) constraints of *no overlap* with existing allocated regions and *alignment* of the start address, implemented through $\neg \text{intrvlInSet}_a()$ check (defined in table 2.1), otherwise error \mathcal{W} is triggered⁸. An allocation adds the allocated interval $[\alpha_b, \alpha_e]$ to the address set $\Sigma_{\mathbb{C}}^z$ of the local; a deallocation empties it. Similar to the start address, the memory contents of the allocated interval are non-deterministically initialized through $\text{upd}_{[\alpha_b, \alpha_e]}(M_{\mathbb{C}}, \theta(\text{i}_{32} \rightarrow \text{i}_8))$. The various ghost variables associated with region z are updated in both cases: an allocation updates the lower bound `lb.z`, upper bound `ub.z`, and last allocation size `lstSz.z` and resets the boolean ghost variable `em.z` that tracks the emptiness of z ; a deallocation simply sets `em.z` to `true`. We use the simultaneous assignment instruction for updating $\Sigma_{\mathbb{C}}^z$, $M_{\mathbb{C}}$, and the ghost variables in a single step.

A (de)allocation instruction generates observable traces using the `wr` instruction at the beginning and end of each execution of that instruction. We will later use these traces to identify a lockstep correlation of (de)allocation events between \mathbb{C} and \mathbb{A} , towards validating a translation.

Modeling procedure calls in \mathbb{C}

The semantics of an LLVM_d procedure-call instruction is given by the rules (CALLV) and (CALL \mathbb{C}) (fig. 2.5). For an LLVM_d call instruction “`call γ ρ ($\vec{\tau}$ \vec{x})`”, we produce a non-silent observable trace event using the `wr` instruction with observables callee name/address ρ , arguments \vec{x} , and callee-accessible regions and memory state (β^*

⁷Recall that Z_l denotes the allocation sites due to the declaration of a local variable or a procedure-call argument

⁸Recall that \mathcal{W} represents an error condition that is external to the procedure and can be assumed to never occur.

and $\pi_{\Sigma_C^{\beta^*}}(M_C)$ in fig. 2.5). A callee may access a memory region iff it is *transitively reachable* from a global variable $g \in G$, the heap hp , or one of the arguments $x \in \vec{x}$. The (transitively) reachable memory regions are over-approximately computed through a reflexive-transitive closure of β_M , denoted β_M^* in fig. 2.5.

To model return values and side-effects to the memory state due to a callee, `rd` instructions are used. A `rd` instruction is used to arbitrarily clobber each the callee-observable state element. Thus, if a callee procedure terminates normally (i.e., without error), `wr` and `rd` instructions over-approximately model the execution of a procedure call. Later, our definition of refinement (section 2.4) caters to the case when a callee procedure may not terminate or terminates with error (i.e., a termination with error is modeled identically to non-termination).

Lastly, a procedure call can potentially be recursive — our modeling does not differentiate between a recursive and a non-recursive call. As we will see later in section 2.5.4, a consequence of this over-approximate modeling is that the tail-call elimination optimization, where a tail recursive call is replaced by a loop, cannot be covered by our refinement definition (i.e., the transformed procedure will not be considered refinement of the original procedure).

2.3.2 Translation of **A**

The rules for translating assembly instructions to graph instructions are shown in figs. 2.6 and 2.7 (and later figs. 2.8, 2.10 and 2.11). We abstract the assembly opcodes to an IR-like syntax for ease of exposition. For example, in (LOAD_A) , a memory read operation is represented by a `load` instruction which is annotated with address p , access size w (in bytes), and required alignment a^9 . Similarly, in (STORE_A) , a memory write operation is represented by a `store` instruction with similar operands. Both (LOAD_A) and (STORE_A) translations update the ghost address sets Σ_A^{rd} and Σ_A^{wr} , in the same manner as done in (LOAD_C) and (STORE_C) . A memory access error due to NULL address dereference or an unaligned access or an out-of-bounds access triggers a \mathcal{U} error (indicating a translation error). An access through address p is deemed out-of-bounds if p lies in `free` region ($p \in \Sigma_A^{\text{free}}$) or p lies in that part of `cv` region which does not overlap with assembly-only regions $F \cup S^{10}$. For a `store`, accessing read-only regions $(G_r \cup F_r)$ is also considered as out-of-bounds. Other machine exceptions such

⁹In 32-bit x86 alignment is only mandatory for some instructions (e.g., vector instructions).

¹⁰Recall that the `cv` region represents the inaccessible subset of local memory in call chain of **A**.

$$\begin{array}{c}
 \text{(OP-ESP)} \frac{p_A^j : \text{esp} := \text{op}(\vec{x})}{\begin{array}{l} \text{if } (\text{UB}_A(\text{op}, \vec{x})) \text{ halt}(\mathcal{U}); \\ t := \text{op}(\vec{x}); \\ \text{if } (\text{isPush}(p_A^j, \text{esp}, t)) \{ \\ \quad \text{if } (\neg \text{intrvlInSet}(t, \text{esp} - 1_{i_{32}}, \Sigma_A^{\text{free}} \cup (\Sigma_A^{cv} \setminus \Sigma_A^F))) \\ \quad \quad \text{halt}(\mathcal{W}); \\ \quad \Sigma_A^{stk} := \Sigma_A^{stk} \cup [t, \text{esp} - 1_{i_{32}}]; \\ \quad M_A := \text{upd}_{[t, \text{esp} - 1_{i_{32}}]}(M_A, \theta(i_{32} \rightarrow i_8)); \\ \quad \} \text{ else if } (t \neq \text{esp}) \{ \\ \quad \quad \text{if } (\neg \text{intrvlInSet}(\text{esp}, t - 1_{i_{32}}, \Sigma_A^{stk})) \\ \quad \quad \quad \text{halt}(\mathcal{U}); \\ \quad \quad \Sigma_A^{stk} := \Sigma_A^{stk} \setminus [\text{esp}, t - 1_{i_{32}}]; \\ \quad \quad \} \\ \text{esp} := t; \quad \boxed{\text{sp}.p_A^j} := t; \end{array}} \\
 \\
 \text{(LOAD}_A) \frac{p_A^j : v := \text{load } w, a, p}{\begin{array}{l} \text{if } (\neg \text{isAlignedIntrvl}_a(p, w) \\ \quad \vee \text{ov}([p]_w, \Sigma_A^{\text{free}} \cup (\Sigma_A^{cv} \setminus \Sigma_A^{FUS}))) \\ \quad \text{halt}(\mathcal{U}); \\ v := \text{sel}_w(M_A, p); \\ \boxed{\Sigma_A^{\text{rd}}} := \boxed{\Sigma_A^{\text{rd}}} \cup [p]_w; \end{array}} \quad \text{(OP-NESP)} \frac{p_A^j : r := \text{op}(\vec{x}) \quad r \neq \text{esp}}{\begin{array}{l} \text{if } (\text{UB}_A(\text{op}, \vec{x})) \text{ halt}(\mathcal{U}); \\ r := \text{op}(\vec{x}); \end{array}} \\
 \\
 \text{(STORE}_A) \frac{p_A^j : \text{store } w, a, v, p}{\begin{array}{l} \text{if } (\neg \text{isAlignedIntrvl}_a(p, w) \\ \quad \vee \text{ov}([p]_w, \Sigma_A^{\{\text{free}\} \cup G_r \cup F_r} \cup (\Sigma_A^{cv} \setminus \Sigma_A^{FUS}))) \\ \quad \text{halt}(\mathcal{U}); \\ M_A := \text{st}_w(M_A, p, v); \\ \boxed{\Sigma_A^{\text{wr}}} := \boxed{\Sigma_A^{\text{wr}}} \cup [p]_w; \end{array}}
 \end{array}$$

Figure 2.6: Translation rules for converting pseudo-assembly instructions to graph instructions. op represents an arithmetic, logical, or relational operator.

as division-by-zero are also modeled as \mathcal{U} errors in \mathbb{A} through the abstract $\text{UB}_A(\dots)$ operation (used in rules (OP-ESP) and (OP-NESP)).

(OP-ESP) (fig. 2.6) shows the translation of an instruction that updates the stackpointer register esp . An assignment to stackpointer esp may indicate allocation (stack push) or deallocation (stack pop) of stack space. An assignment that corresponds to a stackpointer decrement (push) is identified through predicate $\text{isPush}(p_A^j, \iota_b, \iota_a)$ where ι_b and ι_a are the values of esp before and after the execution of the instruction. We use thresholding on the update distance $(\iota_b - \iota_a)$ in our definition of isPush :¹¹

$$\text{isPush}(p_A^j, \iota_b, \iota_a) \Leftrightarrow (\iota_b \neq \iota_a) \wedge ((\iota_b - \iota_a) \leq_u (2^{31} - 1))$$

¹¹ $2^{31} - 1 = \text{INT_MAX}$ in our 32-bit setting.

While this choice of `isPush` suffices for most TV settings, we show in chapter A that if the translation is performed by an adversarial compiler, discriminating a stack push from a pop is trickier and may require external trusted guidance from the user.

For a stackpointer decrement (a push), a failure to allocate stack space, either due to wraparound or overlap with other allocated space, triggers \mathcal{W} , i.e., we expect the environment (e.g., OS) to ensure that the required stack space is available to A to prevent a wraparound or overlap; however, an overlap with region cv is permitted — we defer a discussion on this exception to section 2.6. For a stackpointer increment (a pop), it is a translation error if the stackpointer moves out of current stack frame bounds (captured by error code \mathcal{U}). The stackpointer value at the end of an assignment instruction at PC p_A^j is saved in a ghost variable named $\boxed{\text{sp}.p_A^j}$. These ghost variables help with inference of invariants that relate a local variable’s bounds with stack addresses (invariant inference in our algorithm is discussed in section 4.2). During push, the initial contents of the newly allocated stack region are chosen non-deterministically using θ — this admits the possibility of arbitrary clobbering of the unallocated stack region below the stackpointer due to asynchronous external interrupts, before it is allocated again.

The (OP-NESP) rule in fig. 2.6 gives translation for an instruction ‘ $r := \text{op}(\vec{x})$ ’ that does not update the `esp` register ($r \neq \text{esp}$). The $\text{UB}_A(\text{op}, \vec{x})$ operation abstracts the condition for a machine exception during execution of `op` (e.g., a zero second operand for division).

(ENTRY $_A$) (in fig. 2.7) shows the initialization of state elements of procedure A at entry. For a region $r \in B^{12}$, the initialization of address set Σ_A^r and memory region $\pi_{\Sigma_A^r}(M_{\tilde{A}})$ is same as (ENTRY $_C$). For an assembly-only region $f \in F$, the address set Σ_A^f is initialized using A ’s symbol table (abstracted through `addrSets $_F$ ()`). The memory contents of a read-only global variable $r \in G_r \cup F_r$ are initialized using $\text{ROM}_A^r(i_A^r)$ (defined in table 2.1) — recall that the validator verifies the equality of memory contents of a common read-only global $g \in G$ so that $\text{ROM}_A^g(i_A^g) = \text{ROM}_C^g(i_C^g)$ for $i_A^g = i_C^g$.

The machine registers are initialized with arbitrary contents using θ . The x86 stack of an assembly procedure includes the stack frame Σ_A^{stk} of the currently executing procedure A , the parameters Σ_A^Y of A , and the remaining space which includes caller-stack Σ_A^{cs} and, possibly, the locals Σ_A^{cl} defined in the call chain of A . Due to the calling conventions, we assume (through `stkIsWF()`) that:

¹²Recall that B is the set of regions common to both C and A .

$$\begin{array}{c}
\text{(ENTRY}_A) \frac{p_A^j : \text{def } A(\vec{\tau})}{\begin{array}{l}
\Sigma_A^{hp}, \Sigma_A^{cl}, \Sigma_A^{cv}, \dots, i_A^g, \dots, i_A^y, \dots, \Sigma_A^{vrdc} := \text{rd}(2^{i_{32}}, 2^{i_{32}}, \dots, 2^{i_{32}}); \\
\dots, \Sigma_A^f, \dots := \text{addrSets}_F(); \\
\Sigma_A^{rd}, \Sigma_A^{wr}, \dots, \Sigma_A^z, \dots := \emptyset, \emptyset, \dots, \emptyset, \dots; \\
\text{if } (\neg \text{addrSetsAreWF}(\Sigma_A^{hp}, \Sigma_A^{cl}, \Sigma_A^{cv}, \dots, i_A^g, \dots, \Sigma_A^f, \dots, i_A^y, \dots, \Sigma_A^{vrdc})) \\
\quad \text{halt}(\mathcal{W}); \\
M_A := \theta(i_{32} \rightarrow i_8); \quad M_A := \text{upd}_{\Sigma_B \setminus G_r}(M_A, \text{rd}(i_{32} \rightarrow i_8)); \\
\text{for } r \text{ in } G_r \cup F_r \{ M_A := \text{upd}_r(M_A, \text{ROM}'_A(i'_A)); \} \\
\text{for } x \text{ in } \overline{\text{regs}} \{ x := \theta(T(x)); \} \\
\text{stk}_e := \Sigma_A^Y \neq \emptyset ? \text{ub}(\Sigma_A^Y) : \text{esp} + 3i_{32}; \\
\text{cs}_e := \theta(i_{32}); \\
\text{if } (\neg \text{stkIsWF}(\text{esp}, \text{stk}_e, \text{cs}_e, \vec{\tau}, \Sigma_A^{hp}, \Sigma_A^{cl}, \Sigma_A^{GUF}, \dots, i_A^y, \dots, \Sigma_A^{vrdc})) \\
\quad \text{halt}(\mathcal{W}); \\
\Sigma_A^{stk} := [\text{esp}, \text{stk}_e] \setminus \Sigma_A^Y; \\
\Sigma_A^{cs} := [\text{stk}_e + 1i_{32}, \text{cs}_e] \setminus \Sigma_A^{cl}; \\
\text{sp.entry} := \text{esp}; \quad M^{cs} := \pi_{\Sigma_A^{cs}}(M_A); \\
\text{ebp}, \text{esi}, \text{edi}, \text{ebx}, \text{eip} := \text{ebp}, \text{esi}, \text{edi}, \text{ebx}, \text{sel}_4(M_A, \text{esp}); \\
\text{for } f \text{ in } F \{ \\
\quad \text{sz.f}, \text{em.f}, \text{lb.f}, \text{ub.f} := |\Sigma_A^f|, |\Sigma_A^f| = 0_{i_{32}}, \text{lb}(\Sigma_A^f), \text{ub}(\Sigma_A^f); \\
\} \\
\end{array}} \\
\text{(RET}_A) \frac{p_A^j : \text{ret } \tau}{\begin{array}{l}
\text{if } (\text{sp.entry} \neq \text{esp} \vee \text{ebp} \neq \text{ebp} \\
\vee \text{esi} \neq \text{esi} \vee \text{edi} \neq \text{edi} \vee \text{ebx} \neq \text{ebx} \\
\vee \text{eip} \neq \text{sel}_4(M_A, \text{esp}) \\
\vee \neg (M^{cs} =_{\Sigma_A^{cs}} M_A)) \\
\quad \text{halt}(\mathcal{W}); \\
\text{IF}\{\tau = \text{void}\}\{ \\
\quad \text{wr}(\pi_{\Sigma_B}(M_A)); \\
\} \text{ ELSE } \{ \\
\quad \text{wr}(\text{ret}(\Delta_\tau(\text{eax}, \text{edx}), \pi_{\Sigma_B}(M_A))); \\
\} \\
\text{halt}(\emptyset);
\end{array}}
\end{array}$$

Figure 2.7: Translation rules for converting pseudo-assembly instructions to graph instructions.

1. The parameters are laid out at addresses above the stackpointer (esp) as per calling conventions (abstracted through $\text{obeyCC}()$ in $\text{stkIsWF}()$).
2. The value $\text{esp} + 4i_{32}$ (stackpointer value at the callsite in caller) is 16-byte aligned.
3. The callers' stack (region cs) is above A 's stack frame and laid out parameters (if any).
4. Stack stk and callers' stack cs do not overlap with other allocated space.

A violation of these conditions trigger \mathcal{W} , i.e., we expect the environment to ensure that these conditions are satisfied. The ghost variables $\boxed{\text{stk}_e}$ and $\boxed{\text{cs}_e}$ represent the largest addresses in $\Sigma_A^{Y \cup \{\text{stk}\}}$ and $\Sigma_A^{Y \cup \{\text{stk}, \text{cs}, \text{cl}\}}$ respectively, so that at entry, $\Sigma_A^{\text{stk}} = [\text{esp}, \boxed{\text{stk}_e}] \setminus \Sigma_A^Y$ and $\Sigma_A^{\text{cs}} = [\boxed{\text{stk}_e} + 1_{i_{32}}, \boxed{\text{cs}_e}] \setminus \Sigma_A^{\text{cl}}$. If there are no parameters, $\boxed{\text{stk}_e} = \text{esp} + 3_{i_{32}}$ represents the end of the region that holds the return address of procedure A . The rule concludes with initialization of ghost variables for: stack pointer at entry ($\boxed{\text{sp.entry}}$); memory contents of the caller's stack ($\boxed{M^{\text{cs}}}$); return address of A ($\boxed{\text{eip}}$); callee-saved registers ($\boxed{\text{ebp}}, \boxed{\text{esi}}, \boxed{\text{edi}}, \boxed{\text{ebx}}$); and the address set of an assembly-only global variable $f \in F$ ($\boxed{\text{sz.f}}, \boxed{\text{em.f}}, \boxed{\text{lb.f}}, \boxed{\text{ub.f}}$).

Upon procedure-return (rule (RET_A) in fig. 2.7), we require, as per the calling conventions, that the return address, callers' stack and the callee-save registers remain preserved — a violation of these conditions trigger \mathcal{U} . We use the ghost variables, $\boxed{\text{sp.entry}}, \boxed{\text{ebp}}, \boxed{\text{esi}}, \boxed{\text{edi}}, \boxed{\text{eip}}$, and $\boxed{M^{\text{cs}}}$, set at procedure entry for this check. Validating the calling conventions at procedure return enable us to assume them at a procedure call. For simplicity, we only tackle scalar return values, and ignore aggregate return values that need to be passed in memory.

Notice that unlike region $r \in B$, region cv may potentially overlap with assembly-only regions $F \cup S$. Thus, while an address $\alpha \in \Sigma_C^{cv}$ is inaccessible in C , it is potentially accessible in A if $\alpha \in F \cup S$. We explain the rationale for this in section 2.6.1 when we discuss virtual allocation.

2.4 Observable traces and Refinement Definition

Recall that a procedure $P \in \{C, A\}$ execution yields an observable trace T_P containing silent and non-silent events.

Definition 2.4.1 ($e(T)$). *The error code of a trace T , written $e(T)$, is either \emptyset (indicating either non-termination or error-free termination), or one of $\mathcal{r} \in \{\mathcal{U}, \mathcal{W}\}$ (indicating termination with error code \mathcal{r}).*

For $\mathcal{r} \in \{\mathcal{U}, \mathcal{W}\}$, we call a trace T \mathcal{r} -terminating iff $e(T) = \mathcal{r}$.

Definition 2.4.2 ($\tilde{e}(T)$). *The non-error part of a trace T , written $\tilde{e}(T)$, is T when $e(T) = \emptyset$ and T' such that $T = T' \cdot e(T)$ otherwise.*

Definition 2.4.3 ($(P \downarrow_{\Omega} T)$). *$(P \downarrow_{\Omega} T)$ denotes the condition that for an initial outside*

world Ω , the execution of a procedure P may produce an observable trace T for some sequence of non-deterministic choices.

A compiler must ensure that if A exhibits UB (Undefined Behavior) then there must exist a sequence of non-deterministic choices such that C also exhibits UB, i.e., $(A \downarrow_{\Omega} T \cdot \mathcal{U})$ implies $(C \downarrow_{\Omega} T' \cdot \mathcal{U})$. Further, for an error-free execution of A , C should either be able to produce a trace containing identical sequence of non-silent events or trigger \mathcal{U} — the latter case admitting the *anything is permissible* clause of UB in C . We define a refinement relation between C and A that ensures these properties.

Definition 2.4.4 ($T =_{st} T'$). *Traces T and T' are **stuttering equivalent**, written $T =_{st} T'$, iff they differ only by finite sequences of silent events \perp .*

For example, $T =_{st} T'$ holds for $T = (\text{rd}(\dots), \perp, \perp, \text{rd}(\dots), \perp, \perp, \perp, \text{wr}(\dots), \text{exit})$ and $T' = (\text{rd}(\dots), \perp, \text{rd}(\dots), \perp, \text{wr}(\dots), \text{exit})$.

Definition 2.4.5 ($T \leq_{st} T'$). *A trace T is a **stuttering prefix** of trace T' , written $T \leq_{st} T'$, iff $(T =_{st} T') \vee (\exists T_r : (T \cdot T_r) =_{st} T')$.*

For example, $T \leq_{st} T'$ holds for $T = (\text{rd}(\dots), \perp, \perp, \text{rd}(\dots), \perp, \perp, \perp)$ and $T' = (\text{rd}(\dots), \perp, \text{rd}(\dots), \perp, \text{wr}(\dots), \text{exit})$.

Definition 2.4.6 ($W_{\text{pre}}^{\Omega, T_A}(C)$). $W_{\text{pre}}^{\Omega, T_A}(C)$ denotes the condition:

$$(e(T_A) = \mathcal{W}) \wedge (\exists T_C : (C \downarrow_{\Omega} T_C) \wedge (\tilde{e}(T_A) \leq_{st} T_C))$$

Definition 2.4.7 ($U_{\text{pre}}^{\Omega, T_A}(C)$). $U_{\text{pre}}^{\Omega, T_A}(C)$ denotes the condition:

$$\exists T_C : (C \downarrow_{\Omega} T_C \cdot \mathcal{U}) \wedge (T_C \leq_{st} T_A)$$

Definition 2.4.8 ($C \sqsupseteq A$). $C \sqsupseteq A$, read A refines C (or C is refined by A), iff:

$$\begin{aligned} \forall \Omega : (A \downarrow_{\Omega} T_A) \Rightarrow & W_{\text{pre}}^{\Omega, T_A}(C) \\ & \vee U_{\text{pre}}^{\Omega, T_A}(C) \\ & \vee \exists T_C : (C \downarrow_{\Omega} T_C) \wedge (T_A =_{st} T_C) \end{aligned}$$

The definition of $C \sqsupseteq A$ admits three possibilities for an execution of A : either (1) A triggers \mathcal{W} ($W_{\text{pre}}^{\Omega, T_A}(C)$), or (2) C triggers \mathcal{U} ($U_{\text{pre}}^{\Omega, T_A}(C)$), or (3) C may produce a trace with identical sequence of non-silent events to the one produced by A ($T_A =_{st} T_C$).

In the first case, $W_{\text{pre}}^{\Omega, T_A}(\mathbf{C})$ encodes the condition that \mathbf{A} terminates with error \mathscr{W} ($e(T_A) = \mathscr{W}$) and the sequence of non-silent events in the trace T_A produced by \mathbf{A} (before terminating) is equivalent to trace T_C produced by \mathbf{C} up till that point ($\tilde{e}(T_A) \leq_{st} T_C$) — we call the latter a *trace prefix requirement*. Recall that we do not care about the case where \mathbf{A} terminates with \mathscr{W} because we assume that the external environment (e.g., OS) will ensure that this event never occurs. For example, we do not care for the case where a stack space allocation may fail in \mathbf{A} (OP-ESP). Thus, in the example in fig. 2.1c, we assume that the stackpointer decrement instruction at **A5** successfully allocates stack space. The trace prefix requirement $\tilde{e}(T_A) \leq_{st} T_C$ caters to the situation where a callee procedure may not terminate: if a callee in \mathbf{A}^{13} does not terminate (before \mathbf{A} could halt with \mathscr{W}), then the trace prefix requirement ensures that the corresponding callee in \mathbf{C} will also not terminate (because identical trace events have been recorded for each procedure call till that point). For example, in fig. 2.1c, consider the case when the stackpointer manipulation instructions at **A20** in the assembly procedure \mathbf{A}_{fib} trigger \mathscr{W}^{14} : it is possible that before \mathscr{W} may be triggered at **A20**, the preceding call to `printf` never terminates (or encounters error resulting in termination of program execution). In this situation, \mathbf{A}_{fib} would not actually trigger \mathscr{W} at runtime (because **A20** will never execute). Our execution semantics do not explicitly model the possibility of non-termination (or termination with error) of `printf` and so they make it appear that **A20** will be executed and so \mathscr{W} will be triggered. Thus, it is not enough to simply ignore the case where \mathscr{W} is triggered; we also need to ensure that all prior procedure calls have identical (potentially non-terminating or termination with error) behavior. The trace prefix requirement, $\tilde{e}(T_A) \leq_{st} T_C$, ensures this as it requires the traces due to a procedure call to be identical up till the point of error in \mathbf{A} . If the trace prefix requirement was absent, we may unsoundly admit a translation that passes differing arguments to a procedure call (`printf` in fig. 2.1) in \mathbf{C} and \mathbf{A} but triggers \mathscr{W} in \mathbf{A} thereafter.

In the second case, $U_{\text{pre}}^{\Omega, T_A}(\mathbf{C})$ encodes the condition that \mathbf{C} may terminate with error \mathscr{U} (i.e., it may exhibit UB). Due to UB semantics, we therefore do not care about \mathbf{A} 's behavior in this situation. For example, in fig. 2.1b, for `v[*i]` at **I6**, we do not care for the out-of-bounds case when `*i > (*n + 2)`. The stuttering prefix condition, $T_C \leq_{st} T_A$,

¹³We discuss the exact execution semantics of a procedure call in \mathbf{A} in next section. For now, let's assume that the semantics are similar to \mathbf{C} with `wr` instruction for over-approximately observing and `rd` for arbitrarily but deterministically mutating the callee-observable state.

¹⁴Strictly speaking, \mathscr{W} will not be triggered in fig. 2.1c at **A20**, because the execution would have halted earlier in this situation (for this example). However, for argument's sake, assume the error may get triggered at the mentioned PC.

ensures that all procedure calls (before C exhibits UB) produce identical traces in both C and A , thus ensuring identical termination behavior of prior procedure calls (similar to the first case).

In the third case, C may produce a stuttering equivalent trace as A , and all non-silent events (including non-termination) will be identical.

In each of the cases, a stuttering requirement, stuttering prefix in the first two cases and stuttering equivalent in the third case, ensures that both procedures execute at similar speeds up till the termination of either. It is to be noted that stuttering requirement is not related to preservation of termination behavior — the latter is due to observation of termination through the trace-producing `halt` instruction. A terminating execution has a finite trace ending with an error code (for erroneous termination) or the special `exit` event (for error-free termination). Due to the silent trace event \perp generated by execution of every instruction, a non-terminating execution has an infinite trace.

In the absence of local variables and procedure calls in C , $C \sqsupseteq A$ implies a correct translation from C to A .

In our work, $C \sqsupseteq A$ is our definition of a *correct translation*. In subsequent sections, this formal definition will become more sophisticated, as it will involve annotating the assembly program with custom instructions introduced by us. These formal definitions, which include the custom instructions introduced by us, have been carefully constructed so that they capture the notion of observable equivalence as people usually understand it. We expect that these definitions are succinct and simple enough so that the reader can convince herself of their validity. We help the reader by introducing these formal definitions incrementally, with supporting explanations.

2.5 Refinement Definition in the presence of local variables and procedure calls when all local variables are allocated on the stack in A

For each local variable (de)allocation and for each procedure call, our execution semantics generate a `wr` trace event in C (fig. 2.5). Thus, to reason about refinement, we require correlated and equivalent trace events to be generated in A . For this, we annotate A with two types of annotations to obtain \dot{A} :

1. *Special `allocs` and `deallocs` instructions:* `allocs` and `deallocs` instructions are added to explicitly indicate the (de)allocation of a local variable $z \in Z$ from stack. An address interval in the stack region may be marked as belonging to z through the `allocs` instruction and subsequently marked as belonging to stack again through the `deallocs` instruction.
2. *Annotations to assembly procedure-call instruction:* A procedure call is annotated with the types and addresses of the arguments and the set of memory regions observable by the callee. These annotations augment the assembly call instruction such that its semantics can be defined in a manner similar to the `call` instruction in `C` (section 2.3.1).

These annotations are intended to encode the correlations with the corresponding allocation, deallocation, and procedure-call events in the source procedure `C`. For now, we assume that the locations and values of these annotations in \dot{A} are coming from an oracle — later in chapter 4, we present an algorithm to identify these annotations automatically in a best-effort manner.

Figure 2.8 presents the translations of the three new assembly instructions — `allocs`, `deallocs`, and `call` — to graph instructions.

2.5.1 (De)Allocation indicating `allocs` and `deallocs` instructions

An instruction ‘ $p_{\dot{A}}^j : \text{alloc}_s e_v, e_w, a, z$ ’ represents the *stack allocation* of a local variable identified by allocation site z (fourth argument) at PC $p_{\dot{A}}^j$. e_v is the expression for the start address, e_w is the expression for the allocation size, and a is the required alignment of the start address. (ALLOCS) in fig. 2.8 presents the translation of `allocs` to graph instructions. During the stack allocation of z , the allocated interval $i = [v, v + w - 1_{i32}]$, identified by start address $v = e_v$ and allocation size $w = e_w$, must satisfy the required well-formedness (WF) constraints for separation and alignment, or else \mathcal{U} is triggered. Recall that \mathcal{U} is used for signaling translation errors in A , so a correctly translated assembly procedure must never falsify the separation and alignment constraints for i . A stack allocation removes i from the stack address set $\Sigma_{\dot{A}}^{stk}$ and adds it to the address set $\Sigma_{\dot{A}}^z$ of z . Thus, the separation WF constraint requires i to lie entirely within $\Sigma_{\dot{A}}^{stk}$ (encoded through `intrvlInSeta(..., $\Sigma_{\dot{A}}^{stk}$)` in fig. 2.8). Further,

$$\begin{array}{c}
 \text{(ALLOCS)} \frac{p_{\dot{A}}^j : \text{alloc}_s e_v, e_w, a, z}{\text{wr}(\text{allocBegin}(z, e_w, a));} \\
 v, w := e_v, e_w; \\
 \text{if } (\neg \text{intrvlInSet}_a(v, v + w - 1_{i_{32}}, \Sigma_{\dot{A}}^{stk})) \\
 \quad \text{halt}(\mathcal{U}); \\
 \text{if } (\text{ov}([v]_w, \Sigma_{\dot{A}}^{cv})) \\
 \quad \text{halt}(\mathcal{W}); \\
 \Sigma_{\dot{A}}^z, \Sigma_{\dot{A}}^{stk} := \Sigma_{\dot{A}}^z \cup [v]_w, \Sigma_{\dot{A}}^{stk} \setminus [v]_w; \\
 \text{wr}(\text{allocEnd}(z, [v]_w, \pi_{[v]_w}(M_{\dot{A}})));
 \end{array}
 \qquad
 \begin{array}{c}
 \text{(DEALLOCS)} \frac{p_{\dot{A}}^j : \text{dealloc}_s z}{\Sigma_{\dot{A}}^z := \emptyset, \Sigma_{\dot{A}}^{stk} := \Sigma_{\dot{A}}^{stk} \cup \Sigma_{\dot{A}}^z;} \\
 \text{wr}(\text{dealloc}(z));
 \end{array}$$

$$\begin{array}{c}
 \text{(CALL}_{\dot{A}}) \frac{p_{\dot{A}}^j : \text{call } \gamma \rho(\vec{\tau} \vec{x}) \beta^*}{\text{if } (\neg \text{aligned}_{16}(\text{esp}) \vee \neg \text{obeyCC}(\text{esp}, \vec{\tau}, \vec{x}))} \\
 \quad \text{halt}(\mathcal{U}); \\
 \text{wr}(\text{fcall}(\rho, \vec{x}, \beta^*, \pi_{\Sigma_{\dot{A}}^{stk}}(M_{\dot{A}}))); \\
 M_{\dot{A}} := \text{upd}_{\Sigma_{\dot{A}}^{stk} \setminus G_r}(M_{\dot{A}}, \text{rd}(i_{32} \rightarrow i_8)); \\
 \text{ecx} := \theta(i_{32}); \\
 \text{IF}\{\gamma = \text{void}\}\{ \\
 \quad \text{eax, edx} := \theta(i_{32}, i_{32}); \\
 \} \text{ ELSE } \{ \\
 \quad \text{eax, edx} := \nabla_{\gamma}(\text{rd}(\gamma)); \\
 \}
 \end{array}$$

Figure 2.8: Additional translation rules for converting pseudo-assembly instructions to graph instructions for procedures with only stack-allocated locals.

the allocated interval must be separate from region cv ¹⁵, otherwise \mathcal{W} is triggered; we explain the rationale for triggering \mathcal{W} in this case in the next section when we discuss virtual allocation.

An instruction ‘ $p_{\dot{A}}^j : \text{dealloc}_s z$ ’ represents the deallocation of z and empties the address set $\Sigma_{\dot{A}}^z$, adding the removed addresses to $\Sigma_{\dot{A}}^{stk}$. This action reverses the transfer from the stack to z performed by execution of an alloc_s instruction. (DEALLOCS) in fig. 2.8 shows the translation of dealloc_s to graph instructions.

The alloc_s and dealloc_s instructions do not change the accessible address set $\Sigma_{\dot{A}}^{BUFUS}$ of \dot{A} : alloc_s transfers addresses from stk to z and dealloc_s transfers them back to stk . Thus, no “real allocation”¹⁶ is performed, instead a subset of the stack frame of \dot{A} is identified distinctly for the purpose of validation.

¹⁵Recall that cv may potentially overlap with stk unlike a region $r \in B$.

¹⁶In the sense of stack allocation through stackpointer decrement in \dot{A} or local allocation through alloc in C .

Similar to `alloc` and `dealloc` in \mathbb{C} , `allocs` and `deallocs` in \mathring{A} produce non-silent trace events via `wr` instructions. Both use the same value constructors, so that identical observables (passed as arguments) produce identical trace events.

Figure 2.1c shows an assembly procedure annotated with `allocs` instruction at line A5¹ and `deallocs` instruction at line A19¹. The start address of allocation is `esp`, the allocated size is `4*(eax+2)`, and the required alignment on start address is 4. It can be observed that the stack allocated interval meets the WF constraints in this case: the preceding instruction A5 allocates at least `4*(eax+2)` bytes on stack and aligns `esp` by 16. The `allocs` instruction uses the allocation site I2 from the IR procedure in fig. 2.1b for identifying the allocated interval. The `deallocs` instruction at A19¹ refers to the same allocation site I2 for transferring the allocated addresses due to execution of `allocs` back to stack.

2.5.2 Annotated procedure-call instruction

We annotate an assembly procedure-call instruction ‘ $p_{\mathring{A}}^j : \text{call } \rho$ ’ for a call to callee ρ as ‘ $p_{\mathring{A}}^j : \text{call } \gamma \rho(\vec{\tau} \vec{x}) \beta^*$ ’ to explicitly specify:

- The start addresses \vec{x} of the address regions belonging to the arguments.
- The types $\vec{\tau}$ of the arguments and the return type γ .
- The callee-observable regions β^* .

The address region of an argument should have previously been demarcated using an `allocs` instruction to match the `alloc` instructions for allocating arguments to a procedure call in \mathbb{C} . Additionally, these address regions should satisfy the constraints imposed by the calling conventions — represented through `obeyCC()` in rule (CALL \mathring{A}) of fig. 2.8. The calling conventions also require the stackpointer `esp` to be 16-byte aligned. A failure to meet the calling conventions requirements trigger \mathcal{U} to indicate a translation error. A procedure call is recorded as an observable event, along with the observation of the callee name (or address) ρ , the addresses of the arguments \vec{x} , the callee-observable regions and their memory contents (β^* and $\pi_{\Sigma\beta^*}(M_{\mathring{A}})$), same as (CALL \mathbb{C}). The returned values, modeled through `rd(i32 → i8)` and `rd(γ)`, include the contents of the callee-observable memory regions and the scalar values returned by the callee (in registers `eax`, `edx`). The callee additionally clobbers the caller-save registers (`eax`, `ecx`, `edx`) using θ . We use the translation-selecting `IF{ $\gamma = \text{void}$ }{...}ELSE{...}` construct in (CALL \mathring{A}) for selecting translation, whether to clobber or read, based on

return type γ .

Notice that for an annotated call instruction in \dot{A} to produce identical observables as a call instruction in C , it must not only have identical callee but identical argument addresses, callee-observable regions, and memory contents of callee-observable regions — the latter three being determined by the annotations. While the address regions, identified from $\vec{\tau} \vec{x}$, are validated against the calling conventions (through `obeyCC()`), we do not “validate” the callee-observable regions β^* and the return type γ , both of which are consequential for defining the semantics of `call`. Instead, we rely on correlation with corresponding `call` instruction in C : an instruction `call ρ` in \dot{A} , say `call \dot{A}` , with stack-allocated arguments that match in address and memory contents with instruction `call $\rho \dots$` in C , say `call C` , is *assumed* to have identical mutation behavior, i.e., `call \dot{A}` mutates \dot{A} 's state identically as `call C` mutates C 's state. In other words, the callee-observable regions (and return value) of `call \dot{A}` in \dot{A} are determined based on the behavior of a correlated `call C` in C (with identical arguments).

Figure 2.1c shows an annotated `call` instruction at line A18; the annotated callee-observable regions $\{hp, cl, I9, I10\}$ indicate that `printf` may potentially access and mutate the heap (`hp`), accessible callers' locals (`cl`), and the stack subregion corresponding to the two arguments (`I9` and `I10`) but not the rest of the stack.

2.5.3 Refinement Definition with only stack-allocated locals and procedure calls

With the new instructions and annotations enabling presence of stack-allocated locals and procedure calls in \dot{A} , we define refinement between C and \dot{A} in terms of existence of an annotated \dot{A} such that \dot{A} refines C :

Definition 2.5.1 (Refinement with only stack-allocated locals and procedure calls).

$C \dot{\sqsubseteq} \dot{A}$ iff: $\exists \dot{A} : C \sqsubseteq \dot{A}$

$C \dot{\sqsubseteq} \dot{A}$ encodes the property that it is possible to annotate \dot{A} to obtain \dot{A} so that the local variable (de)allocation and procedure-call events of C and the annotated \dot{A} can be correlated in lockstep.

Soundness of an annotation

It must not be possible to annotate A to produce \dot{A} such that $C \sqsupseteq \dot{A}$ holds but the two procedures (C and \dot{A}) have different observable behavior. Our annotated instructions are carefully constructed and generate observable events such that the refinement definition will never admit an incorrect translation. An informal argument in favor of soundness of \dot{A} is as follows. An assembly procedure \dot{A} , produced by annotating A with alloc_s , dealloc_s , and call instructions, may have executions that are not present in the unannotated procedure A such that these additional executions either terminate with error \mathcal{U} (due to alloc_s and call) or terminate with error \mathcal{W} (due to alloc_s). We consider each case of error separately below.

- If \dot{A} terminates with error \mathcal{U} , due to either alloc_s or call instruction, producing a trace $T_{\dot{A}}$, then $C \sqsupseteq \dot{A}$ requires C to have an execution with trace T_C such that either T_C and $T_{\dot{A}}$ are stuttering equivalent ($T_C =_{st} T_{\dot{A}}$ holds), i.e., C also terminates with \mathcal{U} , or T_C ends with \mathcal{U} and the non-error part of T_C is stuttering prefix of $T_{\dot{A}}$ ($\tilde{e}(T_C) \leq_{st} T_{\dot{A}}$ holds). In both cases, C must exhibit UB before \dot{A} does.
- If \dot{A} terminates with error \mathcal{W} , producing a trace $T_{\dot{A}}$, then $C \sqsupseteq \dot{A}$ holds if there exists an execution of C with trace T_C such that C produces identical sequence of observables as \dot{A} before the latter's termination ($\tilde{e}(T_{\dot{A}}) \leq_{st} T_C$).

Recall that \dot{A} 's execution terminates with error \mathcal{W} due to $\text{ov}([v]_w, \Sigma_{\dot{A}}^{cv})$ condition in (ALLOCS) (fig. 2.8). Before this, the execution must have produced an observable event through $\text{wr}(\text{allocBegin}(\dots))$ that must be present in T_C as well. This indicates that C also executed an alloc instruction ((ALLOC) in fig. 2.5) with identical allocation size, alignment, and region identifier (all part of observables produced through allocBegin). Further, because the execution of \dot{A} did not trigger \mathcal{U} , the region to be allocated, $[v]_w$, belonged to stack $\Sigma_{\dot{A}}^{stk}$ ($\text{intrv1InSet}()$ check). The execution semantics of \dot{A} prohibit overlap of region stk with B and because observable events in T_C and $T_{\dot{A}}$ match till occurrence of error, $[v]_w$ must belong to $\text{comp}(\Sigma_C^B) \cup \Sigma_C^{cv}$. Consequently, there exists a choice of interval $[\alpha_b, \alpha_e]$ such that the execution of C also triggers \mathcal{W} and terminates. Thus, it is possible for C and \dot{A} to produce identical observables.

Notice that the exchange of addresses between $\Sigma_{\dot{A}}^{stk}$ and $\Sigma_{\dot{A}}^z$ (for $z \in Z$) in alloc_s and dealloc_s instructions does not affect the ‘‘out-of-bounds’’ access checks in (LOAD $_{\dot{A}}$) and (STORE $_{\dot{A}}$). For an error-free execution of alloc_s , dealloc_s , and call in \dot{A} , $C \sqsupseteq \dot{A}$ will require identical observable events, allocBegin and allocEnd for alloc_s , dealloc

for `deallocs`, and `fcall` for `call` to be produced in C as well. This concludes our soundness argument.

In the presence of stack-allocated local variables and procedure calls, $C \dot{\sqsupseteq} A$ implies a correct translation from C to A . In the absence of local variables and procedure calls, $C \dot{\sqsupseteq} A$ reduces to $C \sqsupseteq A$ with $\dot{A} = A$.

2.5.4 Capabilities and Limitations of $C \dot{\sqsupseteq} A$

Limitation on relative order of (de)allocation and procedure calls

Because our logical encoding observes each (de)allocation event (due to the `wr` instruction), a fundamental limitation of $C \dot{\sqsupseteq} A$ is that for allocations and procedure calls that *reuse the same stack space*, their relative order remains preserved. This requirement is sound but may be too strict for certain (arguably rare) compiler transformations that may reorder the (de)allocation instructions that reuse the same stack space.

Figure 2.9 shows an example of such a transformation where stack space is reused between allocations and procedure calls. The hypothetical assembly procedure A_{baz} (shown in fig. 2.9b) is a correct translation¹⁷ of the C procedure C_{baz} (in fig. 2.9a) but will not be admitted under $C \dot{\sqsupseteq} A$. This is because the relative order of allocation of the variable `x` and the procedure call `foo` is not preserved in the transformation from C_{baz} to A_{baz} . In the latter, the stack allocation of the local `x` is performed after the procedure call to `foo`. The stack deallocation at A3 prohibits an earlier placement of `allocs`, making an annotation \dot{A}_{baz} that meets $C_{\text{baz}} \sqsupseteq \dot{A}_{\text{baz}}$ impossible. Notice that the stack subregion used for allocating `x` was previously potentially used by `foo`. Figure 2.9c shows another possible compilation (generated by an optimizing compiler) of C_{baz} in which case an annotation is possible and refinement can be established.

Limited handling of interprocedural transformations

Recall that our translation rules for C and A associate production of a non-silent observable trace event with a procedure-call instruction (figs. 2.5 and 2.8). The name (or address) of the callee is also observed. Thus, an interprocedural transformation which eliminates a procedure call, inlines it (including partial inlining), specializes it (through procedure cloning), transforms the control-flow (e.g., tail-recursion elimination), uses

¹⁷As per the C standard and calling conventions for 32-bit

<pre> int baz() { int x; // alloc foo(); return bar(&x); } </pre>	<pre> A0: baz: A1: esp -= 12 ; for alignment A2: call foo A3: esp += 12 ; undo -12 above A4: esp -= 8 ; alloc 'x' A5: push (esp+4) ; setup &x A6: call bar A7: esp += 12 A8: ret </pre>	<pre> baz: esp -= 12 ; alloc 'x' call foo mem4[esp] = esp+8 call bar esp += 12 ret </pre>
(a) C program with address-taken local	(b) Hypothetical (abstracted) 32-bit x86 assembly	(c) Compiler generated (abstracted) 32-bit assembly

Figure 2.9: Example of transformation where relative order of (de)allocations and procedure calls is not preserved. The refinement definition will not admit the hypothetical assembly but will admit the compiler generated one.

a different calling convention (e.g., if the called procedure is not externally visible then compiler may pass some arguments through registers instead of stack), reuses the arguments of the caller procedure (so that `allocs` annotation is not possible for the callee arguments), or performs some other transformation that makes use of information not encoded in the semantics of `C` and `A` will not be admitted.

Key transformations admitted under $C \dot{\sqsubseteq} A$:

Merging of multiple allocations: $C \dot{\sqsubseteq} A$ supports *merging* of multiple allocations into a single stackpointer decrement instruction. Let p_A^s be the PC of a single stackpointer decrement instruction that implements multiple allocations. Merging can be encoded by adding multiple `allocs` instructions to `A`, in the same order as they appear in `C`, to obtain \dot{A} , so that these `allocs` instructions execute only after p_A^s executes. The C and assembly code fragments below illustrate the construction (C on the left and abstracted assembly on the right):

<pre> int x, y, z // alloc 'x' // alloc 'y' // alloc 'z' </pre>	<pre> esp -= 20 alloc_s 'x' alloc_s 'y' alloc_s 'z' </pre>
---	--

Similarly, the corresponding `deallocs` instructions must execute before a stackpointer

increment instruction deallocates this stack space.

CompCert[29]’s preallocation is a special case of merging where stack space for all local variables and procedure call arguments is allocated in the assembly procedure’s prologue and deallocated in the epilogue with no reuse of stack space. In this case, our approach annotates A with `allocs` and `deallocs` instructions, potentially in the middle of the procedure body, such that they execute in lockstep with the allocations and deallocations in C .

Reallocation of stack space: A compiler may *reallocate* stack space by reusing the same space for two or more local variables with non-overlapping lifetimes (potentially without an intervening stackpointer increment instruction). If the relative order of (de)allocations is preserved, reallocation can be encoded by annotating A with a `deallocs` instruction (for deallocating the first variable) immediately followed by an `allocs` instruction, such that the allocated region potentially overlaps with the previously deallocated region. The C fragment (on the left) and corresponding assembly code fragment (on the right) below illustrate the construction:

<pre> { int x; // alloc ... // dealloc } { int y; // alloc ... // dealloc } </pre>	<pre> L0: esp -= 4 alloc_s 'x' L1: edi = esp ... ; esp preserved dealloc_s 'x' alloc_s 'y' L2: edi = esp ; same value as L1 above ... L3: dealloc_s </pre>
--	--

The variables x and y have non-overlapping lifetimes. Common stack space for both x and y is allocated at location L0 and then later used by x between locations L1 and L2 and used by y between locations L2 and L3.

Our refinement definition may not be able to cater to a translation that changes the relative order of (de)allocation instructions during reallocation. The following C source and assembly pair demonstrate an incompleteness example¹⁸:

¹⁸See <https://godbolt.org/z/6rMT7z5re> for a GCC compilation.

<pre> int x, y; // alloc 'x'; alloc 'y' if (...) { /* use only 'x' */ } else { /* use only 'y' */ } // dealloc 'y'; dealloc 'x' </pre>	<pre> L0: esp -= 4 ; alloc stack space ... if (...) jmp L2 L1: edi = esp ; use as 'x' ... L2: edi = esp ; use as 'y' ... </pre>
--	---

Evidently, the variables `x` and `y` do not have non-overlapping lifetimes. However, their uses are limited to disjoint scopes. In the generated assembly, both share the same stack space allocated at L0. In contrast to the previous example, a valid annotation is impossible in this scenario because the same stack space cannot be allocated twice (separation WF constraint).

Dynamic allocations: $C \dot{\sqsubseteq} A$ notably supports *dynamic allocations*, a capability missing in CompCert[29] due to its preallocation strategy, which performs all allocation in the assembly procedure's prologue. Dynamic allocations, necessary for enabling variable-length arrays (VLA) and `alloca()`, and used for allocating procedure-call arguments by most production compilers, allow variable-sized local allocations (similar to `malloc()`) that are automatically deallocated at the end of scope (for variable declarations) or at the end of containing procedure (for `alloca()` allocations). Our modeling does not require any special handling of dynamic allocations. The C source and assembly pair below present an example of dynamic (de)allocation using VLA in a loop:

<pre> for (i = 1; i < n; ++i) { int v[i]; // alloc (4*i), int, // dealloc 'v' } </pre>	<pre> eax = 4 ; = 4*i (i = 1) L0: edi = esp ; save 'esp' esp = esp - eax ; variable decrement alloc, esp, (4*eax), dealloc, esp = edi ; restore 'esp' eax = eax+4 ; corresponds to i++ </pre>
---	---

In the C source, VLA `v` is declared to have `i` elements, where `i` is the for-loop index. Thus, `v` is allocated at the start of every loop body execution and deallocated at the end. In the assembly, the register `eax` holds value corresponding to `4*i`. At the start of

loop body (location L_0), current stackpointer `esp` is saved in `edi` and stack allocation of size `eax` is performed. At the end of loop body, `esp` is restored using `edi` and `eax` is incremented by 4. The annotations are made just before (after) the allocation (deallocation).

Lastly, $C \dot{\sqsubseteq} A$ admits intermittent register-allocation of (parts of) a local variable, but not complete register-allocation or elimination. We address this limitation in the next section.

2.6 Refinement in the presence of potentially register-allocated or eliminated local variables in A

If a local variable $z \in Z$ is either register-allocated or eliminated in A , there may not exist a region in stack of A that can be associated with z . However, recall that our execution model observes each allocation event in C through the `wr` instruction. Thus, for a successful refinement check, a correlated allocation event still needs to be annotated in A . It may not be possible to use an `allocs` instruction for this annotation, as `allocs` requires us to specify a region in stack and such a region may not be available¹⁹. To tackle this, we *pretend* that a correlated allocation occurs in A by introducing the notion of a *virtual allocation* instruction, called `allocv`, in A . An `allocv` instruction allocates a virtual region in A and a `deallocv` instruction deallocates it.

Figure 2.10 shows the graph translations of the virtual (de)allocation instructions `allocv` and `deallocv` and fig. 2.11 shows the revised translations of other assembly instructions to incorporate the notion of virtual allocations. We update and annotate A with the translations and instructions in figs. 2.8, 2.10 and 2.11 to obtain \ddot{A} .

2.6.1 Virtual (de)allocations through `allocv` and `deallocv` instructions

An instruction $\langle p_{\ddot{A}}^j : v := \text{alloc}_v\ e_w, a, zl \rangle$ non-deterministically chooses the start address (using choose instruction $\theta(i_{32})$) of a local variable zl of size e_w and required alignment a , performs a virtual allocation, and returns the start address in v . (`ALLOCV`)

¹⁹We use *may not* here instead of *is not* because in some situations, due to alignment requirements, it may be possible to find a large enough region in stack that is separate from other allocations but is not *used* for storing the local. We present an example of this later.

$$\begin{array}{c}
\text{(ALLOCV)} \frac{p_{\ddot{A}}^j : v := \text{alloc}_v e_w, a, zl}{\text{wr}(\text{allocBegin}(zl, e_w, a)); \\ v, w := \theta(i_{32}), e_w; \\ \text{if } (\neg \text{intrvlInSet}_a(v, v + w - 1_{i_{32}}, \text{comp}(\Sigma_{\ddot{A}}^{B \cup \{cv\}}))) \\ \text{halt}(\mathcal{W}); \\ \Sigma_{\ddot{A}}^{zl|v} := \Sigma_{\ddot{A}}^{zl|s} \cup [v]_w; \\ \text{wr}(\text{allocEnd}(zl, [v]_w, \pi_{[v]_w}(M_{\ddot{A}})))}; \\
\text{(DEALLOCV)} \frac{p_{\ddot{A}}^j : \text{dealloc}_v zl}{\text{if } (\Sigma_{\ddot{A}}^{zl|s} \neq \emptyset) \\ \text{halt}(\mathcal{U}); \\ \Sigma_{\ddot{A}}^{zl|v} := \emptyset; \\ \text{wr}(\text{dealloc}(zl))};
\end{array}$$

Figure 2.10: Translation rules for converting the alloc_v and dealloc_v instructions to graph instructions.

in fig. 2.10 shows the graph translation of alloc_v . The chosen start address v , together with interval e_w , is *assumed* to satisfy the desired WF constraints of separation (no overlap) and alignment (through $\text{intrvlInSet}_a(\dots)$); error \mathcal{W} is triggered otherwise. Notice that this is in contrast to alloc_s , where error \mathcal{U} is triggered on WF violation to indicate that it is the compiler's responsibility to ensure the satisfaction of WF constraints. Unlike a stack allocation where the compiler chooses the allocated region (and the validator identifies it through an alloc_s annotation), a virtual allocation is only a validation construct (the compiler is not involved) that is used only to enforce a lockstep correlation of allocation events. By triggering \mathcal{W} on a failure during a virtual allocation, we effectively assume that allocation through alloc_v satisfies the required WF conditions.

We put two restrictions on an alloc_v annotation to keep our automatic algorithm (for construction of a witness of refinement) simple and tractable.

1. alloc_v restricted to local variable declarations

We support virtual allocations only for a variable declaration $zl \in Z_l$. Thus, we expect a call to $\text{alloca}()$ at allocation site $za \in Z_a$ to always be stack-allocated in \ddot{A} .

In our modeling of \ddot{A} , we replace the single variable $\Sigma_{\ddot{A}}^{zl}$ for address set of zl with two variables $\Sigma_{\ddot{A}}^{zl|s}$ and $\Sigma_{\ddot{A}}^{zl|v}$ that represent the address sets corresponding to the stack-allocations and virtual-allocations due to allocation site zl respectively. We compute $\Sigma_{\ddot{A}}^{zl} = \Sigma_{\ddot{A}}^{zl|s} \cup \Sigma_{\ddot{A}}^{zl|v}$ but do not maintain a separate variable $\Sigma_{\ddot{A}}^{zl}$. For convenience, we define $\Sigma_{\ddot{A}}^{Z_l|v} = \bigcup_{zl \in Z_l} (\Sigma_{\ddot{A}}^{zl|v})$. In (ALLOCV) and (DEALLOCV), we use zl (instead of z) and a virtual (de)allocation updates the address set $\Sigma_{\ddot{A}}^{zl|v}$ (instead of $\Sigma_{\ddot{A}}^{zl}$).

2. A local variable may either be stack-allocated or virtual-allocated, not both

We do not tackle path-specializing transformations that may require, for a single variable declaration zI , a stack-allocation on one assembly path and a virtual-allocation on another. Thus, we assume that a variable declaration zI in \mathbb{C} may either correlate with only stack-allocations (through alloc_s) or only virtual-allocations (through alloc_v) in $\ddot{\mathbb{A}}$, i.e., $\Sigma_{\ddot{\mathbb{A}}}^{zI|s} \cap \Sigma_{\ddot{\mathbb{A}}}^{zI|v} = \emptyset$ holds at all times and if $\Sigma_{\ddot{\mathbb{A}}}^{zI|s} \neq \emptyset$ (resp. $\Sigma_{\ddot{\mathbb{A}}}^{zI|v}$) at any point, then $\Sigma_{\ddot{\mathbb{A}}}^{zI|v} = \emptyset$ (resp. $\Sigma_{\ddot{\mathbb{A}}}^{zI|s}$) holds throughout $\ddot{\mathbb{A}}$'s execution. This assumption simplifies the SMT encoding of proof obligations generated by our algorithm.

Note that while this restriction may appear quite constraining at first glance, it does not impose substantial practical limitations. This is because if a variable zI is *not* address-taken, then it can be virtually-allocated, as its address is never observed (this is possible even if zI is allocated on stack by the compiler, i.e., there exists a region in stack of \mathbb{A} that can be associated with zI).

We discuss the limitations arising from these restrictions in detail in section 2.6.3.

An instruction $\langle p_{\ddot{\mathbb{A}}}^j : \text{dealloc}_v \ zI \rangle$ in $\ddot{\mathbb{A}}$ empties the address set $\Sigma_{\ddot{\mathbb{A}}}^{zI|v}$ and produces an observable event through wr instruction. (DEALLOCV) in fig. 2.10 shows the graph translation of dealloc_v . An execution of dealloc_v where $\Sigma_{\ddot{\mathbb{A}}}^{zI|s}$ is non-empty triggers error \mathcal{U} , i.e., we require an error-free execution of dealloc_v to “empty” the address set $\Sigma_{\ddot{\mathbb{A}}}^{zI}$ (defined as $\Sigma_{\ddot{\mathbb{A}}}^{zI} = \Sigma_{\ddot{\mathbb{A}}}^{zI|s} \cup \Sigma_{\ddot{\mathbb{A}}}^{zI|v}$). Thus, we ensure the emptiness of $\Sigma_{\ddot{\mathbb{A}}}^{zI}$ before producing the observable trace for deallocation of zI (similar to dealloc in \mathbb{C}). Unlike dealloc_s , the deallocation of a virtually-allocated memory region, does not return the freed memory to stack stk (it instead goes back to the implicitly-defined region free).

Effectively, a lockstep correlation of virtual allocations in $\ddot{\mathbb{A}}$ with allocations in \mathbb{C} ensures that the allocation states of both procedures always agree for regions $B \cup \{cv\}$.

Figure 2.1c shows an assembly procedure with annotated alloc_v and dealloc_v instructions. Instruction $\text{A3}^1 : v_{I1} := \text{alloc}_v \ 4, 4, I1$ performs virtual-allocation of a region of size 4 identified by $I1$ and returns the start address in v_{I1} . Because alloc_v is a validator-only construct, the return address is not used anywhere in rest of the procedure. The allocated region is deallocated by instruction $\text{A19}^2 : \text{dealloc}_v \ I1$.

2.6.2 Revised semantics for assembly procedure instructions

$$\begin{array}{c}
\text{(ALLOCS')} \frac{p_{\check{A}}^j : \text{alloc}_s e_v, e_w, a, z}{\dots} \quad \text{(DEALLOCS')} \frac{p_{\check{A}}^j : \text{dealloc}_s z}{\dots} \\
\begin{array}{l}
\dots \\
\text{if } (\text{ov}([v]_w, \Sigma_{\check{A}}^{cv} \cup \Sigma_{\check{A}}^{Z_l|v})) \text{ halt } (\mathcal{W}); \\
\cancel{\Sigma_{\check{A}}^{stk}, \Sigma_{\check{A}}^z} := \Sigma_{\check{A}}^{stk} \setminus [v]_w, \Sigma_{\check{A}}^z \cup [v]_w; \\
\text{IF}\{z \in Z_l\}\{ \\
\quad \Sigma_{\check{A}}^{stk}, \Sigma_{\check{A}}^z|^s := \Sigma_{\check{A}}^{stk} \setminus [v]_w, \Sigma_{\check{A}}^z|^s \cup [v]_w; \\
\} \text{ ELSE } \{ \\
\quad \Sigma_{\check{A}}^{stk}, \Sigma_{\check{A}}^z := \Sigma_{\check{A}}^{stk} \setminus [v]_w, \Sigma_{\check{A}}^z \cup [v]_w; \\
\} \\
\dots
\end{array}
\quad
\begin{array}{l}
\cancel{\Sigma_{\check{A}}^z, \Sigma_{\check{A}}^{stk}} := \emptyset, \Sigma_{\check{A}}^{stk} \cup \Sigma_{\check{A}}^z; \\
\text{IF}\{z \in Z_l\}\{ \\
\quad \text{if } (\Sigma_{\check{A}}^z|^v \neq \emptyset) \\
\quad \quad \text{halt } (\mathcal{W}); \\
\quad \Sigma_{\check{A}}^z|^s, \Sigma_{\check{A}}^{stk} := \emptyset, \Sigma_{\check{A}}^{stk} \cup \Sigma_{\check{A}}^z|^s; \\
\} \text{ ELSE } \{ \\
\quad \Sigma_{\check{A}}^z, \Sigma_{\check{A}}^{stk} := \emptyset, \Sigma_{\check{A}}^{stk} \cup \Sigma_{\check{A}}^z; \\
\} \\
\text{wr}(\text{dealloc}(z)); \\
\dots
\end{array}
\end{array}$$

$$\begin{array}{c}
\text{(OP-ESP')} \frac{p_{\check{A}}^j : \text{esp} := \text{op}(\vec{x})}{\dots} \\
\dots \\
\text{intrvlInSet}(t, \text{esp} - 1_{i_{32}}, \Sigma_{\check{A}}^{\text{free}} \cup ((\Sigma_{\check{A}}^{cv} \cup \Sigma_{\check{A}}^{Z_l|v})) \setminus \Sigma_{\check{A}}^F) \\
\dots
\end{array}$$

$$\begin{array}{c}
\text{(ENTRY}_{\check{A}}) \frac{p_{\check{A}}^j : \text{def } \check{A}(\vec{\tau})}{\dots} \\
\dots \\
\text{(same as fig. 2.7)} \\
\dots \\
\cancel{\dots, \Sigma_{\check{A}}^z, \dots} := \dots, \emptyset, \dots; \\
\dots, \Sigma_{\check{A}}^{z^a}, \dots := \dots, \emptyset, \dots; \\
\text{for } z_l \text{ in } Z_l \{ \Sigma_{\check{A}}^{z_l|^s}, \Sigma_{\check{A}}^{z_l|^v} := \emptyset, \emptyset; \}
\end{array}$$

$$\begin{array}{c}
\text{(LOAD}_{\check{A}}) \frac{p_{\check{A}}^j : v := \text{load } w \ a \ p}{\dots} \\
\dots \\
\text{ov}([p]_w, \Sigma_{\check{A}}^{\text{free}} \cup ((\Sigma_{\check{A}}^{cv} \cup \Sigma_{\check{A}}^{Z_l|v})) \setminus \Sigma_{\check{A}}^{F \cup S}) \\
\dots
\end{array}$$

$$\begin{array}{c}
\text{(STORE}_{\check{A}}) \frac{p_{\check{A}}^j : \text{store } w \ a \ p \ v}{\dots} \\
\dots \\
\text{ov}([p]_w, \Sigma_{\check{A}}^{\{\text{free}\} \cup G_r \cup F_r} \cup ((\Sigma_{\check{A}}^{cv} \cup \Sigma_{\check{A}}^{Z_l|v})) \setminus \Sigma_{\check{A}}^{F \cup S}) \\
\dots
\end{array}$$

Figure 2.11: Revised translation rules for converting pseudo-assembly instructions to graph instructions. The $\text{IF}\{z \in Z_l\}\{\dots\}\text{ELSE}\{\dots\}$ construct selects one of the translation depending on the result of syntactic predicate $z \in Z_l$.

Figure 2.11 shows the revised semantics for the procedure entry ($\text{(ENTRY}_{\check{A}})$), alloc_s and dealloc_s instructions ((ALLOCS') and (DEALLOCS')), load and store instructions ($\text{(LOAD}_{\check{A}})$ and $\text{(STORE}_{\check{A}})$), and esp -modifying instruction ((OP-ESP')) of the assembly

procedure. Instead of reproducing the full translations, we only show the changes with appropriate context: the additions have a highlighted background and deletions are ~~striked out~~.

The execution semantics of \ddot{A} maintain the important invariant of separation of a virtual region from other common (with C) regions $B \cup \{cv\}$, thereby mirroring the execution semantics of C . However, a virtual region may potentially overlap with assembly-only regions $F \cup S$. This is because virtual-allocation is a validator-only construct, used by the validator solely for identifying a lockstep correlation of allocation states; a virtual-allocated region is never accessed otherwise in \ddot{A} . Thus, in the revised semantics of a stackpointer updating instruction, shown in (OP-ESP') of fig. 2.11, a stack push is allowed to overstep a virtually-allocated region (represented through $\Sigma_{\ddot{A}}^{Z_l|v}$ ²⁰). Notice the similarity in treatment of $\Sigma_{\ddot{A}}^{cv}$ and $\Sigma_{\ddot{A}}^{Z_l|v}$, we will expound on this in a bit.

The revised semantics of the `allocs` instruction in (ALLOCS') *assume* that the stack-allocated local memory is separate from virtually-allocated regions (and region $\Sigma_{\ddot{A}}^{cv}$), similar to separation assumption in `allocv`, error \mathcal{W} is triggered otherwise. Further, both revised rules (ALLOCS') and (DEALLOCS') now use variable $\Sigma_{\ddot{A}}^{z_l|s}$ (instead of $\Sigma_{\ddot{A}}^{z_l}$) for an address set of region $z_l \in Z_l$, reflecting the separate tracking of stack-allocated and virtual-allocated address sets of region z_l . Similarly to `deallocv`, `deallocs` triggers \mathcal{U} if $\Sigma_{\ddot{A}}^{z_l|v}$ ($z_l \in Z_l$) is non-empty, ensuring the execution of `deallocs` empties $\Sigma_{\ddot{A}}^{z_l}$ ($= \Sigma_{\ddot{A}}^{z_l|s} \cup \Sigma_{\ddot{A}}^{z_l|v}$).

The revised (ENTRY $_{\ddot{A}}$) rule initializes both address sets $\Sigma_{\ddot{A}}^{z_l|s}$ and $\Sigma_{\ddot{A}}^{z_l|v}$ for each $z_l \in Z_l$ to empty.

The revised semantics of memory access instructions (LOAD $_{\ddot{A}}$) and (STORE $_{\ddot{A}}$) enforce that a virtually-allocated region must never be accessed in \ddot{A} , unless it also happens to belong to the assembly-only regions $F \cup S$. Notice that this is similar to treatment of region cv which is similarly inaccessible in \ddot{A} .

The purpose of the cv or callers' virtual region should be clear now: cv or callers's virtual region of an assembly procedure \ddot{A} is the set of virtually-allocated addresses in \ddot{A} 's call chain. At a procedure-call, the address set $\Sigma_{\ddot{A}}^{cv}$ for a callee is computed as $\Sigma_{\ddot{A}}^{cv} \cup \Sigma_{\ddot{A}}^{Z_l|v}$. The lockstep correlation of allocation states (due to observation of (de)allocation) enables us to define Σ_C^{cv} for a callee in C using $\Sigma_{\ddot{A}}^{cv}$. As a virtual allocation is supposed to correspond to a register-allocated or an eliminated local, region cv is assumed to be

²⁰Recall that $\Sigma_{\ddot{A}}^{Z_l|v} = \bigcup_{z_l \in Z_l} (\Sigma_{\ddot{A}}^{z_l|v})$.

inaccessible in the callee²¹. This is sound because the set of observable regions for a callee constitute an observable in the caller and the equality of observables is required for establishing refinement.

2.6.3 Refinement Definition with both stack-allocated and register-allocated or eliminated locals

We define refinement in the presence of both stack-allocated and register-allocated or eliminated locals through the existence of an annotation \ddot{A} of A such that \ddot{A} refines C . \ddot{A} is obtained through addition of $(de)alloc_v$ and $(de)alloc_s$ instructions (with semantics as described in sections 2.5.1 and 2.6.1) and annotation of a procedure-call instruction (with semantics as described in section 2.5.2) and use of revised semantics for other assembly instructions in A (as described in section 2.6.2).

Definition 2.6.1 (Refinement with stack and virtually-allocated locals). $C \ddot{\sqsubseteq} A$ iff: $\exists \ddot{A} : C \sqsupseteq \ddot{A}$

Recall that $C \sqsupseteq \ddot{A}$ requires that *for all* non-deterministic choices of a virtually allocated local variable address in \ddot{A} (v in $(ALLOCV)$), there *exists* a non-deterministic choice for the correlated local variable address in C (v in $(ALLOC)$) such that: if \ddot{A} 's execution is well-formed (does not trigger \mathscr{W}), and C 's execution is UB-free (does not trigger \mathscr{U}), then the two allocated intervals are identical (the observable values created through `allocBegin` and `allocEnd` value constructors must be equal).

In the presence of potentially register-allocated and eliminated local variables, $C \ddot{\sqsubseteq} A$ implies a correct translation from C to A . If all local variables are allocated in stack, $C \ddot{\sqsubseteq} A$ reduces to $C \sqsupseteq A$ with $\ddot{A} = \dot{A}$. Figure 2.1c is an example of an annotated \ddot{A} .

Because our execution model observes each (de)allocation event (due to the `wr` instruction), a successful refinement check ensures that the allocation states of \ddot{A} and C are identical at every correlated callsite for each procedure $C \in \mathbb{C}$. A coinductive argument over \mathbb{C} and A is thus used to show that the address sets for the callers' locals — identified by `cl` and `cv` — are identical at the beginning of each correlated pair of procedures C and A , as modeled through identical reads from the outside world in $(ENTRY_P)$ ($P \in \{C, A\}$) of figs. 2.5 and 2.7. A successful refinement check for each procedure-pair

²¹For a caller local to be accessible in a callee, it should have its address taken. An address-taken local cannot be register-allocated or eliminated.

\mathbb{C}, \mathbb{A} including the `main` procedure enables a coinductive proof of refinement from \mathbb{C} to \mathbb{A} .

Capabilities and Limitations of $\mathbb{C} \dot{\sqsubseteq} \mathbb{A}$

$\mathbb{C} \dot{\sqsubseteq} \mathbb{A}$ inherits the limitation of requirement of preservation of the relative order of allocations and procedure calls that reuse stack space and the inability to handle interprocedural transformations from $\mathbb{C} \dot{\sqsubseteq} \mathbb{A}$. Notably, the example presented in fig. 2.9 is still not admitted under the new definition. However, virtual allocations enable admitting transformations involving register-allocation or elimination of a local — transformations that an optimizing compiler may and, in most cases, does perform. In Figure 2.1, the assembly procedure $\check{\mathbb{A}}_{\text{fib}}$ has register-allocated the local `i`. The `(de)allocv` annotations enable the annotated $\check{\mathbb{A}}_{\text{fib}}$ to produce identical observable events as \mathbb{C}_{fib} and establish $\mathbb{C}_{\text{fib}} \dot{\sqsubseteq} \mathbb{A}_{\text{fib}}$ ²².

Recall that we imposed two restrictions on annotations for virtual allocations (section 2.6.1): (1) `allocv` annotation may only be added for local variable declarations, (2) a local variable may exclusively be either stack-allocated or virtual-allocated. While these restrictions make the execution semantics and our automatic algorithm simpler, they preclude supporting certain transformations. In particular, (2) appears to be quite limiting. We argue below that this is not the case and that the scope of transformations prohibited by (2) is relatively limited.

Recall that we need to add `allocs` only in cases where a variable `zl` is address-taken (`allocv` can be added in rest of the cases). So, unless variable `zl` is address-taken, we can always “virtual-allocate” it. For the case where stack-allocation of `zl` through `allocs` is required (the case of an address-taken `zl`), $\mathbb{C} \dot{\sqsubseteq} \mathbb{A}$ prohibits stack space reuse transformations where the stack space reserved for `zl` is reallocated during its lifetime. Such kind of transformations may be introduced due to *live-range splitting* [11] which is utilized by optimizing compilers for improving register allocation. Thus, if a local variable is address-taken, and yet the compiler has performed live-range splitting on it, then this definition does not support it. Fortunately, such transformations on address-

²²Curiously, in this particular example, it is possible to establish $\mathbb{C}_{\text{fib}} \sqsubseteq \check{\mathbb{A}}_{\text{fib}}$ with `(de)allocs` annotations for `i`. Line `A3` allocates 12 bytes on stack (presumably for aligning `esp` by 16) which are not used in any manner later. The 4 bytes required for `i` can be “allocated” out of these 16 bytes, i.e., the annotation at `A3`¹ can be `allocs esp, 4, 4, I1`. Similarly, the annotation at `A19`² can be replaced with `deallocs I1`. The resulting annotated $\check{\mathbb{A}}_{\text{fib}}$ will satisfy $\mathbb{C}_{\text{fib}} \sqsubseteq \check{\mathbb{A}}_{\text{fib}}$, and consequently establish $\mathbb{C}_{\text{fib}} \dot{\sqsubseteq} \mathbb{A}_{\text{fib}}$.

taken variables are rare. It is also worth pointing out that Static Single Assignment (SSA) [12] is a form of live-range splitting but almost all compilers implement SSA only for pseudo-registers, which are not address-taken, and our definition supports such transformations.

We give examples of unsupported transformations below.

Refinement failure due to elimination of `alloca()`: Recall that we restrict `allocv` annotation to allocations due to a local variable declaration. If an `alloca()` in `C` is eliminated in `A`, then it may not be possible to annotate `A` with `allocs` to produce \check{A} such that $C \sqsupseteq \check{A}$ holds. The following example demonstrates such `C` and `A` where the observable behaviors of both procedures are identical but $C \not\sqsupseteq A$ does not hold.

<pre>int foo() { int *p = alloca(sizeof(int)*10); return 0; }</pre>	<pre>foo: eax = 0 ret</pre>
---	-------------------------------------

The pointer `p` returned by `alloca()` is not used in the `C` procedure (shown on the left) and hence eliminated by the compiler in the generated assembly (shown on the right). Our execution semantics forbid use of `allocv` for this eliminated local allocation and, because the assembly procedure does not allocate any stack space, an `allocs` annotation such that $C_{\text{foo}} \sqsupseteq \check{A}_{\text{foo}}$ holds cannot be made either. Thus, $C_{\text{foo}} \not\sqsupseteq A_{\text{foo}}$ cannot hold in this case.

Because the `alloca()` operator is not a part of the C standard, its use is, arguably, rare; also compilers are usually not aggressive about register-allocating the memory allocated by `alloca()`; and so, this limitation is not practically significant.

Refinement failure due to live-range splitting of an address-taken variable: Consider the C source and assembly code shown below.

<pre> void bar() { int x; // alloc scanf("*/...*/", &x); for (...) { // 'x' used } printf("*/...*/", x); } // dealloc </pre>	<pre> bar: ... esp -= 4 ; 'x' stack-allocated ... call scanf ... eax = mem4[esp]; ; 'x' register-allocated esp += 4 ; 'x' stack-deallocated ... ; for loop push eax ; value of 'x' as argument to printf ... call printf </pre>
--	---

In the C source, the variable `x` is address-taken in the call to `scanf`, then used in a loop, and, at the end, in the call to `printf`. In the assembly code, `x` is first stack-allocated so that the stack address can be passed to `scanf`. Then, before the loop, `x` is register-allocated to register `eax` and the allocated stack space is reclaimed. Lastly, `x`, present in register `eax`, is passed as argument to `printf`. Such an assembly code may be generated by a live-range splitting transformation where the live range of variable `x` is split into two variables `x1` and `x2` and `x1` is stack-allocated while `x2` is register-allocated with an assignment `eax = mem4[esp]` used to connect the two at the splitting point.

This procedure-pair cannot be admitted by our refinement definition because `x` requires an `allocs` annotation (by virtue of being address-taken) but no `(de)allocs` annotation can be made to produce \ddot{A}_{bar} such that $C_{\text{bar}} \sqsupseteq \ddot{A}_{\text{bar}}$ holds — the stack space for `x` is deallocated *before* the call to `printf` so a `deallocs` must be inserted before the call but this would break the trace requirement because `dealloc` happens *after* the call to `printf` in C_{bar} .

It is worth pointing out production compilers (Clang/LLVM, GCC) usually do not reclaim the stack space as eagerly as shown in the assembly above. Instead, the stack space for locals is reserved once in the prologue of the assembly procedure and reclaimed in the epilogue; in which case, our refinement definition will admit the transformation.

Refinement failure due to path specialization involving register-allocation on one path and stack-allocation on another for a local: Consider the C source and assembly code fragments shown below:

<pre> void baz() { int x; // alloc if (...) { // used as '&x' } else { // used as 'x' } } // dealloc </pre>	<pre> baz: ... if (...) jmp L1 L0: esp -= 4 ; allocation of 'x' ... esp += 4 ; deallocation of 'x' jmp L2 L1: ... ; no stack allocation ... ; on this branch L2: ... </pre>
---	---

On the `if` branch, the address of the local `x` is taken such that a stack allocation would be required for it. But, on the `else` branch, the address of `x` is not taken so that a register allocation may suffice. The assembly on the right implements this strategy. Because our execution semantics disallow using both `allocs` and `allocv` annotations for `x`, it is impossible to annotate the assembly procedure such that refinement holds.

2.7 Towards A More General Refinement Definition and Execution Semantics

In this section, we explore execution semantics and a corresponding refinement definition to remove the limitations described in previous section. In these new semantics we retain the observation of (de)allocation in `C` and `A`, but define revised semantics for `allocv` and replace `(de)allocs` with new instructions.

We eliminate `allocs` and `deallocs`, and instead introduce instructions `v2s` and `s2v` that do not produce non-silent trace events. We make minor changes to the semantics of `allocv` and no longer restrict `allocv` to an allocation due to local variable declaration. Figure 2.12 shows graph translations for these new instructions (`(V2S)` and `(S2V)`) and revised semantics of `allocv` (`(ALLOCV')`) and procedure entry (`(ENTRY \bar{A})`). Let \bar{A} be obtained by annotating `A` using `v2s`, `s2v`, and `(de)allocv` instructions, using revised semantics for procedure entry, and annotating a procedure-call with semantics as described in section 2.5.2.

`(ALLOCV')` shows the revised semantics for `allocv` instruction. An instruction `'v := allocv ew, a, z'` now additionally sets three new ghost variables: `lstVSz.z`, `align.z`,

$$\begin{array}{c}
\text{(ALLOCV')} \frac{p_{\ddot{A}}^j : v := \text{alloc}_v e_w, a, z}{\dots} \\
\quad \Sigma_{\ddot{A}}^z |^v := \Sigma_{\ddot{A}}^z |^v \cup [v]_w; \\
\quad \text{lstVSz.z, align.z, avail.z} := w, a, \text{true}; \\
\quad \text{wr}(\text{allocEnd}(z, [v]_w, \pi_{[v]_w}(M_{\ddot{A}}))); \\
\\
\text{(V2S)} \frac{p_{\ddot{A}}^j : v2s e_v, e_w, a, z}{v, w := e_v, e_w; \\ \text{if } (\neg \text{intrvlInSet}_a(v, v + w - 1_{i_{32}}, \Sigma_{\ddot{A}}^{stk})) \\ \quad \text{halt}(\mathcal{W}); \\ \text{if } (\neg \text{avail.z} \vee (w > \text{lstVSz.z}) \vee (a \neq \text{align.z})) \\ \quad \text{halt}(\mathcal{W}); \\ \text{if } ([v]_w \notin \Sigma_{\ddot{A}}^z |^v) \text{halt}(\mathcal{W}); \\ \Sigma_{\ddot{A}}^z |^s, \Sigma_{\ddot{A}}^z |^s \cup [v]_w, \\ \Sigma_{\ddot{A}}^z |^v, := \Sigma_{\ddot{A}}^z |^v \setminus [v]_w, ; \\ \Sigma_{\ddot{A}}^{stk}, \Sigma_{\ddot{A}}^{stk} \setminus [v]_w; \\ \text{avail.z} := \text{false};} & \text{(S2V)} \frac{p_{\ddot{A}}^j : s2v z}{\Sigma_{\ddot{A}}^z |^s, \quad \emptyset, \\ \Sigma_{\ddot{A}}^z |^v, \quad := \Sigma_{\ddot{A}}^z |^v \cup \Sigma_{\ddot{A}}^z |^s, \\ \Sigma_{\ddot{A}}^{stk} \quad \Sigma_{\ddot{A}}^{stk} \cup \Sigma_{\ddot{A}}^z |^s;} \\
\\
\text{(ENTRY}_{\ddot{A}}) \frac{p_{\ddot{A}}^j : \text{def } \ddot{A}(\vec{\tau})}{\dots} \\
\quad \text{(same as fig. 2.7)} \\
\quad \dots \\
\quad \dots, \Sigma_{\ddot{A}}^z, \dots := \dots, \emptyset, \dots, \\
\quad \text{for } z \text{ in } Z \{ \\
\quad \quad \Sigma_{\ddot{A}}^z |^s, \Sigma_{\ddot{A}}^z |^v := \emptyset, \emptyset; \\
\quad \quad \text{avail.z, lstVSz.z, align.z} := \text{false}, \theta(i_{32}), \theta(i_{32}); \\
\quad \}
\end{array}$$

Figure 2.12: Translation rules for the converting pseudo-assembly instructions to graph instructions for \ddot{A} . (ALLOCV') is derived from (ALLOCV) in fig. 2.10.

and `avail.z`. These ghost variables keep track of the parameters of the *last* virtual allocation.

- `lstVSz.z` tracks the size of last virtual allocation due to z . In (ALLOCV'), `lstVSz.z` is assigned the size w ($= e_w$) of allocation. Each execution of `allocv` thus updates `lstVSz.z` to the new value of allocation size.
- `align.z` tracks the required alignment of an allocation due to z — unlike `lstVSz.z`, the required alignment does not change across multiple executions of same `allocv` instruction. Thus, the value of `align.z`, once assigned, does not change as \ddot{A} executes.
- `avail.z` tracks the availability of virtual region for *materialization* into a stack region. An execution of an `allocv` instruction sets `avail.z` to `true` and an execution of an

`v2s` instruction resets it (shown in (V2S)).

Similar to the semantics defined in (ALLOCV), an error-free execution of `allocv` requires the allocated interval to be separate from common regions $B \cup \{cv\}$ (but overlap with assembly-only regions $F \cup S$ is permitted).

An instruction ‘`v2s e_v, e_w, a, z` ’ converts (part of) virtual allocation region due to z to a (partial) stack-allocation of z . The parameters e_v, e_w, a, z of `v2s` have identical semantics as `allocs`. The interval of size e_w , starting at address e_v , must be aligned by a and belong to stack $\Sigma_{\bar{A}}^{stk}$. Additionally, the execution of `v2s` must be sequenced after an `allocv` instruction without any intervening `v2s` instruction for the same z ($\neg \text{avail.z}$) such that the interval size is no greater than last virtual allocation size ($w > \text{lstVSz.z}$) and alignment is exactly same as last virtual allocation ($a \neq \text{align.z}$). A failure in meeting any of the above requirements triggers \mathcal{U} indicating a translation error in this stack materialization of z , i.e., the compiler failed to meet the WF constraints for this stack allocation of z .

A valid execution of `v2s` requires the allocated interval to belong to *both* $\Sigma_{\bar{A}}^{stk}$ and $\Sigma_{\bar{A}}^z|^\nu$ so that the address set of stack allocations due to z , $\Sigma_{\bar{A}}^z|^\nu$, is populated by transferring addresses from $\Sigma_{\bar{A}}^{stk}$ and $\Sigma_{\bar{A}}^z|^\nu$ to $\Sigma_{\bar{A}}^z|^\nu$ (recall that $\Sigma_{\bar{A}}^{stk}$ and $\Sigma_{\bar{A}}^z|^\nu$ are allowed to overlap). We trigger \mathcal{W} if the allocated interval does not belong to $\Sigma_{\bar{A}}^z|^\nu$, i.e, we assume that the non-deterministic choice made in `allocv` is always such that a valid stack materialization is contained inside it. We define $\Sigma_{\bar{A}}^z|^\nu = \bigcup_{z \in Z} \Sigma_{\bar{A}}^z|^\nu$ and $\Sigma_{\bar{A}}^z = \Sigma_{\bar{A}}^z|^\nu \cup \Sigma_{\bar{A}}^z|^\nu$ for a $z \in Z$ (as done in section 2.6.1 for $zl \in Z_l$). An execution of `v2s` does not change $\Sigma_{\bar{A}}^z$.

An instruction ‘`s2v z` ’ transfers *all* the stack allocations for a local z back to the virtual region and returns the corresponding addresses to stack. Unlike `deallocs`, the address set $\Sigma_{\bar{A}}^z$ is not emptied due to execution of `s2v`; on the contrary, as with `v2s`, the execution of `s2v` does not affect $\Sigma_{\bar{A}}^z$ — instead, it transfers the addresses in $\Sigma_{\bar{A}}^z|^\nu$ back to $\Sigma_{\bar{A}}^z|^\nu$ and $\Sigma_{\bar{A}}^{stk}$. Similar to `v2s`, an execution of `s2v` does not change $\Sigma_{\bar{A}}^z$. Neither `v2s` nor `s2v` produces any (non-silent) observable trace event.

(ENTRY $_{\bar{A}}$) shows the revised semantics for procedure entry: the revision is limited to initializing address sets $\Sigma_{\bar{A}}^z|^\nu$ and $\Sigma_{\bar{A}}^z|^\nu$ to empty, setting ghost variable `avail.z` to `false`, and assigning non-deterministic values to ghost variables `lstVSz.z` and `align.z`, for each $z \in Z$.

As with previous definitions, refinement in this case is also defined through the existence

of an annotated \ddot{A} , annotated with $(\text{de})\text{alloc}_v$, s2v and v2s instructions with semantics as described above and procedure-call with annotations as described in section 2.5.2, such that $C \sqsupseteq \ddot{A}$ holds.

Definition 2.7.1 (Refinement with stack and virtually-allocated locals). $C \sqsupseteq A$ iff: $\exists \ddot{A} : C \sqsupseteq \ddot{A}$

2.7.1 Comparison with $C \sqsupseteq A$

$C \sqsupseteq A$ admits strictly more procedure-pairs than $C \sqsupseteq A$

We claim that $C \sqsupseteq A$ admits at least as many procedure-pairs as $C \sqsupseteq A$.

Let \ddot{A} be obtained by annotating A as described in section 2.6. Let $\ddot{\ddot{A}}$ be obtained from \ddot{A} by the following substitution:

- Substitute an ‘ $\text{alloc}_s e_v, e_w, a, z$ ’ instruction with instruction sequence ‘ $\text{alloc}_v e_w, a, z; \text{v2s } e_v, e_w, a, z$ ’.
- Substitute a ‘ $\text{dealloc}_s z$ ’ instruction with instruction sequence ‘ $\text{s2v } z; \text{dealloc}_v z$ ’.

Let the semantics defined in fig. 2.12 be used for the instructions in $\ddot{\ddot{A}}$.

Assume machine states σ and $\dot{\sigma}$ for \ddot{A} and $\ddot{\ddot{A}}$ at nodes $n_{\ddot{A}} \in \mathcal{N}_{\ddot{A}}^{\text{DW}}$ and $n_{\ddot{\ddot{A}}} \in \mathcal{N}_{\ddot{\ddot{A}}}^{\text{DW}}$ respectively such that $n_{\ddot{A}}$ and $n_{\ddot{\ddot{A}}}$ are not error nodes. Let $=_\delta$ be a relation between σ and $\dot{\sigma}$ such that $\sigma =_\delta \dot{\sigma}$ holds iff:

- The state elements present in both σ and $\dot{\sigma}$ have identical values.
- $\Sigma_{\ddot{A}}^z = \Sigma_{\ddot{\ddot{A}}}^z$, $\Sigma_{\ddot{A}}^z|^s = \Sigma_{\ddot{\ddot{A}}}^z|^s$, and $\Sigma_{\ddot{A}}^z|^v = \Sigma_{\ddot{\ddot{A}}}^z|^v$ hold where:
 - $\Sigma_{\ddot{A}}^{zl} = \Sigma_{\ddot{\ddot{A}}}^{zl}|^s \cup \Sigma_{\ddot{\ddot{A}}}^{zl}|^v$ for $zl \in Z_l$.
 - $\Sigma_{\ddot{A}}^{za}|^s = \Sigma_{\ddot{\ddot{A}}}^{za}$ and $\Sigma_{\ddot{\ddot{A}}}^{za}|^v = \emptyset$ for $za \in Z_a$.
 - $\Sigma_{\ddot{A}}^z = \Sigma_{\ddot{\ddot{A}}}^z|^s \cup \Sigma_{\ddot{\ddot{A}}}^z|^v$ for $z \in Z$.

Lemma 2.7.2. *If $\sigma =_\delta \dot{\sigma}$ holds, and an execution of ‘ $\text{alloc}_s e_v, e_w, a, z$ ’ on σ triggers error, then an execution of instructions (in sequence) ‘ $\text{alloc}_v e_w, a, z; \text{v2s } e_v, e_w, a, z$ ’ on $\dot{\sigma}$ will also trigger an error (for some sequence of non-deterministic choices).*

Further, if the erroneous execution on σ does not trigger error \mathcal{U} , then the erroneous execution on $\dot{\sigma}$ also does not trigger error \mathcal{U} .

Proof. An execution of alloc_s on σ may trigger either error \mathcal{U} due to evaluation of

$\neg\text{intrvlInSet}(\dots)$ check ((ALLOCS)) to true, call it case 1. Or, it may terminate with error \mathcal{W} due to evaluation of $\text{ov}(\dots)$ check ((ALLOCS')) to true, call it case 2.

1. *When error \mathcal{U} is triggered:* The $\neg\text{intrvlInSet}_a(v, v + w - 1_{i_{32}}, \Sigma_{\ddot{A}}^{stk})$ check due to alloc_s ((ALLOCS)) is structurally identical to the \mathcal{U} -triggering check due to v2s ((V2S)). Further, the arguments to both instructions alloc_s and v2s above are identical and due to $\sigma =_{\delta} \dot{\sigma}$, $\Sigma_{\ddot{A}}^{stk} = \Sigma_{\ddot{A}}^{stk}$ holds. Also, the execution of alloc_v does not affect $\Sigma_{\ddot{A}}^{stk}$. Therefore, if execution of alloc_v completes without triggering error \mathcal{W} , then execution of v2s will terminate with error \mathcal{U} due to evaluation of $\neg\text{intrvlInSet}_a(\dots)$ check to false.
2. *When error \mathcal{W} is triggered:* An execution of instruction alloc_v on $\dot{\sigma}$ may trigger \mathcal{W} due to the non-deterministic choice for v in (ALLOCV') — to distinguish the v in (ALLOCV') from the v in (V2S), we will refer to the former as v_v and the latter as simply v ²³.

If the execution of alloc_s on σ triggers \mathcal{W} due to evaluation of $\text{ov}([v]_w, \Sigma_{\ddot{A}}^{cv} \cup \Sigma_{\ddot{A}}^Z |^v)$ check to false, then, for some sequence of non-deterministic choices, execution of alloc_v ; v2s on $\dot{\sigma}$ may also trigger \mathcal{W} . We prove this by contradiction. Assume that the execution of alloc_v on $\dot{\sigma}$ does not trigger \mathcal{W} , then consider the subsequent execution of v2s on resulting $\dot{\sigma}$. The evaluation of $\neg\text{intrvlInSet}(\dots)$ on $\dot{\sigma}$ will not result in triggering \mathcal{U} as it did not trigger \mathcal{U} in execution of σ (same reasoning as case 1 above). If the execution of v2s on $\dot{\sigma}$ does not trigger \mathcal{W} due to evaluation of $([v]_w \not\subseteq \Sigma_{\ddot{A}}^z)$ to false, then we have $[v]_w \subseteq \Sigma_{\ddot{A}}^z$ or $[v]_w \subseteq \Sigma_{\ddot{A}}^z |^s \cup \Sigma_{\ddot{A}}^z |^v$ at the end of alloc_v 's execution (recall that v refers to the v in (V2S) here which holds identical value as v at the end of alloc_s 's execution). Due to $\sigma =_{\delta} \dot{\sigma}$ and the fact that execution of alloc_s did not change $\Sigma_{\ddot{A}}^z |^s$ (due to early termination), we can conclude $[v]_w \subseteq (\Sigma_{\ddot{A}}^z |^v \cup [v_v]_w) \cup \Sigma_{\ddot{A}}^z |^s$ (notice the change from $\Sigma_{\ddot{A}}^z |^v$ to $\Sigma_{\ddot{A}}^z |^v$). As execution of v2s did not terminate with error \mathcal{U} , $[v]_w \subseteq \Sigma_{\ddot{A}}^{stk}$ must hold prior to execution of v2s and, therefore, $[v]_w \not\subseteq \Sigma_{\ddot{A}}^z |^s$ must hold at the same point²⁴. Consequently, $[v]_w \subseteq \Sigma_{\ddot{A}}^z |^v \cup [v_v]_w$. If execution on σ triggered \mathcal{W} , then $\text{ov}([v]_w, \Sigma_{\ddot{A}}^{cv} \cup \Sigma_{\ddot{A}}^Z |^v)$ must hold. As v in v2s has identical value as v in alloc_s , $\text{ov}(\Sigma_{\ddot{A}}^z |^v \cup [v_v]_w, \Sigma_{\ddot{A}}^{cv} \cup \Sigma_{\ddot{A}}^Z |^v)$ must hold. $\neg\text{ov}([v_v]_w, \Sigma_{\ddot{A}}^{cv} \cup \Sigma_{\ddot{A}}^Z |^v)$ holds for error-free execution of alloc_v and $\neg\text{ov}(\Sigma_{\ddot{A}}^z |^v, \Sigma_{\ddot{A}}^{cv})$ holds for σ at error-free node, yielding $\text{ov}(\Sigma_{\ddot{A}}^z |^v, \Sigma_{\ddot{A}}^Z |^v)$. Recall that \ddot{A} restricts alloc_v to local variable declaration z and for a local z , \ddot{A} may either have

²³We do not similarly distinguish between w because, unlike v , any execution over $\dot{\sigma}$ will have identical values for both w .

²⁴This is due to the semantics of v2s and s2v , which prohibit an address α from simultaneously satisfying $\alpha \in \Sigma_{\ddot{A}}^z |^s$ and $\alpha \in \Sigma_{\ddot{A}}^{stk}$.

alloc_s or alloc_v annotation but not both. This gives us $\Sigma_{\bar{A}}^z|^\nu = \emptyset$, resulting in the contradiction $\text{ov}(\emptyset, \Sigma_{\bar{A}}^z|^\nu)$. Thus, either of our assumption of error-free execution of alloc_v or v2s was incorrect and execution of $\hat{\sigma}$ must also terminate with error \mathcal{W} .

Due to case 2, it follows that the erroneous execution on $\hat{\sigma}$ will trigger \mathcal{W} (and not \mathcal{U}) if the erroneous execution on σ triggers \mathcal{W} . \square

Lemma 2.7.3. *If $\sigma =_\delta \hat{\sigma}$ holds, and an execution of ‘ $\text{alloc}_s e_v, e_w, a, z$ ’ on σ completes without triggering error, then there exists a sequence of non-deterministic choices such that an execution of instructions (in sequence) ‘ $\text{alloc}_v e_w, a, z; \text{v2s } e_v, e_w, a, z$ ’ on $\hat{\sigma}$ also completes without triggering error.*

Proof. If the execution on σ completed without triggering error, then neither of the conditions $\neg\text{intrvlInSet}(\dots)$ and $\text{ov}(\dots)$ in (ALLOCS’) evaluated to **true** during σ ’s execution. Thus, both $\neg\text{ov}([v]_w, \Sigma_{\bar{A}}^{cv} \cup \Sigma_{\bar{A}}^z|^\nu)$ and $\neg\text{ov}([v]_w, \Sigma_{\bar{A}}^B \setminus (\Sigma_{\bar{A}}^{cv} \cup \Sigma_{\bar{A}}^z|^\nu))$ ²⁵ should hold on σ for $v = e_v$ and $w = e_w$ before execution of alloc_s begins, yielding $\neg\text{ov}([v]_w, \Sigma_{\bar{A}}^{B \cup \{cv\}})$ or $\neg\text{ov}([v]_w, \Sigma_{\bar{A}}^{B \cup \{cv\}})$ (due to $\sigma =_\delta \hat{\sigma}$). During the execution on $\hat{\sigma}$, choose v in (ALLOCV’) to be identical to v above. It can be observed that the execution on $\hat{\sigma}$ will complete without triggering error. \square

Lemma 2.7.4. *If $\sigma =_\delta \hat{\sigma}$ holds, then after an error-free execution of instruction ‘ $\text{alloc}_s e_v, e_w, a, z$ ’ on σ to obtain σ' and an error-free execution of instructions (in sequence) ‘ $\text{alloc}_v e_w, a, z; \text{v2s } e_v, e_w, a, z$ ’ on $\hat{\sigma}$ to obtain $\hat{\sigma}'$, $\sigma' =_\delta \hat{\sigma}'$ holds.*

As $\Omega_{\bar{A}} \in \sigma'$ and $\Omega_{\bar{A}} \in \hat{\sigma}'$, the executions must produce identical traces as well.

Proof. An error-free execution of alloc_s on σ will mutate the address sets: $\Sigma_{\bar{A}}^{stk}$ to $\Sigma_{\bar{A}}^{stk} \setminus [v]_w$ and either of $\Sigma_{\bar{A}}^z$ to $\Sigma_{\bar{A}}^z \cup [v]_w$ (if $z \in Z_a$) or $\Sigma_{\bar{A}}^z|^\nu$ to $\Sigma_{\bar{A}}^z|^\nu \cup [v]_w$ (if $z \in Z_l$). Similarly, an error-free execution of $\text{alloc}_v; \text{v2s}$ in sequence will mutate the address sets: $\Sigma_{\bar{A}}^{stk}$ to $\Sigma_{\bar{A}}^{stk} \setminus [v_s]_w$, $\Sigma_{\bar{A}}^z|^\nu$ to $\Sigma_{\bar{A}}^z|^\nu \cup [v_s]_w$, and $\Sigma_{\bar{A}}^z|^\nu$ to $(\Sigma_{\bar{A}}^z|^\nu \cup [v_v]_w) \setminus [v_s]_w$, where v_s refers to the v in (V2S) (which is identical to the v in alloc_s above) and v_v refers to the v in (ALLOCV’). Choosing v_v (of (ALLOCV’)) to be identical to v_s , it can be observed that $\sigma' =_\delta \hat{\sigma}'$ will hold. Proof of lemma 2.7.3 shows that this choice is always feasible for an error-free execution. \square

²⁵Due to semantics of (OP-ESP’), (ALLOCS’) and (DEALLOCS’), and (ALLOCV) and (DEALLOCV), because of which $\neg\text{ov}(\Sigma_{\bar{A}}^{stk}, \Sigma_{\bar{A}}^B \setminus (\Sigma_{\bar{A}}^{cv} \cup \Sigma_{\bar{A}}^z|^\nu))$ holds.

It is easy to state and prove similar claims as lemmas 2.7.2 to 2.7.4 for ‘`deallocs z`’ in \ddot{A} and instruction sequence ‘`s2v z; deallocv z`’ in \ddot{A} . Similarly, for instruction ‘`allocv`’ (resp. `deallocv`) in \ddot{A} and instruction ‘`allocv`’ (resp. `deallocv`) in \ddot{A} , where the semantics are identical (barring the ghost variables `lstVSz.z`, `align.z`, and `avail.z`).

Theorem 2.7.5. *If there exists an annotation \ddot{A} such that $C \sqsupseteq \ddot{A}$ holds, then it is possible to construct an annotation \ddot{A} such that $C \sqsupseteq \ddot{A}$ holds.*

Proof. Construct \ddot{A} as described above by replacing the `allocs` and `deallocs`, instructions in \ddot{A} with instruction sequences ‘`allocv; v2s`’ and ‘`s2v; deallocv`’ respectively.

The proof follows from induction on lockstep execution of \ddot{A} and \ddot{A} and lemmas 2.7.2 to 2.7.4. □

Recall that $C \sqsupseteq A$ does not admit certain procedure-pairs that have identical observable behavior according to the C standard (section 2.6.3). We demonstrate below that those procedure-pairs will be admitted under $C \sqsupseteq A$.

- Elimination of `alloca()`:

<pre>int foo() { int *p = alloca(sizeof(int)*10); return 0; }</pre>	<pre>foo: alloc_v 40, ... eax = 0 dealloc_v ret</pre>
---	---

Without the restriction on placement of `allocv`, it is easy to add the required (de)`allocv` annotations (shown in red) such that $C_{\text{foo}} \sqsupseteq \ddot{A}_{\text{foo}}$ holds.

- Live-range splitting of variable so that it is both stack-allocated and register-allocated in same path:

```

void bar()
{
  int x; // alloc
  scanf(/*...*/, &x);
  for (...) {
    // 'x' used
  }
  printf(/*...*/, x);
} // dealloc

```

```

bar:
  allocv 4, ...
  esp -= 4 ; 'x' stack-allocated
  v2s esp, 4, ...
  ... call scanf ...
  eax = mem4[esp]; ; 'x' register-allocated
  s2v ...
  esp += 4 ; 'x' stack-deallocated
  ... ; for loop
  push eax ; value of 'x' as argument to printf
  ... call printf ...
  ...
  deallocv

```

The $v2s$ and $s2v$ annotations in \ddot{A}_{bar} trace the stack-allocation and register-allocation performed by the compiler. As neither $v2s$ nor $s2v$ produce any observable trace event that needs to be correlated with C_{bar} , an identical behavior in the \dots fragments in both C_{bar} and A_{bar} with annotations as shown in red, $C_{\text{bar}} \sqsupseteq \ddot{A}_{\text{bar}}$ will hold.

- Path specialization involving register-allocation on one path and stack-allocation on another:

```

void baz()
{
  int x; // alloc
  if (...) {
    // used as 'ℓx'
  } else {
    // used as 'x'
  }
} // dealloc

```

```

baz:
  ...
  allocv 4, ...
  if (...) jmp L1
L0: esp -= 4 ; true branch
  v2s esp, 4, ...
  ...
  s2v ...
  esp += 4
  jmp L2
L1: ... ; false branch: no stack
  ... ; allocation
L2: ... ; meet point
  deallocv

```

The use of $v2s$ and $s2v$ annotations in \ddot{A}_{baz} enable treatment of a stack region as

belonging to a previously “allocated” (using alloc_v) local without producing any observable trace event that needs to be correlated with C_{baz} . Assuming identical behaviors in the \dots fragments in both C_{baz} and A_{baz} and annotations as shown in red, $C_{\text{baz}} \sqsupseteq \ddot{A}_{\text{baz}}$ will hold.

Thus, $C \sqsupseteq \ddot{A}$ admits strictly more procedure-pairs than $C \sqsupseteq A$.

Practicality of $C \sqsupseteq \ddot{A}$

While $C \sqsupseteq \ddot{A}$ is strictly more *powerful* than $C \sqsupseteq A$, the restrictions on $C \sqsupseteq \ddot{A}$ make it more amenable to a *simpler* algorithmic construction such that $C \sqsupseteq \ddot{A}$ can be witnessed (our construction is detailed in chapters 3 and 4). The two restrictions defined over annotation of alloc_v (section 2.6.1) are consequential in realizing an efficient SMT encoding of verification conditions over C and \ddot{A} . In particular, the restrictions on limiting annotation of alloc_v to a variable declaration enable use of an address interval (defined by ghost variables $\boxed{\text{lb.z}}$ and $\boxed{\text{ub.z}}$ in C) for tracking $\Sigma_{\ddot{A}}^{z|v}$. Use of such *interval encoding* results in measurable performance improvements in runtime of our algorithm (section 6.2). The restriction on exclusivity (using exclusively either an annotation of alloc_s or an annotation alloc_v) enables a rewrite of the predicate $(\alpha \in \Sigma_{\ddot{A}}^{z|})$ to either $(\alpha \in \Sigma_{\ddot{A}}^{z|s})$ or $(\alpha \in \Sigma_{\ddot{A}}^{z|v})$, depending on whether an alloc_s or an alloc_v annotation is used respectively. $(\alpha \in \Sigma_{\ddot{A}}^{z|v})$ can be encoded (in SMT) using simple (bitvector) comparisons by taking advantage of the first restriction. We leave the exploration of similar construction for $C \sqsupseteq \ddot{A}$ for future work.

Chapter 3

Witnessing Refinement through a Determinized Cross-Product

In the previous chapter, we defined refinement between a source procedure C and an annotated assembly procedure \ddot{A} as a relation over their traces. In this chapter, we present our proof method for witnessing refinement. Our proof method involves constructing a cross-product or product program that puts C and \ddot{A} in lockstep. While product program construction is a well-known technique for establishing bisimulation, we propose a *determinized product program*, an extension that accommodates non-determinism and thus can be used for witnessing refinement. We describe a set of requirements over a determinized product program between C and \ddot{A} such that the existence of a program meeting these requirements implies the refinement relation $C \sqsupseteq \ddot{A}$.

The chapter is organized as follows: we start by defining program paths in section 3.1; in section 3.2, we define the product program as a determinized product graph and in section 3.3, we state the requirements over the determinized product graph that enable it to witness refinement. In sections 3.4 and 3.5, we describe new *callers' virtual smallest* and *safety-relaxed* semantics for C and A that enable a more efficient SMT encoding of verification conditions.

3.1 Program Paths

Let $P \in \{\mathbf{C}, \mathbf{\ddot{A}}\}$. Let $e_P = (n_P \rightarrow n_P^t) \in \mathcal{E}_P$ represent an edge from node n_P to node n_P^t , both drawn from nodes \mathcal{N}_P of P . A *path* ξ_P from n_P to n_P^t , written $\xi_P = (n_P \twoheadrightarrow n_P^t)$, is a finite sequence of $m \geq 0$ edges $(e_P^1, e_P^2, \dots, e_P^m)$ with $\forall_{1 \leq j \leq m} : e_P^j = (n_P^{f,j} \rightarrow n_P^{t,j}) \in \mathcal{E}_P$, such that $n_P^{f,1} = n_P$, $n_P^{t,m} = n_P^t$, and $\bigwedge_{j=1}^{m-1} (n_P^{t,j} = n_P^{f,j+1})$. An empty sequence ($m = 0$), written ϵ , represents the *empty path*. Nodes n_P and n_P^t are called the *source* and *sink* nodes of ξ_P respectively. ξ_P is said to originate at n_P and end at n_P^t . Edge e_P^j (for some $1 \leq j \leq m$) is said to be present in ξ_P , written $e_P^j \in \xi_P$.

A path $\xi_P^x = (e_P^{x_1}, e_P^{x_2}, \dots, e_P^{x_m})$ ($m \geq 0$) is a prefix of path $\xi_P^y = (e_P^{y_1}, e_P^{y_2}, \dots, e_P^{y_n})$ ($n \geq 0$), written $\xi_P^x \preceq \xi_P^y$, iff $m \leq n$ and $\bigwedge_{i=1}^m (e_P^{x_i} = e_P^{y_i})$.

Definition 3.1.1 (Mutually exclusive paths). *Two paths, $\xi_P^1 = (n_P \twoheadrightarrow n_P^{t_1})$ and $\xi_P^2 = (n_P \twoheadrightarrow n_P^{t_2})$, both originating at node n_P are **mutually-exclusive**, written $\xi_P^1 \approx \xi_P^2$, iff neither is a prefix of the other, i.e., both $\neg(\xi_P^1 \preceq \xi_P^2)$ and $\neg(\xi_P^2 \preceq \xi_P^1)$ hold.*

Definition 3.1.2 (Pathset). *A **pathset** $\langle \xi \rangle_P$ is a set of pairwise mutually-exclusive paths $\langle \xi \rangle_P = \{\xi_P^1, \xi_P^2, \dots, \xi_P^m\}$ originating at the same node n_P , i.e., $\forall_{1 \leq j \leq m} : \xi_P^j = (n_P \twoheadrightarrow n_P^j)$ and $\forall_{1 \leq j_1 < j_2 \leq m} : (\xi_P^{j_1} \approx \xi_P^{j_2})$.*

The execution of a path $\xi_P = (n_P \twoheadrightarrow n_P^t) = (e_P^1, e_P^2, \dots, e_P^m)$ ($m \geq 0$) is the sequential execution of edges $e_P^1, e_P^2, \dots, e_P^m$ starting at node n_P . The *path condition* of a path ξ_P , written $\text{pathcond}(\xi_P)$, is a conjunction of the edge conditions of the constituent edges. Starting at n_P , $\text{pathcond}(\xi_P)$ represents the condition that ξ_P executes to completion.

Definition 3.1.3 (I/O path). *A sequence of edges corresponding to a shaded statement in the translations (figs. 2.4 to 2.8 and 2.11) is distinguished and identified as an **I/O path**. An I/O path must contain either a single **rd** or a single **wr** instruction.*

A **rd/wr** instruction is always part of an I/O path. The sequence of edges corresponding to “ $\text{wr}(\text{fcall}(\rho, \vec{x}, \beta^*, \pi_{\Sigma_C}^{\beta^*}(M_C)))$ ” and “ $M_C := \text{upd}_{\Sigma_C^{\beta^*} \setminus G_r}(M_C, \text{rd}(i_{32} \rightarrow i_8))$ ” in (CALL_C) of fig. 2.5 refer to two separate I/O paths. A path without any **rd** or **wr** instructions is called an *I/O-free path*.

3.2 Determinized Product Graph as a Transition Graph

A product program, represented as a *determinized product graph*, also called a comparison graph or a cross-product, $X = \ddot{A} \boxtimes C = (\mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X)$, is a directed multigraph with finite sets of nodes \mathcal{N}_X and edges \mathcal{E}_X and a *deterministic choice map* \mathcal{D}_X . X is used to encode a lockstep execution of \ddot{A} and C , such that $\mathcal{N}_X \subseteq \mathcal{N}_{\ddot{A}} \times \mathcal{N}_C$. The start node of X is $n_X^s = (n_{\ddot{A}}^s, n_C^s)$ and all nodes in \mathcal{N}_X must be reachable from n_X^s . A node $n_X = (n_{\ddot{A}}, n_C)$ is an error node iff either $n_{\ddot{A}}$ or n_C is an error node¹. $\mathcal{N}_X^{\text{DW}}$ denotes the set of error-free nodes in X , such that $n_X = (n_{\ddot{A}}, n_C) \in \mathcal{N}_X^{\text{DW}} \Leftrightarrow (n_{\ddot{A}} \in \mathcal{N}_{\ddot{A}}^{\text{DW}} \wedge n_C \in \mathcal{N}_C^{\text{DW}})$. A node $n_X = (n_{\ddot{A}}, n_C)$ is a terminating node iff either of $n_{\ddot{A}}$ or n_C is a terminating node. An error node is always a terminating node.

Let $n_X = (n_{\ddot{A}}, n_C)$ and $n_X^t = (n_{\ddot{A}}^t, n_C^t)$ be nodes in \mathcal{N}_X . Let $\xi_{\ddot{A}} = (n_{\ddot{A}} \rightarrow n_{\ddot{A}}^t)$ be a path in \ddot{A} and let $\xi_C = (n_C \rightarrow n_C^t)$ be a path in C . Each edge, $e_X = (n_X \xrightarrow{\xi_{\ddot{A}}; \xi_C} n_X^t) \in \mathcal{E}_X$, is defined as a sequential execution of $\xi_{\ddot{A}}$ followed by ξ_C . The execution of e_X transfers control of X from n_X to n_X^t .

The machine state σ_X of X is the concatenation of the machine states of \ddot{A} and C . The outside world of X , written Ω_X , is a pair of the outside worlds of \ddot{A} and C , i.e., $\Omega_X = (\Omega_{\ddot{A}}, \Omega_C)$. Similarly, the trace generated by X , written T_X , is a pair of the traces generated by \ddot{A} and C , i.e., $T_X = (T_{\ddot{A}}, T_C)$.

Let $n_X = (n_{\ddot{A}}, n_C) \in \mathcal{N}_X$ and $n_X^t = (n_{\ddot{A}}^t, n_C^t) \in \mathcal{N}_X$. During an execution of $e_X = (n_X \xrightarrow{\xi_{\ddot{A}}; \xi_C} n_X^t) \in \mathcal{E}_X$, let $\vec{x}_{\ddot{A}}$ be variables in \ddot{A} just at the end of the execution of path $\xi_{\ddot{A}}$ (at $n_{\ddot{A}}^t$) but before the execution of path ξ_C (recall, $\xi_{\ddot{A}}$ executes before ξ_C). $\mathcal{D}_X : (\mathcal{E}_X \times \mathcal{E}_C \times \mathbb{N}) \rightarrow \text{ExprList}$, called a *deterministic choice map*, is a partial function that maps edge $e_X \in \mathcal{E}_X$, and the n^{th} (for $n \in \mathbb{N}$) occurrence of an edge $e_C^\theta \in \xi_C$ labeled with instruction ‘ $\vec{v} := \theta(\vec{\tau})$ ’ to a list of expressions $E(\vec{x}_{\ddot{A}})$. The semantics of \mathcal{D}_X are such that, if $\mathcal{D}_X(e_X, e_C^\theta, n)$ is defined, then during an execution of e_X , an execution of the n^{th} occurrence of edge $e_C^\theta \in \xi_C$ labeled with ‘ $\vec{v} := \theta(\vec{\tau})$ ’ is semantically equivalent to an execution of ‘ $\vec{v} := \mathcal{D}_X(e_X, e_C^\theta, n)$ ’; otherwise, the original semantics of θ are used.

\mathcal{D}_X determinizes (or refines) the non-deterministic choices in C . For example, in a product graph X that correlates the programs in fig. 2.1b and fig. 2.1c, let $e_X^2 \in \mathcal{E}_X$ correlate

¹Recall that there are two error nodes in $P \in \{C, \ddot{A}\}$: \mathcal{U}_P and \mathcal{W}_P (section 2.2.7).

single instructions I2 and A5¹ (corresponding to `alloc` and `allocs` respectively). Let e_C^{I2, θ_a} represent the edge labeled with ‘ $\alpha_b := \theta(i_{32})$ ’ as a part of the translation of the `alloc` instruction at I2, as seen in (ALLOC) (fig. 2.5). Then, $\mathcal{D}_X(e_X^2, e_C^{I2, \theta_a}, 1) = (\text{esp})$ is identified by the first operand of the annotated `allocs` instruction at A5¹. Similarly, if another edge e_C^{I2, θ_m} (in the translation of `alloc` at I2) is labeled with $\theta(i_{32} \rightarrow i_8)$ (due to ‘ $M_C := \text{upd}_{[\alpha_b, \alpha_e]}(M_C, \theta(i_{32} \rightarrow i_8))$ ’), then $\mathcal{D}_X(e_X^2, e_C^{I2, \theta_m}, 1) = (M_{\check{A}})$, i.e., the initial contents of the newly-allocated region in \check{C} are based on the contents of the correlated uninitialized stack region in \check{A} . Similarly, let $e_X^1 \in \mathcal{E}_X$ correlate single instructions I1 and A3¹ so that $\mathcal{D}_X(e_X^1, e_C^{I1, \theta_a}, 1) = (v_{I1})$ and $\mathcal{D}_X(e_X^1, e_C^{I1, \theta_m}, 1) = (M_{\check{A}})$.

Definition 3.2.1 (Determinized Path). *For a path ξ_C in \check{C} , edge $e_X = (n_X \xrightarrow{\xi_{\check{A}}; \xi_C} n_X^t) \in \mathcal{E}_X$ and deterministic choice map \mathcal{D}_X , $[\xi_C]_{\mathcal{D}_X}^{ex}$ denotes a **determinized path** that is identical to ξ_C except that: if $\mathcal{D}_X(e_X, e_C^\theta, n)$ is defined, then the n^{th} occurrence of edge $e_C^\theta \in \xi_C$, labeled with ‘ $\vec{v} := \theta(\vec{\tau})$ ’, is replaced with a new edge $e_C^{\theta'}$ labeled with ‘ $\vec{v} := \mathcal{D}_X(e_X, e_C^\theta, n)$ ’.*

Execution of a product graph X must begin at node n_X^s in an initial machine state where $\Omega_{\check{A}} = \Omega_C$ and $T_{\check{A}} =_{st} T_C$ hold. Execution of an edge $e_X = (n_X \xrightarrow{\xi_{\check{A}}; \xi_C} n_X^t) \in \mathcal{E}_X$ is execution of $\xi_{\check{A}}$ followed by execution of potentially determinized (using \mathcal{D}_X) ξ_C . Thus, X is a transition graph with its execution semantics derived from the semantics of \check{A} and \check{C} , and the map \mathcal{D}_X .

3.3 Analysis of the determinized product graph

Let $X = \check{A} \boxtimes \check{C} = (\mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X)$ be a determinized product graph. At each error-free node $n_X \in \mathcal{N}_X^{\text{DFW}}$, we infer an inductively-provable node invariant ϕ_{n_X} which is a first-order logic predicate over state elements of X at node n_X that holds for all possible executions of X . A node invariant ϕ_{n_X} relates the values of state elements of \check{C} and \check{A} that can be observed at n_X . Let Φ_X be an *invariant network* that maps n_X to its node invariant ϕ_{n_X} so that $\Phi_X(n_X) = \phi_{n_X}$.

A path $\xi_{\check{A}}$ (similarly, ξ_C) is said to originate at node $n_X = (n_{\check{A}}, n_C) \in \mathcal{N}_X$ iff $\xi_{\check{A}}$ (ξ_C) originates at $n_{\check{A}}$ (n_C).

Definition 3.3.1 (Hoare triple). *Let $n_X = (n_{\check{A}}, n_C) \in \mathcal{N}_X^{\text{DFW}}$. Let $\xi_{\check{A}} = (n_{\check{A}} \rightarrow n_{\check{A}}^t)$ and $\xi_C = (n_C \rightarrow n_C^t)$ be paths in \check{A} and \check{C} . A **Hoare triple**, written $\{pre\}(\xi_{\check{A}}; \xi_C)\{post\}$, denotes the statement: if execution starts at node n_X in state σ such that predicate $pre(\sigma)$ holds, and if paths $\xi_{\check{A}}; \xi_C$ are executed in sequence to completion finishing in*

state σ' , then predicate $\text{post}(\sigma')$ holds.

We define *path infeasibility* and *path cover* in terms of Hoare triple(s).

Definition 3.3.2 (Path infeasibility). *At a node $n_X = (n_{\bar{A}}, n_C) \in \mathcal{N}_X^{\text{DW}}$, a path $\xi_{\bar{A}} = (n_{\bar{A}} \rightarrow n_{\bar{A}}^t)$ is **infeasible** at n_X iff the Hoare triple $\{\phi_{n_X}\}(\xi_{\bar{A}}; \epsilon)\{\text{false}\}$ holds.*

Definition 3.3.3 (Path cover). *At a node $n_X = (n_{\bar{A}}, n_C) \in \mathcal{N}_X^{\text{DW}}$, for a path $\xi_{\bar{A}} = (n_{\bar{A}} \rightarrow n_{\bar{A}}^t)$, let $\forall 1 \leq j \leq m : e_X^j = (n_X \xrightarrow{\xi_{\bar{A}}; \xi_C^j} n_X^{t_j})$ be all edges in \mathcal{E}_X , such that $n_X^{t_j} = (n_{\bar{A}}^t, n_C^{t_j})$. The set of edges $\{e_X^1, e_X^2, \dots, e_X^m\}$ **covers path** $\xi_{\bar{A}}$, written $\{e_X^1, e_X^2, \dots, e_X^m\} \langle \mathcal{D}_X, \xi_{\bar{A}} \rangle$, iff the Hoare triple $\{\phi_{n_X}\}(\xi_{\bar{A}}; \epsilon) \{ \bigvee_{j=1}^m \text{pathcond}([\xi_C^j]_{\mathcal{D}_X}^{e_X^j}) \}$ holds.*

In other words, the set of edges $\{e_X^1, \dots, e_X^j, \dots, e_X^m\}$ covers the path $\xi_{\bar{A}}$ iff: whenever an execution starting at n_X in state σ such that $\phi_{n_X}(\sigma)$ holds completes the execution of $\xi_{\bar{A}}$, a subsequent execution starting at n_C must execute at least one of the determinized path $[\xi_C^j]_{\mathcal{D}_X}^{e_X^j}$ ($1 \leq j \leq m$) to completion. Thus, if path cover at n_X for path $\xi_{\bar{A}}$ holds, then at least one of the outgoing edges at n_X will execute to completion.

3.3.1 X requirements

Let $_$ be a wildcard character for a node. We define the following requirements on X so it may witness $C \sqsupseteq \bar{A}$:

1. (Mutex \bar{A}): For each node $n_X \in \mathcal{N}_X$, with *all* outgoing edges $\{e_X^1, e_X^2, \dots, e_X^m\}$ such that $e_X^j = (n_X \xrightarrow{\xi_{\bar{A}}; \xi_C^j} n_X^{t_j}) \in \mathcal{E}_X$ and $n_X^{t_j} \in \mathcal{N}_X$ (for $1 \leq j \leq m$), the following holds:

$$\forall 1 \leq j_1, j_2 \leq m : (\xi_{\bar{A}}^{j_1} = \xi_{\bar{A}}^{j_2}) \vee (\xi_{\bar{A}}^{j_1} \approx \xi_{\bar{A}}^{j_2})$$

In other words, two \bar{A} paths $\xi_{\bar{A}}^{j_1}, \xi_{\bar{A}}^{j_2}$ ($1 \leq j_1, j_2 \leq m$) originating at node n_X are either identical or mutually exclusive.

2. (Mutex C): At each node $n_X = (n_{\bar{A}}, n_C) \in \mathcal{N}_X$, for some path $\xi_{\bar{A}} = (n_{\bar{A}} \rightarrow n_{\bar{A}}^t)$, let $\{e_X^1, e_X^2, \dots, e_X^m\}$ be a set of *all* outgoing edges such that $e_X^j = (n_X \xrightarrow{\xi_{\bar{A}}; \xi_C^j} n_X^{t_j}) \in \mathcal{E}_X$ (for $1 \leq j \leq m$) and $n_X^{t_j} = (n_{\bar{A}}^t, n_C^{t_j}) \in \mathcal{N}_X$. Then, the set $\{\xi_C^1, \xi_C^2, \dots, \xi_C^m\}$ must be a pathset, i.e., the paths are pairwise mutually-exclusive.

In other words, the set of paths correlated with a path $\xi_{\bar{A}}$ originating at n_X are pairwise mutually exclusive. Because mutually-exclusive paths cannot be executed simultaneously (by virtue of their pairwise complementary path conditions), together

(Mutex \ddot{A}) and (Mutex C) make execution of X *deterministic* such that at a node n_X , at most one outgoing edge e_X^j may be executed to completion².

3. (Termination) For each error-free node $n_X = (n_{\ddot{A}}, n_C) \in \mathcal{N}_X^{\text{UW}}$, $n_{\ddot{A}}$ is a terminating node iff n_C is a terminating node.

(Termination) ensures error-free termination of both executions (C and \ddot{A}) happens simultaneously.

4. (SingleIO): For each edge $e_X = (n_X \xrightarrow{\xi_{\ddot{A}}; \xi_C} n_X^t) \in \mathcal{E}_X$, either both $\xi_{\ddot{A}}$ and ξ_C are I/O paths or both $\xi_{\ddot{A}}$ and ξ_C are I/O-free.

(SingleIO) enforces lockstep execution of non-silent trace-producing instructions in $\xi_{\ddot{A}}$ and ξ_C .

5. (Similar-speed): Let $(e_X^1, e_X^2, \dots, e_X^m)$ be a cyclic path, so that $\forall_{1 \leq j \leq m} : e_X^j = (n_X^{f,j} \xrightarrow{\xi_{\ddot{A}}^j; \xi_C^j} n_X^{t,j}) \in \mathcal{E}_X$; $n_X^{f,1} = n_X^{t,m}$; and $\bigwedge_{j=1}^{m-1} (n_X^{t,j} = n_X^{f,j+1})$. For each cyclic path, $\neg \bigwedge_{j=1}^m (\xi_{\ddot{A}}^j = \epsilon)$ and $\neg \bigwedge_{j=1}^m (\xi_C^j = \epsilon)$ holds.

(SimilarSpeed) enforces divergence preservation: \ddot{A} diverges (i.e., does not terminate) iff C diverges.

6. (Well-formedness): If a node of the form $n_X = (-, \mathcal{W}_C)$ exists in \mathcal{N}_X , then n_X must be $(\mathcal{W}_{\ddot{A}}, \mathcal{W}_C)$.

(Well-formedness) ensures that if a well-formedness (WF) constraint is violated in C (indicated by transition to \mathcal{W}_C), then \ddot{A} must transition to $\mathcal{W}_{\ddot{A}}$ — recall that \mathcal{W} is used to signal violation of a WF constraint in C , and violation of a condition that can be assumed to never happen, e.g. stack overflow, in \ddot{A} (section 2.2.4). In other words, an error-free execution in \ddot{A} must never falsify the WF constraints.

7. (Safety): If a node of the form $n_X = (\mathcal{U}_{\ddot{A}}, -)$ exists in \mathcal{N}_X , then n_X must be $(\mathcal{U}_{\ddot{A}}, \mathcal{U}_C)$.

(Safety) ensures that \ddot{A} triggers \mathcal{U} only if C triggers \mathcal{U} — recall that \mathcal{U} is used to signal occurrence of an undefined behavior (UB) in C , and occurrence of UB or translation error in \ddot{A} (section 2.2.4). With (Safety), we ensure that \ddot{A} may have translation errors (or may trigger UB) only if a lockstep execution of C triggers UB.

²Recall that execution of X edge e_X^j is defined as execution of $\xi_{\ddot{A}}$ followed by execution of ξ_C^j .

8. (Coverage \ddot{A}): For each error-free node $n_X = (n_{\ddot{A}}, n_C) \in \mathcal{N}_X^{\text{DW}}$ and for each possible outgoing path $\xi_{\ddot{A}}^o = (n_{\ddot{A}} \rightarrow n_{\ddot{A}}^o)$, either $\xi_{\ddot{A}}^o$ is infeasible at n_X or there exists $e_X = (n_X \xrightarrow{\xi_{\ddot{A}}; \xi_C} n_X^t) \in \mathcal{E}_X$ such that either $\xi_{\ddot{A}}$ is a prefix of $\xi_{\ddot{A}}^o$ or $\xi_{\ddot{A}}^o$ is a prefix of $\xi_{\ddot{A}}$.

(Coverage \ddot{A}) ensures all executable paths of \ddot{A} are present or *covered* in X .

9. (Coverage C): At each node $n_X = (n_{\ddot{A}}, n_C) \in \mathcal{N}_X$, for some $\xi_{\ddot{A}} = (n_{\ddot{A}} \rightarrow n_{\ddot{A}}^t)$, let $\{e_X^1, e_X^2, \dots, e_X^m\}$ be the set of *all* outgoing edges such that $e_X^j = (n_X \xrightarrow{\xi_{\ddot{A}}; \xi_C^j} n_X^{t_j}) \in \mathcal{E}_X$ (for $1 \leq j \leq m$) and $n_X^{t_j} = (n_{\ddot{A}}^t, n_C^{t_j}) \in \mathcal{N}_X$. Then, $\{e_X^1, e_X^2, \dots, e_X^m\} \langle \mathcal{D}_X, \xi_{\ddot{A}} \rangle$ holds.

(Coverage C) requires that at least one of the outgoing edge e_X^j must execute to completion. (Mutex \ddot{A}), (Mutex C), and (Coverage C) together ensure that at a node n_X exactly one outgoing edge (if any) may be executed to completion.

10. (Inductive): For each error-free edge (an edge with error-free destination node) $e_X = (n_X \xrightarrow{\xi_{\ddot{A}}; \xi_C} n_X^t) \in \mathcal{E}_X$ such that $n_X, n_X^t \in \mathcal{N}_X^{\text{DW}}$, the Hoare triple $\{\phi_{n_X}\}(\xi_{\ddot{A}}; [\xi_C]_{\mathcal{D}_X}^{\text{ex}})\{\phi_{n_X^t}\}$ holds.

(Inductive) ensures that the invariant network Φ_X is inductively provable starting at start node $n_X^s \in \mathcal{N}_X$

11. (Equivalence): For each error-free node $n_X = (n_{\ddot{A}}, n_C) \in \mathcal{N}_X^{\text{DW}}$, $\Omega_{\ddot{A}} = \Omega_C$ must belong to ϕ_{n_X} .

(Equivalence) ensures that \ddot{A} and C produce identical non-silent traces in an error-free execution. Because our execution semantics observe (de)allocation in C and A (figs. 2.5, 2.8 and 2.10), (Equivalence) ensures that the allocation state (address sets) of common regions B in C and \ddot{A} is identical.

12. (Memory Access Correspondence) or (MAC): For each edge $e_X = (n_X \xrightarrow{\xi_{\ddot{A}}; \xi_C} n_X^t) \in \mathcal{E}_X$, such that $n_X^t \neq (-, \mathcal{U}_C) \in \mathcal{N}_X$, the following Hoare triples hold:

$$(a) \ \{\phi_{n_X} \wedge (\Sigma_{\ddot{A}}^{\text{rd}} = \Sigma_C^{\text{rd}} = \emptyset)\}(\xi_{\ddot{A}}; [\xi_C]_{\mathcal{D}_X}^{\text{ex}})\{(\Sigma_{\ddot{A}}^{\text{rd}} \setminus \Sigma_C^{\text{rd}}) \subseteq \Sigma_{\ddot{A}}^{GUF} \cup [\text{esp}, \text{stk}_e]\}$$

$$(b) \ \{\phi_{n_X} \wedge (\Sigma_{\ddot{A}}^{\text{wr}} = \Sigma_C^{\text{wr}} = \emptyset)\}(\xi_{\ddot{A}}; [\xi_C]_{\mathcal{D}_X}^{\text{ex}})\{(\Sigma_{\ddot{A}}^{\text{wr}} \setminus \Sigma_C^{\text{wr}}) \subseteq \Sigma_{\ddot{A}}^{G_wUF_w} \cup [\text{esp}, \text{stk}_e]\}$$

Recall that the ghost address sets $\Sigma_{\ddot{A}}^{\text{rd}}$ (Σ_C^{rd}) and $\Sigma_{\ddot{A}}^{\text{wr}}$ (Σ_C^{wr}) accumulate the memory accesses performed by $\ddot{A}(C)$ during its execution. (MAC) effectively requires that for every (unique) memory access made to address α belonging to region $r \in \{hp, cl, cs\}$ during execution of path $\xi_{\ddot{A}}$, there exists an access to α of the same read/write type on determinized path $[\xi_C]_{\mathcal{D}_X}^{\text{ex}}$ that also executes to completion.

This requirement enables a sound over-approximation of the set of addresses belonging to *hp*, *cl*, and *cs* for a faster SMT encoding (described later in theorem 3.5.3 and section 5.3.2). For (MAC) to be meaningful, $\Sigma_{\ddot{A},C}^{\text{rd}}$ and $\Sigma_{\ddot{A},C}^{\text{wr}}$ must not be included in X 's state elements over which a node invariant ϕ_{n_X} is inferred.

13. (MemEq): For each error-free node $n_X \in \mathcal{N}_X^{\text{DW}}$, $M_{\ddot{A}} =_{\Sigma_{\ddot{A}}^B \setminus (\Sigma_{\ddot{A}}^{Zl}|^v)} M_C$ must belong to ϕ_{n_X} .

(MemEq) requires the memory state of common regions modulo virtually-allocated locals $(\Sigma_{\ddot{A}}^B \setminus \Sigma_{\ddot{A}}^{Zl}|^v)^3$ of C and A to be identical at a node n_X so that these regions are mutated identically in a lockstep fashion in both C and \ddot{A} .

This requirement enables an efficient search algorithm which trades some completeness (by rejecting *sound* product graphs which do not respect this requirement) for a more efficient incremental exploration for the required product graph.

The first seven requirements are constraints on the graph structure of X and are referred to as *structural requirements*. The remaining six require discharge of proof obligations (in the form of Hoare triples) and are referred to as *semantic requirements*.

The first eleven requirements are *soundness requirements* that are required for witnessing refinement through X (section 3.3.2). The first twelve requirements are *fast-encoding requirements* that enable a faster SMT encoding (chapter 5). All thirteen are *search-algorithm requirements* that enable product graph search optimizations. Excluding (Coverage \ddot{A}) and (Coverage C), the remaining eleven are called *non-coverage requirements*.

3.3.2 Soundness of X requirements

Let $X = \ddot{A} \boxtimes C = (\mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X)$ be a determinized product graph that satisfies the soundness requirements (first eleven requirements in section 3.3.1).

Lemma 3.3.4 (X 's execution). *The following holds for an execution of X :*⁴

$$\begin{aligned} \forall \Omega, T'_{\ddot{A}}, T'_C : (X \downarrow_{\Omega} (T'_{\ddot{A}}, T'_C)) \Rightarrow & T'_{\ddot{A}} =_{st} T'_C \\ & \vee ((e(T'_{\ddot{A}}) = \mathcal{W}) \wedge (\tilde{e}(T'_{\ddot{A}}) \leq_{st} T'_C)) \\ & \vee ((e(T'_C) = \mathcal{U}) \wedge (\tilde{e}(T'_C) \leq_{st} T'_{\ddot{A}})) \end{aligned} \quad (3.1)$$

³Recall that $\Sigma_{\ddot{A}}^{Zl}|^v$ is defined as the union of address sets of virtually-allocated locals in \ddot{A} (section 2.6.1)

⁴The relations (\downarrow) , $=_{st}$, and \leq_{st} are defined in section 2.4.

Proof. The proof proceeds through a coinduction on the number of edges executed by X . We prove that the execution of a single edge $e_X = (n_X \xrightarrow{\xi_{\bar{A}}; \xi_C} n_X^t) \in \mathcal{E}_X$, starting at an error-free node $n_X \in \mathcal{N}_X^{\text{DFW}}$ in a state that satisfies $T'_{\bar{A}} =_{st} T'_C$, either reaches a terminating node n_X^t , such that final state satisfies the RHS of the \Rightarrow in eq. (3.1), or reaches a non-terminating node n_X^t , such that $T'_{\bar{A}} =_{st} T'_C$ holds at the end of execution of e_X .

Let $\{e_X^1, e_X^2, \dots, e_X^m\}$ be the set of *all* outgoing edges at error-free node $n_X \in \mathcal{N}_X^{\text{DFW}}$ such that $\forall 1 \leq j \leq m : e_X^j = (n_X \xrightarrow{\xi_{\bar{A}}^j; \xi_C^j} n_X^j) \in \mathcal{E}_X$. Due to (SingleIO) there can be two cases:

1. $\xi_{\bar{A}}^j$ and ξ_C^j are I/O paths. Because I/O paths are straight-line sequences of instructions (with no branching), it must be true that $j = m = 1$. Further, an I/O path can only end at an error-free node n_X^j that must satisfy (Equivalence) requirement. (Equivalence) (through $\Omega_{\bar{A}} = \Omega_C$) implies production of identical non-silent trace events. Therefore, the claim holds.
2. $\xi_{\bar{A}}^j$ and ξ_C^j are I/O-free. Due to (Mutex \bar{A}) and (Coverage \bar{A}), it must be possible to execute a path $\xi_{\bar{A}}^j$ to completion. Due to (Coverage C), there exists some outgoing edge $e_X^j = (n_X \xrightarrow{\xi_{\bar{A}}^j; \xi_C^j} n_X^j) \in \mathcal{E}_X$ that is executed to completion. Further, due to (Mutex C), such an edge e_X^j must be unique. The execution of $\xi_{\bar{A}}^j$ followed by execution of ξ_C^j effectively causes X to execute e_X^j and reach node $n_X^j = (n_{\bar{A}}^j, n_C^j)$.

The execution of $\xi_{\bar{A}}^j$ may end at either: (1) the error node $\mathcal{W}_{\bar{A}}$, (2) the error node $\mathcal{U}_{\bar{A}}$, (3) an error-free node $n_{\bar{A}}^j$.

- In case (1), the execution ends at the error node $\mathcal{W}_{\bar{A}}$. Because the traces were stuttering equivalent before the execution of e_X^j (coinduction hypothesis), and the execution of $\xi_{\bar{A}}^j$ must only produce the \mathcal{W} trace event (as $\xi_{\bar{A}}^j$ is I/O-free and cannot contain rd/wr instructions), $(e(T'_{\bar{A}}) = \mathcal{W} \wedge \tilde{e}(T_{\bar{A}}) \leq_{st} T'_C)$ will hold.
- In case (2), due to the (Safety) requirement, execution of e_X^j must reach node $n_X^j = (\mathcal{U}_{\bar{A}}, \mathcal{U}_C)$. Moreover, the execution $\xi_{\bar{A}}^j$ and ξ_C^j must only generate the error code \mathcal{U} as a trace event (as both $\xi_{\bar{A}}^j$ and ξ_C^j are I/O-free). Because the traces were stuttering equivalent before the execution of e_X^j (coinduction hypothesis), $(e(T'_C) = \mathcal{U} \wedge \tilde{e}(T'_C) \leq_{st} T'_{\bar{A}})$ will hold.
- In case (3), we analyze each possibility of n_X^j separately. n_X^j must be of one of the following forms: (a) $(n_{\bar{A}}^j, \mathcal{W}_C)$, (b) $(n_{\bar{A}}^j, \mathcal{U}_C)$, or (c) an error-free node $(n_{\bar{A}}^j, n_C^j)$, where n_C^j is an error-free node (recall that $n_{\bar{A}}^j$ is also an error-free node in this case).

Case (a) cannot occur due to the (Well-formedness) requirement. In case (b), $(e(T'_C) = \mathcal{U} \wedge \tilde{e}(T'_C) \leq_{st} T'_A)$ holds due to $\xi_{\tilde{A}}^j$ and ξ_C^j being I/O-free and coinduction hypothesis (similar reasoning as case (2) above). In case (c), due to $\xi_{\tilde{A}}^j$ and ξ_C^j being I/O-free, their execution cannot produce any non-silent trace event. Thus, due to coinduction hypothesis, $T'_A =_{st} T'_C$ must hold at n_X^j .

Finally, consider the case when ξ_C^j is ϵ and $\xi_{\tilde{A}}^j$ is not. Due to (Similar-speed), there exists a finite sequence of edges $(e_X^{x_1}, e_X^{x_2}, \dots, e_X^{x_n})$ such that $\forall_{1 \leq i \leq n} : e_X^{x_i} = (n_X^{f,i} \xrightarrow{\xi_{\tilde{A}}^{x_i}; \xi_C^{x_i}} n_X^{t,i}) \in \mathcal{E}_X$, $e_X^j = e_X^{x_1}$, $\forall_{1 \leq i < n} : \xi_C^{x_i} = \epsilon$, and $\xi_C^{x_n} \neq \epsilon$. Similar argument can be used when $\xi_{\tilde{A}}^j$ is ϵ and ξ_C^j is not. (Similar-speed) thus ensures that the silent events in both traces differ only by a finite amount, thereby upholding $T'_A =_{st} T'_C$. \square

Lemma 3.3.5 (X 's trace is derived from C 's trace). *The following holds for an execution of X :*

$$\begin{aligned} \forall \Omega, T'_A, T'_C : (X \downarrow_{\Omega} (T'_A, T'_C)) &\Rightarrow \exists T_C : (C \downarrow_{\Omega} T_C) \\ &\wedge (T'_C =_{st} T_C \\ &\vee ((e(T'_A) = \mathcal{W}) \wedge (\tilde{e}(T'_A) \leq_{st} T_C))) \end{aligned}$$

Proof. The proof proceeds through a coinduction on the number of edges executed by X . Suppose X and C start execution with states $\sigma_X = (\sigma_{\tilde{A}}, \sigma_C)$ and σ_C at error-free nodes $n_X = (n_{\tilde{A}}, n_C) \in \mathcal{N}_X^{\text{DFW}}$ and $n_C \in \mathcal{N}_C^{\text{DFW}}$ respectively, such that $T_C =_{st} T'_C$, where $T_C \in \sigma_C$ and $(T'_A, T'_C) \in \sigma_X$, holds.

Consider the execution of edge $e_X = (n_X \xrightarrow{\xi_{\tilde{A}}; \xi_C} n_X^t) \in \mathcal{E}_X$, starting at n_X on state σ_X . If ξ_C is executed, as part of e_X 's execution, using some sequence of non-deterministic choices determined by \mathcal{D}_X , the same path ξ_C can be executed in C for the same sequence of non-deterministic choices. As both executions start in identical states, they will produce identical sequence of trace events till execution reaches the sink node n_C^t where $T_C =_{st} T'_C$ will hold (note that execution of $\xi_{\tilde{A}}$ may not modify the state elements of C in σ_X as both have disjoint state space).

If ξ_C is ϵ , then, due to (Similar-speed), there exists a finite sequence of edges $(e_X^1, e_X^2, \dots, e_X^m)$ such that $e_X = e_X^1$, $\forall_{1 \leq j \leq m} : e_X^j = (n_X^{f,j} \xrightarrow{\xi_{\tilde{A}}^j; \xi_C^j} n_X^{t,j}) \in \mathcal{E}_X$, $\forall_{1 \leq j < m} : \xi_C^j = \epsilon$ and $\xi_C^m \neq \epsilon$. Let n_X^t be $n_X^{t,m}$ in this case.

If $n_X^t = (n_{\tilde{A}}^t, n_C^t)$ is a non-terminating node, then the claim holds due to the coinduction

hypothesis. Similarly, if both n_X^t and n_C^t are terminating nodes, then the claim holds by definition.

Consider the case when $n_X^t = (n_{\ddot{A}}^t, n_C^t)$ is a terminating node (due to $n_{\ddot{A}}^t$ being a terminating node) but n_C^t is not a terminating node. There are three possibilities for $n_{\ddot{A}}^t$ in this case:

- $n_{\ddot{A}}^t = \mathcal{W}_{\ddot{A}}$: Due to (Equivalence), $T'_{\ddot{A}} =_{st} T'_C$ holds at n_X . Further, due to (SingleIO), $\xi_{\ddot{A}}$ cannot produce any non-silent trace event other than \mathcal{W} . Hence, $T'_{\ddot{A}} \leq_{st} T_C$ holds due to coinduction hypothesis.
- $n_{\ddot{A}}^t = \mathcal{U}_{\ddot{A}}$: Due to (Safety), $n_X^t = (n_{\ddot{A}}^t, n_C^t)$ must be of the form $(\mathcal{U}_{\ddot{A}}, \mathcal{U}_C)$. However, this violates the assumption that n_C^t is a non-terminating node. Therefore, this case is not possible.
- $n_{\ddot{A}}^t$ is an error-free terminating node: This case is not possible due to (Termination) requiring n_C^t to be error-free terminating node whenever $n_{\ddot{A}}^t$ is an error-free terminating node.

□

Lemma 3.3.6 (\ddot{A} 's traces are in X). *The following holds for an execution of \ddot{A} :*

$$\begin{aligned} \forall \Omega, T_{\ddot{A}} : (\ddot{A} \downarrow_{\Omega} T_{\ddot{A}}) &\Rightarrow \exists T'_{\ddot{A}}, T'_C : (X \downarrow_{\Omega} (T'_{\ddot{A}}, T'_C)) \\ &\wedge (T_{\ddot{A}} =_{st} T'_{\ddot{A}} \\ &\quad \vee ((e(T'_C) = \mathcal{U}) \wedge (e(T'_{\ddot{A}}) \neq \mathcal{W}) \wedge (\tilde{e}(T'_C) \leq_{st} T_{\ddot{A}}))) \end{aligned} \quad (3.2)$$

Proof. Consider an execution of X that is currently at an error-free node $n_X = (n_{\ddot{A}}, n_C) \in \mathcal{N}_X^{\text{WF}}$. We show by coinduction on the number of edges executed in \ddot{A} , starting at $n_{\ddot{A}}$, that eq. (3.2) holds. The proof of the lemma follows by using $n_X = n_X^s = (n_{\ddot{A}}^s, n_C^s) \in \mathcal{N}_X$.

Due to (Mutex \ddot{A}), (Mutex C), (Coverage \ddot{A}), and (Coverage C), there exists exactly one $e_X = (n_X \xrightarrow{\xi_{\ddot{A}}; \xi_C} n_X^t) \in \mathcal{E}_X$ such that $\xi_{\ddot{A}}$ and ξ_C execute to completion to reach $n_X^t = (n_{\ddot{A}}^t, n_C^t) \in \mathcal{N}_X$. Consider the following two cases for $\xi_{\ddot{A}}$:

- $\xi_{\ddot{A}} \neq \epsilon$: If n_X^t is an error-free node, then for each non-deterministic choice that an execution of \ddot{A} can make to execute to completion, an execution of \ddot{A} as part of X can make as well. Thus, due to coinduction hypothesis, both executions can produce

identical observable events such that $T_{\ddot{A}} =_{st} T'_{\ddot{A}}$ holds. If $n_{\mathbb{C}}^t = \mathcal{W}_{\mathbb{C}}$, then $n_{\ddot{A}}^t$ must also be $\mathcal{W}_{\ddot{A}}$ due to (Well-formedness), and $T_{\ddot{A}} =_{st} T'_{\ddot{A}}$ holds due to (SingleIO) and coinduction hypothesis. If $n_{\mathbb{C}}^t = \mathcal{U}_{\mathbb{C}}$ and $n_{\ddot{A}}^t = \mathcal{W}_{\ddot{A}}$, $T_{\ddot{A}} =_{st} T'_{\ddot{A}}$ holds due to (SingleIO) and coinduction hypothesis. If $n_{\mathbb{C}}^t = \mathcal{U}_{\mathbb{C}}$ and $n_{\ddot{A}}^t \neq \mathcal{W}_{\ddot{A}}$, then $\tilde{e}(T'_{\mathbb{C}}) \leq_{st} T_{\ddot{A}}$ holds due to (SingleIO) and coinduction hypothesis. $n_{\mathbb{C}}^t \neq \mathcal{U}_{\mathbb{C}}$ and $n_{\ddot{A}}^t = \mathcal{U}_{\ddot{A}}$ is not possible due to (Safety).

- $\xi_{\ddot{A}} = \epsilon$: Execute k edges in \mathbb{X} before a non- ϵ path is encountered, where k is the length of the longest sequence of edges in \mathbb{X} such that an edge $e_{\mathbb{X}} = (n_{\mathbb{X}} \xrightarrow{\xi_{\ddot{A}}; \xi_{\mathbb{C}}} n'_{\mathbb{X}})$ with $\xi_{\ddot{A}} \neq \epsilon$ is reached; then repeat the coinductive step above. Due to (Similar-speed), k must be defined.

□

Theorem 3.3.7 (\mathbb{X} witnesses $\mathbb{C} \sqsupseteq \ddot{A}$). *If there exists $\mathbb{X} = \ddot{A} \boxtimes \mathbb{C}$ that satisfies the soundness requirements, then $\mathbb{C} \sqsupseteq \ddot{A}$ holds.*

Proof. Consider an execution of \ddot{A} under world Ω . Using lemma 3.3.6, we have:

$$\begin{aligned} \forall \Omega, T_{\ddot{A}} : (\ddot{A} \downarrow_{\Omega} T_{\ddot{A}}) &\Rightarrow \exists T'_{\ddot{A}}, T'_{\mathbb{C}} : (\mathbb{X} \downarrow_{\Omega} (T'_{\ddot{A}}, T'_{\mathbb{C}})) \\ &\wedge (T_{\ddot{A}} =_{st} T'_{\ddot{A}} \\ &\quad \vee ((e(T'_{\mathbb{C}}) = \mathcal{U}) \wedge (e(T'_{\ddot{A}}) \neq \mathcal{W}) \wedge (\tilde{e}(T'_{\mathbb{C}}) \leq_{st} T_{\ddot{A}}))) \end{aligned} \quad (3.3)$$

Instantiating lemma 3.3.5, we obtain,

$$\begin{aligned} \forall \Omega, T_{\ddot{A}} : (\ddot{A} \downarrow_{\Omega} T_{\ddot{A}}) &\Rightarrow \exists T'_{\ddot{A}}, T'_{\mathbb{C}} : (\mathbb{X} \downarrow_{\Omega} (T'_{\ddot{A}}, T'_{\mathbb{C}})) \\ &\wedge (T_{\ddot{A}} =_{st} T'_{\ddot{A}} \\ &\quad \vee ((e(T'_{\mathbb{C}}) = \mathcal{U}) \wedge (e(T'_{\ddot{A}}) \neq \mathcal{W}) \wedge (\tilde{e}(T'_{\mathbb{C}}) \leq_{st} T_{\ddot{A}}))) \\ &\wedge (\exists T_{\mathbb{C}} : (\mathbb{C} \downarrow_{\Omega} T_{\mathbb{C}}) \\ &\quad \wedge (T'_{\mathbb{C}} =_{st} T_{\mathbb{C}} \\ &\quad \quad \vee ((e(T'_{\ddot{A}}) = \mathcal{W}) \wedge (\tilde{e}(T'_{\ddot{A}}) \leq_{st} T_{\mathbb{C}}))) \end{aligned} \quad (3.4)$$

We consider each minterm in the sum-of-products representation of the following terms

in the RHS of eq. (3.4):

$$\begin{aligned} & (T_{\ddot{A}} =_{st} T'_{\ddot{A}} \\ & \quad \vee ((e(T'_C) = \mathcal{U}) \wedge (e(T'_{\ddot{A}}) \neq \mathcal{W}) \wedge (\tilde{e}(T'_C) \leq_{st} T_{\ddot{A}}))) \\ \wedge & (T'_C =_{st} T_C \\ & \quad \vee ((e(T'_{\ddot{A}}) = \mathcal{W}) \wedge (\tilde{e}(T'_{\ddot{A}}) \leq_{st} T_C))) \end{aligned}$$

1. $(T_{\ddot{A}} =_{st} T'_{\ddot{A}}) \wedge (T'_C =_{st} T_C)$ holds.

Instantiating lemma 3.3.4 in eq. (3.4), there are three cases:

- (a) $T'_{\ddot{A}} =_{st} T'_C$ holds.

Due to $=_{st}$ being an equivalence relation, we have $T_{\ddot{A}} =_{st} T_C$ and, therefore, $C \sqsupseteq \ddot{A}$ holds.

- (b) $(e(T'_{\ddot{A}}) = \mathcal{W}) \wedge (\tilde{e}(T'_{\ddot{A}}) \leq_{st} T'_C)$ holds.

As $=_{st}$ is congruent with respect to \leq_{st} , we have $(e(T_{\ddot{A}}) = \mathcal{W}) \wedge (\tilde{e}(T_{\ddot{A}}) \leq_{st} T_C)$, which is equivalent to $W_{\text{pre}}^{\Omega, T_{\ddot{A}}}(C)$. Therefore, $C \sqsupseteq \ddot{A}$ holds.

- (c) $(e(T'_C) = \mathcal{U}) \wedge (e(T'_{\ddot{A}}) \neq \mathcal{W}) \wedge (\tilde{e}(T'_C) \leq_{st} T'_{\ddot{A}})$ holds.

Using congruence of $=_{st}$ with respect to \leq_{st} , we have $(e(T_C) = \mathcal{U}) \wedge (\tilde{e}(T_C) \leq_{st} T_{\ddot{A}})$, which is equivalent to $U_{\text{pre}}^{\Omega, T_{\ddot{A}}}(C)$. Therefore, $C \sqsupseteq \ddot{A}$ holds.

2. $(T_{\ddot{A}} =_{st} T'_{\ddot{A}}) \wedge ((e(T'_{\ddot{A}}) = \mathcal{W}) \wedge (\tilde{e}(T'_{\ddot{A}}) \leq_{st} T_C))$ holds.

Using definition of $=_{st}$ and congruence of $=_{st}$ with respect to \leq_{st} , we have $(e(T_{\ddot{A}}) = \mathcal{W}) \wedge (\tilde{e}(T_{\ddot{A}}) \leq_{st} T_C)$, which is equivalent to $W_{\text{pre}}^{\Omega, T_{\ddot{A}}}(C)$. Therefore, $C \sqsupseteq \ddot{A}$ holds.

3. $((e(T'_C) = \mathcal{U}) \wedge (\tilde{e}(T'_C) \leq_{st} T_{\ddot{A}})) \wedge (T'_C =_{st} T_C)$ holds.

Using definition of $=_{st}$ and congruence of $=_{st}$ with respect to \leq_{st} , we have $(e(T_C) = \mathcal{U}) \wedge (\tilde{e}(T_C) \leq_{st} T_{\ddot{A}})$, which is equivalent to $U_{\text{pre}}^{\Omega, T_{\ddot{A}}}(C)$. Therefore, $C \sqsupseteq \ddot{A}$ holds.

4. $((e(T'_C) = \mathcal{U}) \wedge (e(T'_{\ddot{A}}) \neq \mathcal{W}) \wedge (\tilde{e}(T'_C) \leq_{st} T_{\ddot{A}})) \wedge ((e(T'_{\ddot{A}}) = \mathcal{W}) \wedge (\tilde{e}(T'_{\ddot{A}}) \leq_{st} T_C))$ holds.

This case is not possible due to the mutually unsatisfiable clauses $\dots \wedge (e(T'_{\ddot{A}}) \neq \mathcal{W}) \wedge \dots \wedge (e(T'_{\ddot{A}}) = \mathcal{W}) \wedge \dots$

□

3.3.3 Global Invariants in \mathbb{C} , $\check{\mathbb{A}}$, and \mathbb{X}

Definition 3.3.8 (Non-entry Node). *Let $P \in \{\check{\mathbb{A}}, \mathbb{C}\}$. A node $n_P \in \mathcal{N}_P$ is called a **non-entry node** iff it does not correspond to a node due to $(\text{ENTRY}_{\mathbb{C}})$ and $(\text{ENTRY}_{\check{\mathbb{A}}})$ (figs. 2.5 and 2.7) in P . A node $n_{\mathbb{X}} = (n_{\check{\mathbb{A}}}, n_{\mathbb{C}}) \in \mathcal{N}_{\mathbb{X}}$ is called a non-entry node iff both $n_{\check{\mathbb{A}}}$ and $n_{\mathbb{C}}$ are non-entry nodes.*

Due to the execution semantics of $\check{\mathbb{A}}$ and \mathbb{C} , certain invariants hold by construction in $\check{\mathbb{A}}$ and \mathbb{C} . We call these invariants *global invariants* as they hold at each error-free, non-entry node.

Theorem 3.3.9 (Global Invariants in $\check{\mathbb{A}}$). *The following invariants hold at each error-free, non-entry node $n_{\check{\mathbb{A}}} \in \mathcal{N}_{\check{\mathbb{A}}}^{\text{DW}}$:*

- (*em.f tracks emptiness*) $\Sigma_{\check{\mathbb{A}}}^f = \emptyset \Leftrightarrow \text{em.f}$, for $f \in F$. Note that because $|\Sigma_{\check{\mathbb{A}}}^f| > 0$ by definition, $\text{em.f} = \text{false}$ holds for all $f \in F$.
- (*sz.f tracks size*) $\text{sz.f} = |\Sigma_{\check{\mathbb{A}}}^f| = \text{sz}(\text{T}(f))$ for $f \in F$ ⁵.
- (*Address sets of F are intervals*) $(\text{em.f} \vee [\text{lb.f}, \text{ub.f}] = \Sigma_{\check{\mathbb{A}}}^f)$, for $f \in F$. Or, using $\text{em.f} = \text{false}$, $[\text{lb.f}, \text{ub.f}] = \Sigma_{\check{\mathbb{A}}}^f$.
- (*Alignment of f*) $\text{aligned}_{\text{alignmt}(f)}(\text{lb.f})$, for $f \in F$ ⁶.
- (*Stack bounds*) $\Sigma_{\check{\mathbb{A}}}^{\{\text{stk}\} \cup Y} \cup (\Sigma_{\check{\mathbb{A}}}^Z \setminus (\Sigma_{\check{\mathbb{A}}}^{Z_l}|^v)) = [\text{esp}, \text{stk}_e]$.
- (*cs and cl*) $\Sigma_{\check{\mathbb{A}}}^{\{\text{cs}, \text{cl}\}} = [\text{stk}_e + 1, \text{cs}_e]$
- (*Heap subset*) $\Sigma_{\check{\mathbb{A}}}^{\text{hp}} \subseteq \text{comp}(\Sigma_{\check{\mathbb{A}}}^{G \cup F \cup \{\text{cv}\}} \cup \Sigma_{\check{\mathbb{A}}}^{Z_l}|^v \cup [\text{esp}, \text{cs}_e])$
- (*Disjoint regions in $\check{\mathbb{A}}$*) $\neg \text{ov}(\Sigma_{\check{\mathbb{A}}}^{\text{hp}}, \Sigma_{\check{\mathbb{A}}}^{\text{cl}}, \Sigma_{\check{\mathbb{A}}}^{\text{cv}}, \Sigma_{\check{\mathbb{A}}}^{\text{vrdc}}, \dots, \Sigma_{\check{\mathbb{A}}}^g, \dots, \Sigma_{\check{\mathbb{A}}}^y, \dots, \Sigma_{\check{\mathbb{A}}}^z, \dots)$ and $\neg \text{ov}(\Sigma_{\check{\mathbb{A}}}^{\text{hp}}, \Sigma_{\check{\mathbb{A}}}^{\text{cl}}, \Sigma_{\check{\mathbb{A}}}^{\text{vrdc}}, \dots, \Sigma_{\check{\mathbb{A}}}^g, \dots, \Sigma_{\check{\mathbb{A}}}^y, \dots, \Sigma_{\check{\mathbb{A}}}^z|^\text{s}, \dots, \Sigma_{\check{\mathbb{A}}}^f, \dots, \Sigma_{\check{\mathbb{A}}}^{\text{stk}}, \Sigma_{\check{\mathbb{A}}}^{\text{cs}})$.
- (*Read-only memory in $\check{\mathbb{A}}$*) $M_{\check{\mathbb{A}}} =_{i^r} \text{ROM}_{\check{\mathbb{A}}}^r(i^r)$ for $r \in F_r$.

Proof sketch: By induction on the number of transitions executed in $\check{\mathbb{A}}$ with the base case defined by the first transition out of $(\text{ENTRY}_{\check{\mathbb{A}}})$ in fig. 2.11. □

⁵Recall that $\text{sz}(\text{T}(r))$ gives the size in bytes of the variable named r (table 2.1).

⁶Recall that $\text{aligned}_{\text{alignmt}(r)}(\text{lb.r})$ holds iff lb.r is aligned by the alignment of variable r (table 2.1).

Theorem 3.3.10 (Global Invariants in \mathbb{C}). *The following invariants hold at each error-free, non-entry node $n_{\mathbb{C}} \in \mathcal{N}_{\mathbb{C}}^{\text{UW}}$:*

- ($\boxed{\text{em}.r}$ tracks emptiness) $\Sigma_{\mathbb{C}}^r = \emptyset \Leftrightarrow \boxed{\text{em}.r}$, for $r \in G \cup Y \cup Z$. Note that $\Sigma_{\mathbb{C}}^r = \emptyset$ may hold only for $r \in \{\text{vrdc}\} \cup Z$.
- ($\boxed{\text{sz}.r}$ tracks size) $\boxed{\text{sz}.r} = |\Sigma_{\mathbb{C}}^r|$ for $r \in G \cup Y$. $\boxed{\text{sz}.r} = \text{sz}(\mathbb{T}(r))$, for $r \in G \cup (Y \setminus \{\text{vrdc}\})$.
- ($\boxed{\text{1stSz}.z_l}$ tracks size) $\boxed{\text{1stSz}.z_l} = |\Sigma_{\mathbb{C}}^{z_l}|$ for $z_l \in Z_l$.
- ($\boxed{\text{lb}.r}, \boxed{\text{ub}.r}$ track bounds) $\boxed{\text{em}.r} \vee (\boxed{\text{lb}.r} = \text{lb}(\Sigma_{\mathbb{C}}^r) \wedge \boxed{\text{ub}.r} = \text{ub}(\Sigma_{\mathbb{C}}^r))$, for $r \in G \cup Y \cup Z$.
- (Address sets of G, Y, Z_l are intervals) $\boxed{\text{em}.r} \vee (\boxed{\text{lb}.r}, \boxed{\text{ub}.r}] = \Sigma_{\mathbb{C}}^r$, for $r \in G \cup Y \cup Z_l$.
As a consequence, we have: $\boxed{\text{em}.r} \vee ((\boxed{\text{lb}.r} \leq_u \boxed{\text{ub}.r}) \wedge (\boxed{\text{ub}.r} = \boxed{\text{lb}.r} + \boxed{\text{sz}.r} - 1_{i_{32}}))$,
for $r \in G \cup Y$ and $\boxed{\text{em}.z_l} \vee ((\boxed{\text{lb}.z_l} \leq_u \boxed{\text{ub}.z_l}) \wedge (\boxed{\text{ub}.z_l} = \boxed{\text{lb}.z_l} + \boxed{\text{1stSz}.z_l} - 1_{i_{32}}))$
for $z_l \in Z_l$.
- (Alignment of g and y) $\text{aligned}_{\text{alignmt}(r)}(\boxed{\text{lb}.r})$, for $r \in G \cup (Y \setminus \{\text{vrdc}\})$.
- (Disjoint regions in \mathbb{C}) $\neg \text{ov}(\Sigma_{\mathbb{C}}^{hp}, \Sigma_{\mathbb{C}}^{cl}, \Sigma_{\mathbb{C}}^{cv}, \Sigma_{\mathbb{C}}^{\text{vrdc}}, \dots, \Sigma_{\mathbb{C}}^g, \dots, \Sigma_{\mathbb{C}}^y, \dots, \Sigma_{\mathbb{C}}^z, \dots)$.
- (Read-only memory in \mathbb{C}) $M_{\mathbb{C}} =_{i_r} \text{ROM}_{\mathbb{C}}^r(i_r^r)$ for $r \in G_r$.

Proof sketch. By induction on the number of transitions executed in \mathbb{C} with the base case defined by the first transition out of $(\text{ENTRY}_{\mathbb{C}})$ in fig. 2.5 □

Theorem 3.3.11 (Global Invariants in \mathbb{X}). *Let $\mathbb{X} = \ddot{\mathbb{A}} \boxtimes \mathbb{C}$ be a product graph that satisfies the search-algorithm requirements of section 3.3.1. The following invariants hold at an error-free, non-entry node $n_{\mathbb{X}} = (n_{\ddot{\mathbb{A}}}, n_{\mathbb{C}}) \in \mathcal{N}_{\mathbb{X}}^{\text{UW}}$.*

1. *The invariants stated in theorems 3.3.9 and 3.3.10.*

2. (Stack subset) $\Sigma_{\ddot{\mathbb{A}}}^{\text{stk}} \subseteq \Sigma_{\mathbb{C}}^{\{\text{cv}, \text{free}\}} \cup \Sigma_{\ddot{\mathbb{A}}}^{Z_l} |^v$

Proof sketch. Item 1 follows because $n_{\mathbb{X}}$ is an error-free, non-entry node iff both $n_{\ddot{\mathbb{A}}}$ and $n_{\mathbb{C}}$ are error-free, non-entry nodes.

Item 2 follows from (Disjoint regions in $\ddot{\mathbb{A}}$) of item 1 and the (Equivalence) requirement. □

3.4 Callers' Virtual Smallest Semantics

We are going to introduce two different semantics for \mathbf{C} and \mathbf{A} : callers' virtual smallest semantics (this section) and safety-relaxed semantics (section 3.5). These semantics are amenable to a faster SMT encoding and we separately prove that the theorems proved using these new semantics translate to theorems proved with the original semantics.

We construct \mathbf{C}' and \mathbf{A}' from \mathbf{C} and \mathbf{A} by using new *callers' virtual smallest semantics* where the cv region is made empty, i.e., $\Sigma_{\mathbf{C}'}^{cv} = \Sigma_{\mathbf{A}'}^{cv} = \emptyset$. With an empty cv , we compute the address set of region `free` as $\Sigma_P^{\text{free}} = \text{comp}(\Sigma_P^{B \cup F \cup S})$ for $P \in \{\mathbf{C}', \mathbf{A}'\}$.

Formally, we obtain \mathbf{C}' and \mathbf{A}' from \mathbf{C} and \mathbf{A} by removing assignments to $\Sigma_{\mathbf{C}}^{cv}$ and $\Sigma_{\mathbf{A}}^{cv}$ due to $(\text{ENTRY}_{\mathbf{C}})$ and $(\text{ENTRY}_{\mathbf{A}})$ respectively (figs. 2.5 and 2.7) and replacing uses of $\Sigma_{\mathbf{C}}^{cv}$ and $\Sigma_{\mathbf{A}}^{cv}$ due to $(\text{ENTRY}_{\mathbf{C}})$, $(\text{ENTRY}_{\mathbf{A}})$, $(\text{OP-ESP}')$, $(\text{LOAD}_{\mathbf{A}})$, $(\text{STORE}_{\mathbf{A}})$, (ALLOCS') , and (ALLOCV) (figs. 2.4 and 2.11) with \emptyset :

1. In $(\text{ENTRY}_{\mathbf{C}})$ and $(\text{ENTRY}_{\mathbf{A}})$, $\text{addrSetsAreWF}(\Sigma_P^{hp}, \Sigma_P^{cl}, \Sigma_P^{cv}, \dots, i_P^g, \dots, \Sigma_P^f, \dots, i_P^y, \dots, \Sigma_P^{\text{vrdc}})$ is replaced with $\text{addrSetsAreWF}(\Sigma_P^{hp}, \Sigma_P^{cl}, \dots, i_P^g, \dots, \Sigma_P^f, \dots, i_P^y, \dots, \Sigma_P^{\text{vrdc}})$ for $P \in \{\mathbf{C}, \mathbf{A}\}$.
2. In $(\text{OP-ESP}')$, $\text{intrvlInSet}(t, \text{esp} - 1_{i_{32}}, \Sigma_{\mathbf{A}}^{\{\text{free}\}} \cup ((\Sigma_{\mathbf{A}}^{cv} \cup \Sigma_{\mathbf{A}}^{Zl|v}) \setminus \Sigma_{\mathbf{A}}^F))$ is replaced with $\text{intrvlInSet}(t, \text{esp} - 1_{i_{32}}, \Sigma_{\mathbf{A}}^{\text{free}} \cup (\Sigma_{\mathbf{A}}^{Zl|v} \setminus \Sigma_{\mathbf{A}}^F))$.
3. In $(\text{LOAD}_{\mathbf{A}})$, $\text{ov}([p]_w, \Sigma_{\mathbf{A}}^{\text{free}} \cup ((\Sigma_{\mathbf{A}}^{cv} \cup \Sigma_{\mathbf{A}}^{Zl|v}) \setminus \Sigma_{\mathbf{A}}^{F \cup S}))$ is replaced with $\text{ov}([p]_w, \Sigma_{\mathbf{A}}^{\text{free}} \cup ((\Sigma_{\mathbf{A}}^{Zl|v}) \setminus \Sigma_{\mathbf{A}}^{F \cup S}))$.
4. In $(\text{STORE}_{\mathbf{A}})$, $\text{ov}([p]_w, \Sigma_{\mathbf{A}}^{\{\text{free}\} \cup G_r \cup F_r} \cup ((\Sigma_{\mathbf{A}}^{cv} \cup (\Sigma_{\mathbf{A}}^{Zl|v})) \setminus \Sigma_{\mathbf{A}}^{F_w \cup S}))$ is replaced with $\text{ov}([p]_w, \Sigma_{\mathbf{A}}^{\{\text{free}\} \cup G_r \cup F_r} \cup ((\Sigma_{\mathbf{A}}^{Zl|v}) \setminus \Sigma_{\mathbf{A}}^{F_w \cup S}))$.
5. In (ALLOCS') , $\text{ov}([v]_w, \Sigma_{\mathbf{A}}^{cv} \cup \Sigma_{\mathbf{A}}^{Zl|v})$ is replaced with $\text{ov}([v]_w, \Sigma_{\mathbf{A}}^{Zl|v})$.
6. In (ALLOCV) , $\text{intrvlInSet}_a(v, v+w-1_{i_{32}}, \text{comp}(\Sigma_{\mathbf{A}}^{B \cup \{cv\}}))$ is replaced with $\text{intrvlInSet}_a(v, v+w-1_{i_{32}}, \text{comp}(\Sigma_{\mathbf{A}}^B))$.

The callers' virtual smallest semantics are useful because they allow an over-approximation of the heap (hp) region in \mathbf{C}' , which can be efficiently encoded in SMT to achieve faster discharge (as described in chapter 5).

3.4.1 Soundness of Callers' Virtual Smallest semantics

Let A and C be transition graphs obtained due to original semantics described in figs. 2.4 to 2.8, 2.10 and 2.11. Let A' and C' be obtained from A and C respectively by applying the callers' virtual smallest semantics described in previous section. Let \ddot{A}' be obtained by annotating A' as described in section 2.4. Let \ddot{A} be obtained by annotating A such that annotations made in \ddot{A}' and \ddot{A} are identical. Let $X' = \ddot{A}' \boxtimes C' = (\mathcal{N}_{X'}, \mathcal{E}_{X'}, \mathcal{D}_{X'})$ be a product graph such that X' satisfies the search-algorithm requirements. We prove that there exists a product graph $X = \ddot{A} \boxtimes C = (\mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X)$ such that X satisfies the search-algorithm requirements. In other words, we show that existence of a product graph between procedures with callers' virtual smallest semantics implies existence of a product graph between the same procedures without callers' virtual smallest semantics.

Definition 3.4.1 ((Coverage C) holds for $\xi_{\ddot{A}}$ at n_X in X). *At a node $n_X \in \mathcal{N}_X$, let $\{e_X^1, e_X^2, \dots, e_X^m\}$ be the set of all outgoing edges such that $e_X^j = (n_X \xrightarrow{\xi_{\ddot{A}}, \xi_C^j} n_{\ddot{A}}^t, n_C^t)$ (for $1 \leq j \leq m$). Then, (Coverage C) holds for $\xi_{\ddot{A}}$ at n_X in X iff $\{e_X^1, e_X^2, \dots, e_X^m\} \langle \mathcal{D}_X, \xi_{\ddot{A}} \rangle$ holds.*

Notice that this definition is identical to the (Coverage C) definition in section 3.3.1, except that it defines (Coverage C) for a specific path $\xi_{\ddot{A}}$ starting at a specific node n_X . We define (Coverage \ddot{A}) at node n_X analogously.

Theorem 3.4.2 (Soundness of Callers' Virtual Smallest Semantics). *Given $X' = \ddot{A}' \boxtimes C' = (\mathcal{N}_{X'}, \mathcal{E}_{X'}, \mathcal{D}_{X'})$ that satisfies the search-algorithm requirements (section 3.3.1), it is possible to construct $X = \ddot{A} \boxtimes C = (\mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X)$ that also satisfies the search-algorithm requirements.*

Proof. Construct $X = X'$. Add extra edges in X to nodes $(\mathcal{W}_{\ddot{A}}, n_C)$ where n_C is an error-free node such that (Mutex \ddot{A}) is not violated. These extra edges help in ensuring (Coverage \ddot{A}) in X .

As the use of callers' virtual smallest semantics does not affect the graph structure of A and C (recall that the changes were limited to modifications to instructions of an edge), the seven structural requirements, (Mutex \ddot{A}), (Mutex C), (Termination), (SingleIO), (Well-formedness), (Safety), and (Similar-speed), should continue to hold for X .

Let $n_{X'} = (n_{\ddot{A}'}, n_{C'}) \in \mathcal{N}_{X'}$ be a node in X' and let $n_X = (n_{\ddot{A}}, n_C) \in \mathcal{N}_X$ be its corresponding node in X . Let $\xi_{\ddot{A}'}$ be an outgoing path at $n_{\ddot{A}'}$ in \ddot{A}' and let $\xi_{\ddot{A}}$ be its structurally similar

path originating at $n_{\bar{A}}$ in \bar{A} . Let $\{e_{X'}^1, \dots, e_{X'}^m\}$ be the set of all outgoing edges at $n_{X'}$ such that $\forall_{1 \leq j \leq m} : e_{X'}^j = (n_{X'} \xrightarrow{\xi_{\bar{A}}; \xi_{C'}^j} n_{X'}^j) \in \mathcal{E}_{X'}$. Let the set $\{e_{X'}^1, \dots, e_{X'}^m\}$ be defined analogously for X . Our proof completes by induction on the number of edges executed in X , starting at n_X .

We analyze the instructions in \bar{A} and C affected by the semantics change and consider the case when an edge $e_{\bar{A}} \in \xi_{\bar{A}}$ or $e_C \in \xi_C^j$ corresponds to it ⁷.

- (ENTRY_C) and (ENTRY _{\bar{A}}): The $\neg \text{addrSetsAreWF}(\dots)$ condition is weaker in \bar{A} and C than \bar{A}' and C' respectively. Consequently, the path condition for paths $\xi_{\bar{A}} = (n_{\bar{A}} \rightarrow n_{\bar{A}}^{\cancel{\mathcal{W}}})$ (where $n_{\bar{A}}^{\cancel{\mathcal{W}}} \in \mathcal{N}_{\bar{A}} \setminus \mathcal{W}_{\bar{A}}$) and $\xi_C = (n_C \rightarrow n_C^{\cancel{\mathcal{W}}})$ (where $n_C^{\cancel{\mathcal{W}}} \in \mathcal{N}_C \setminus \mathcal{W}_C$) that do not go to $\mathcal{W}_{\bar{A}}$ and \mathcal{W}_C respectively is stronger in \bar{A} and C than \bar{A}' and C' respectively.

Because the address sets returned by the `rd` instruction are arbitrary and identical across C and \bar{A} , due to (Equivalence), (Coverage_C) holds by construction in this case.

As the results of the `rd` instruction are arbitrary, the difference in infeasibility of $\xi_{\bar{A}'} = (n_{\bar{A}'} \rightarrow \mathcal{W}_{\bar{A}'})$ and structurally similar $\xi_{\bar{A}} = (n_{\bar{A}} \rightarrow \mathcal{W}_{\bar{A}})$ can only be due to the address set of regions in F (see definition of $\text{addrSetsAreWF}(\dots)$ in table 2.1) As $\Sigma_{\bar{A}'}^F = \Sigma_{\bar{A}}^F$, (Coverage _{\bar{A}}) at n_X should continue to hold in this case.

- (ALLOC), (ALLOCV), and (ALLOCS'): As $(\Sigma_{\bar{A}}^{cv} = \Sigma_C^{cv}) \supseteq (\Sigma_{\bar{A}'}^{cv} = \Sigma_{C'}^{cv} = \emptyset)$, the $\neg \text{intrvlInSet}_a(\dots)$ condition of (ALLOC) and (ALLOCV) and $\text{ov}(\dots)$ condition of (ALLOCS') is weaker in \bar{A} and C than \bar{A}' and C' respectively. Consequently, similarly to previous case, the path condition for paths that do not go to $\mathcal{W}_{\bar{A}}$ and \mathcal{W}_C respectively is stronger in \bar{A} and C than \bar{A}' and C' respectively.

Due to (SingleIO), the nodes $n_{\bar{A}}$ and n_C must either correspond to PCs due to: (1) (ALLOCV) and (ALLOC); or (2) (ALLOCS') and (ALLOC). Due to (Equivalence), $\Sigma_{\bar{A}}^{\text{comp}(B \cup \{cv\})} = \Sigma_C^{\text{comp}(B \cup \{cv\})} = \Sigma_C^{\text{free}}$ must hold at n_X . As, for $P \in \{\bar{A}, \bar{A}', C, C'\}$, $\Sigma_P^{\{hp, cl\}}$ is assigned arbitrarily at entry, the set of possible values for $\Sigma_P^{\text{comp}(B \cup \{cv\})}$ (note $\Sigma_{\bar{A}'}^{cv} = \Sigma_{C'}^{cv} = \emptyset$) remain identical in P at an error-free node n_X and $n_{X'}$. Thus, in case (1), the affected $\neg \text{intrvlInSet}_a(\dots)$ condition should have identical semantics in both X' and X and (Coverage _{\bar{A}}) and (Coverage_C) should continue to hold.

In case (2), a path $\xi_{\bar{A}'} = (n_{\bar{A}'} \rightarrow \mathcal{W}_{\bar{A}'})$ with an edge with the $\text{ov}(\dots)$ condition could be provably infeasible at $n_{X'}$ in X' but a similarly structured path $\xi_{\bar{A}}$ could

⁷Note that (LOAD_C), (STORE_C), (CALLV), and (CALL_C), are not affected as the cv region is inaccessible in C and cannot be returned by $\beta(x)$ for any variable x and $\beta_M(r)$ for any region r .

potentially be feasible at n_X in X — e.g., when $\Sigma_{\check{A}}^{Z_l|v} = \emptyset$. To ensure (Coverage \check{A}), we introduce edge $e'_X = ((n_{\check{A}}, n_C) \xrightarrow{\xi_{\check{A}}; \epsilon} (\mathcal{W}_{\check{A}}, n_C))$ for each such path $\xi_{\check{A}}$ in X . Notice that (Coverage C) holds for $\xi_{\check{A}}$ at n_X . Because $\xi_{\check{A}}$ does not contain any memory access, introduction of e'_X would not disturb (MAC).

For a path $(n_{\check{A}} \rightarrow n_{\check{A}}^{\mathcal{W}})$ (where $n_{\check{A}}^{\mathcal{W}} \in \mathcal{N}_{\check{A}} \setminus \{\mathcal{W}_{\check{A}}\}$), (Coverage C) holds due to (Stack subset) invariant (theorem 3.3.11) and by using identical reasoning as case (1) above.

- (OP-ESP'): The condition `intrvlInSet()` is not affected by the semantics change as the address sets $\Sigma_{\check{A}}^{\text{free}} \cup ((\Sigma_{\check{A}}^{cv} \cup \Sigma_{\check{A}}^{Z_l|v}) \setminus \Sigma_{\check{A}}^F)$ and $\Sigma_{\check{A}'}^{\text{free}} \cup (\Sigma_{\check{A}'}^{Z_l|v} \setminus \Sigma_{\check{A}'}^F)$ must evaluate to identical values (on states σ and σ' at nodes n_X and $n_{X'}$ in X and X' resp. such that $\phi_{n_X}(\sigma)$ and $\phi_{n_{X'}}(\sigma')$ hold) due to new definition of $\Sigma_{\check{A}'}^{\text{free}}$ in \check{A}' .
- (LOAD \check{A}) and (STORE \check{A}): Identical reasoning as (OP-ESP') case; the address set expressions should evaluate to identical values. Hence, no change in semantics for this case too.

As the path condition to an error-free node is only stronger (or equivalent) in \check{A} and C , the remaining semantic requirements, (Inductive), (Equivalence), (MAC), and (MemEq) should also continue to hold in X .

□

3.5 Safety-Relaxed Semantics

We define new *safety-relaxed semantics* for the assembly procedure A with callers' virtual smallest semantics. These semantics relax the memory-safety checks in A ; the soundness is retained in the context of product graph because of the (MAC) requirement.

Under the safety-relaxed semantics, we construct A' from A such that

- a $\varphi_l = \text{ov}([p]_w, \Sigma_{\check{A}}^{\text{free}} \cup ((\Sigma_{\check{A}}^{Z_l|v}) \setminus \Sigma_{\check{A}}^{F \cup S}))$ check due to (LOAD \check{A}) in A is replaced with $\varphi'_l = \text{ov}([p]_w, (\Sigma_{\check{A}}^{Z_l|v}) \setminus (\Sigma_{\check{A}}^F \cup [\text{esp}, \text{cs}_e]))$ in A' .
- a $\varphi_s = \text{ov}([p]_w, \Sigma_{\check{A}}^{\{\text{free}\} \cup G_r \cup F_r} \cup ((\Sigma_{\check{A}}^{Z_l|v}) \setminus \Sigma_{\check{A}}^{F_w \cup S}))$ check due to (STORE \check{A}) in A is replaced with $\varphi'_s = \text{ov}([p]_w, (\Sigma_{\check{A}}^{Z_l|v}) \setminus (\Sigma_{\check{A}}^{F_w} \cup [\text{esp}, \text{cs}_e]))$ in A' .
- a $\varphi_r = \neg(\text{M}^{\text{CS}})_{\Sigma_{\check{A}}^{\text{CS}}} M_A$ check due to (RET \check{A}) in A is replaced with $\varphi'_r = \text{false}$ in A' .

We call this construction a *safety-relaxed rewrite* and call A' the *safety-relaxed version* of A .

The callers' virtual smallest and safety-relaxed semantics are useful in enabling an efficient SMT encoding where the *hp* region in C' and A' is over-approximated and the *cs* region in A' is under-approximated (as described in chapter 5).

3.5.1 Soundness of Safety-Relaxed Semantics

Let \check{A}' be obtained by annotating A' as described in section 2.6. Let \check{A} be the annotated version of A such that the annotations made in \check{A} and \check{A}' are identical. Let C be the corresponding unoptimized IR procedure with the callers' virtual smallest semantics. Let $X' = \check{A}' \boxtimes C = (\mathcal{N}_{X'}, \mathcal{E}_{X'}, \mathcal{D}_{X'})$ be a product graph that satisfies the search-algorithm requirements. We prove that it is possible to construct $X = \check{A} \boxtimes C = (\mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X)$ such that X also satisfies search-algorithm requirements (theorem 3.5.3). In other words, we prove that the existence of a product graph between procedures with safety-relaxed semantics implies the existence of a product graph between procedures without safety-relaxed semantics.

Lemma 3.5.1 (Paths containing memory accesses do not modify allocation state of common regions). *Let $e_X = (n_X \xrightarrow{\xi_{\check{A}}; \xi_C} n_X^t) \in \mathcal{E}_X$. If $\xi_{\check{A}}$ contains an edge corresponding to $(\text{LOAD}_{\check{A}})$ or $(\text{STORE}_{\check{A}})$ (i.e., a load or store instruction), then $\xi_{\check{A}}$ does not modify the address sets corresponding to regions in B : $\Sigma_{\check{A}}^g$ (for each $g \in G$), $\Sigma_{\check{A}}^{hp}$, $\Sigma_{\check{A}}^{cl}$, $\Sigma_{\check{A}}^y$ (for each $y \in Y$), and $\Sigma_{\check{A}}^z$ (for each $z \in Z$).*

Proof. Once initialized in $(\text{ENTRY}_{\check{A}})$ in an I/O path that does not contain any load or store instruction (fig. 2.7), the address sets corresponding to regions $B \setminus Z$ are not modified during the entire execution of \check{A} .

The address set corresponding to a region $z \in Z$ may only be modified by the $(\text{de})\text{alloc}_{s,v}$ instructions. Due to (SingleIO) requirement, these $(\text{de})\text{alloc}_{s,v}$ instructions cannot exist as a part of longer paths that may contain load or store instructions (as evident from translations given in figs. 2.8, 2.10 and 2.11). \square

Note the the lemma holds for both original and safety-relaxed version of \check{A} . As a corollary, due to (SingleIO), ξ_C also does not modify the address sets corresponding to regions in B .

Lemma 3.5.2 ($\pi_{\Sigma_{\ddot{A}'}}(M_{\ddot{A}'})$ is not modified in \mathcal{X}'). Let $\mathcal{X}' = \ddot{A}' \boxtimes \mathcal{C}$ be a product graph for a lockstep execution between \ddot{A}' (with safety-relaxed semantics) and \mathcal{C} . If \mathcal{X}' satisfies the search-algorithm requirements, then $\boxed{M_{\mathcal{X}'}}^{cs} =_{\Sigma_{\ddot{A}'}} M_{\ddot{A}'}$ holds at each error-free, non-entry node $n_{\mathcal{X}'} \in \mathcal{N}_{\mathcal{X}'}^{\text{DW}}$.

Proof. For simplicity, let's first assume that there is only one outgoing edge $e_{\mathcal{X}'}^s = (n_{\mathcal{X}'}^s, \xrightarrow{\xi_{\ddot{A}'}^s; \xi_{\mathcal{C}}^s} n_{\mathcal{X}'}^{s2})$ from the start node $n_{\mathcal{X}'}^s$ to an error-free node $n_{\mathcal{X}'}^{s2}$ such that $\xi_{\ddot{A}'}^s$ and $\xi_{\mathcal{C}}^s$ represent the program paths corresponding to $(\text{ENTRY}_{\ddot{A}'})$ and $(\text{ENTRY}_{\mathcal{C}})$ respectively. Let's call this the *start-edge* assumption.

The proof proceeds by induction over the number of edges executed in \mathcal{X}' starting from $n_{\mathcal{X}'}^{s2}$.

$\boxed{M_{\mathcal{X}'}}^{cs} =_{\Sigma_{\ddot{A}'}} M_{\ddot{A}'}$ holds at $n_{\mathcal{X}'}^{s2}$ due to $(\text{ENTRY}_{\ddot{A}'})$, which forms our base case.

Consider a node $n_{\mathcal{X}'}$ such that $\boxed{M_{\mathcal{X}'}}^{cs} =_{\Sigma_{\ddot{A}'}} M_{\ddot{A}'}$ holds at $n_{\mathcal{X}'}$, and let $e_{\mathcal{X}'} = (n_{\mathcal{X}'}, \xrightarrow{\xi_{\ddot{A}'}^t; \xi_{\mathcal{C}}^t} n_{\mathcal{X}'}^t) \in \mathcal{E}_{\mathcal{X}'}$ such that $n_{\mathcal{X}'}^t = (n_{\ddot{A}'}^t, n_{\mathcal{C}}^t) \in \mathcal{N}_{\mathcal{X}'}^{\text{DW}}$ is an error-free node. Two cases for $\xi_{\ddot{A}'}^t$:

1. If $\xi_{\ddot{A}'}^t$ does not contain a **store** instruction, then $\boxed{M_{\mathcal{X}'}}^{cs} =_{\Sigma_{\ddot{A}'}} M_{\ddot{A}'}$ holds trivially at $n_{\mathcal{X}'}^t$.
2. If $\xi_{\ddot{A}'}^t$ contains a **store** instruction, then it cannot modify the allocation state of common regions (\mathcal{B}) in \ddot{A}' (due to lemma 3.5.1). Similarly, $\xi_{\mathcal{C}}^t$ also cannot modify the allocation state of common memory regions in \mathcal{C} (corollary of lemma 3.5.1).

Let α be an address such that a **store** is performed to α in $\xi_{\ddot{A}'}^t$. If $\alpha \in \Sigma_{\ddot{A}'}^{cs}$, then due to (MAC), there must be a **store** to the same address in \mathcal{C} before execution may reach $n_{\mathcal{C}}^t$. Then, due to the global invariants (Disjoint regions in \ddot{A}) and (*cs* and *cl*) (section 3.3.3) and requirement (Equivalence), $\Sigma_{\ddot{A}'}^{cs} \subseteq (\Sigma_{\ddot{A}'}^{Zl}|^v \cup \Sigma_{\mathcal{C}}^{\text{free}}) \cap [\boxed{\text{stk}_e} + 1, \boxed{\text{cs}_e}]$ must hold during the execution of $e_{\mathcal{X}'}$. So, $\alpha \in (\Sigma_{\ddot{A}'}^{Zl}|^v \cup \Sigma_{\mathcal{C}}^{\text{free}}) \cap [\boxed{\text{stk}_e} + 1, \boxed{\text{cs}_e}]$. However, $\alpha \in (\Sigma_{\ddot{A}'}^{Zl}|^v) \setminus (\Sigma_{\ddot{A}'}^{Fw} \cup [\text{esp}, \boxed{\text{cs}_e}])$ is not possible for an error-free node going $\xi_{\ddot{A}'}^t$ due to $(\text{STORE}_{\ddot{A}'})$ with the safety-relaxed semantics. Thus, $\alpha \in (\Sigma_{\mathcal{C}}^{\text{free}} \cap [\boxed{\text{stk}_e} + 1, \boxed{\text{cs}_e}])$ must hold. However, this is not possible for an error-free node going $\xi_{\mathcal{C}}^t$ due to $(\text{STORE}_{\mathcal{C}})$. Thus, by contradiction, a **store** to address $\alpha \in \Sigma_{\ddot{A}'}^{cs}$ is infeasible in $\xi_{\ddot{A}'}^t$. Thus, $\boxed{M_{\mathcal{X}'}}^{cs} =_{\Sigma_{\ddot{A}'}} M_{\ddot{A}'}$ holds at $n_{\mathcal{X}'}^t$.

To generalize beyond the start-edge assumption, we only need to show that for an outgoing edge $e_{\mathcal{X}'} = (n_{\mathcal{X}'}, \xrightarrow{\xi_{\ddot{A}'}^t; \xi_{\mathcal{C}}^t} n_{\mathcal{X}'}^t) \in \mathcal{E}_{\mathcal{X}'}$ such that $n_{\mathcal{X}'}$ is not a non-entry node but

$n_{\mathcal{X}'}^t = (n_{\check{A}'}^t, n_{\check{C}}^t)$ is a non-entry node, $\boxed{M^{cs}} =_{\Sigma_{\check{A}'}} M_{\check{A}'}$ holds at $n_{\mathcal{X}'}^t$. We observe that there must exist a node $n_{\check{A}'}^t$ in $\xi_{\check{A}'}$ where $\boxed{M^{cs}} =_{\Sigma_{\check{A}'}} M_{\check{A}'}$ holds due to (ENTRY $_{\check{A}}$). The rest of the argument remains identical for the path $\xi'_{\check{A}'} = (n_{\check{A}'}^t \rightarrow n_{\check{A}'}^t)$. \square

Theorem 3.5.3 (Soundness of Safety-Relaxed Semantics). *Given $\mathcal{X}' = \check{A}' \boxtimes \check{C}$ that satisfies the search-algorithm requirements, it is possible to construct $\mathcal{X} = \check{A} \boxtimes \check{C}$ that also satisfies the search-algorithm requirements.*

Proof. Construct $\mathcal{X} = \mathcal{X}'$ with some extra edges from nodes in \mathcal{X} to the error-node $(\mathcal{U}_{\check{A}}, \mathcal{U}_{\check{C}})$ such that (Mutex $_{\check{A}}$), (Mutex $_{\check{C}}$) and (SingleIO) are not violated. We later describe what edges are added to \mathcal{X} and why \mathcal{X} continues to satisfy the search-algorithm requirements even after the addition of these edges. It is already possible to see that the structural requirements viz., (Termination), (Similar-speed), (Well-formedness), and (Safety), will hold for \mathcal{X} even after the addition of such edges.

Let $\xi_{\check{A}}$ be a path in \check{A} on which there exists an overlap check $\varphi_l = \text{ov}([p]_w, \Sigma_{\check{A}}^{\text{free}} \cup ((\Sigma_{\check{A}}^{\text{Zl}}|^\nu) \setminus \Sigma_{\check{A}}^{F \cup S}))$ (for triggering \mathcal{U}) due to a (LOAD $_{\check{A}}$) instruction ⁸, In \check{A}' , φ_l is replaced by $\varphi'_l = \text{ov}([p]_w, (\Sigma_{\check{A}}^{\text{Zl}}|^\nu) \setminus (\Sigma_{\check{A}}^F \cup [\text{esp}, \boxed{cs_e}]))$ ⁹ to obtain $\xi'_{\check{A}'}$.

Recall that \check{A} 's translation for (LOAD $_{\check{A}}$) has “if φ_l then halt(\mathcal{U})” ¹⁰ while \check{A}' 's translation has “if φ'_l then halt(\mathcal{U})” ¹¹. Because $\varphi'_l \Rightarrow \varphi_l$, \check{A} may trigger \mathcal{U} when \check{A}' would simply execute the *error-free path* (the path that does not end at an error node) in (LOAD $_{\check{A}}$). Conversely, if \check{A} executes an error-free path (of (LOAD $_{\check{A}}$)) on an initial state σ , then \check{A}' will also execute the same error-free path on σ ¹².

Similarly, let $\varphi_r = \neg(\boxed{M^{cs}} =_{\Sigma_{\check{A}'}} M_{\check{A}'})$ be a check in \check{A} (due to (RET $_{\check{A}}$)), that has been replaced with $\varphi'_r = \text{false}$ in \check{A}' . Again, if \check{A} executes an error-free path of (RET $_{\check{A}}$) on an initial state σ , then \check{A}' will also execute the same error-free path on σ .

Thus, it can be shown through induction that four of the six semantic requirements — (Inductive), (Equivalence), (MAC), (MemEq) — hold on \mathcal{X} if they hold on \mathcal{X}' with $\Phi_{\mathcal{X}} = \Phi_{\mathcal{X}'}$. The common argument in this part of the proof is that the path condition of an error-free path in \mathcal{X} (containing $\neg\varphi_{l,s,r}$ for (LOAD $_{\check{A}}$), (STORE $_{\check{A}}$), and (RET $_{\check{A}}$)) is always stronger than the path condition of an error-free path in \mathcal{X}' (containing $\neg\varphi'_{l,s,r}$).

⁸Or, an overlap check $\varphi_s = \text{ov}([p]_w, \Sigma_{\check{A}}^{\text{free}} \cup ((\Sigma_{\check{A}}^{\text{Zl}}|^\nu) \setminus \Sigma_{\check{A}}^{F \cup S}))$ (for triggering \mathcal{U}) due to a (STORE $_{\check{A}}$) instruction.

⁹ $\varphi'_s = \text{ov}([p]_w, (\Sigma_{\check{A}}^{\text{Zl}}|^\nu) \setminus (\Sigma_{\check{A}}^{F \cup S} \cup [\text{esp}, \boxed{cs_e}]))$ in case of a (STORE $_{\check{A}}$).

¹⁰“if φ_s then halt(\mathcal{U})” for (STORE $_{\check{A}}$)

¹¹“if φ'_s then halt(\mathcal{U})” for (STORE $_{\check{A}}$)

¹² $\varphi'_s \Rightarrow \varphi_s$ in case of (STORE $_{\check{A}}$) — rest of the argument remains identical.

We next show that if (Coverage \mathbf{C}) holds for path $\xi_{\check{A}'}$ starting at node $n_{X'}$ in X' , (Coverage \mathbf{C}) also holds for corresponding path $\xi_{\check{A}}$ starting at corresponding node n_X in X (note: using definition 3.4.1 here). For an edge $e_{X'}^j = (n_{X'} \xrightarrow{\xi_{\check{A}'}^j; \xi_{\check{C}}^j} (n_{\check{A}'}, n_{\check{C}}^j)) \in \mathcal{E}_{X'}$ ($1 \leq j \leq m$), if $\xi_{\check{A}}$ ends at a node $n_{\check{A}}^t \neq \mathcal{U}_{\check{A}}$, then this is easy to show by induction on the number of edges executed on a path: because the path condition of $\xi_{\check{A}}$ in \check{A} is always equal or stronger than the path condition of a corresponding (structurally identical) path $\xi_{\check{A}'}$ in \check{A}' . If (Coverage \mathbf{C}) holds for $\xi_{\check{A}'}$ at a node $n_{X'}$ in X' , it must also hold for $\xi_{\check{A}}$ at the corresponding node n_X in X . We next show that (Coverage \mathbf{C}) holds for a path $\xi_{\check{A}}$ terminating in $\mathcal{U}_{\check{A}}$ ($n_{\check{A}}^t = \mathcal{U}_{\check{A}}$).

Consider a path $\xi_{\check{A}}$ in \check{A} and the corresponding path $\xi_{\check{A}'}$ in \check{A}' . If on a machine state σ both paths $\xi_{\check{A}}$ and $\xi_{\check{A}'}$ transition to $\mathcal{U}_{\check{A}}$ and $\mathcal{U}_{\check{A}'}$, respectively, then because X' satisfies (Coverage \mathbf{C}), σ must execute one of $\xi_{\check{C}}^j$ (for $1 \leq j \leq m$) to completion, thus satisfying (Coverage \mathbf{C}) in X in this case. Thus, we only need to cater to the following two situations where execution on \check{A} may deviate from \check{A}' (i.e., execution does not complete for \check{A} but completes for \check{A}')

- (RET \mathbf{A}): Let $\varphi_r = \neg(\boxed{M^{cs}} =_{\Sigma_{\check{A}}} M_{\check{A}})$ be the check in \check{A} (due to (RET \mathbf{A})), that has been replaced with $\varphi_r' = \text{false}$ in \check{A}' . We show that φ_r must evaluate to **false** in X at procedure return. In other words, the \check{A} path “if φ_r then halt(\mathcal{U})” is infeasible and so \check{A} does not deviate from \check{A}' in this case.

By lemma 3.5.2, $\boxed{M^{cs}} =_{\Sigma_{\check{A}'}} M_{\check{A}'}$ holds at every error-free node $n_{X'} \in \mathcal{N}_{X'}^{\text{DW}}$ and therefore $n_X \in \mathcal{N}_X$. Further, using the (MAC) requirement at the error-free terminating node **exit**, this can be generalized to show that $\boxed{M^{cs}} =_{\Sigma_{\check{A}}} M_{\check{A}}$ holds at the beginning of the path corresponding to (RET \mathbf{A}) in \check{A} . Thus, because the \check{A} path “if φ_r then halt(\mathcal{U})” is infeasible, (Coverage \mathbf{C}) holds trivially for this path at n_X in X .

- (LOAD \check{A}) or (STORE \check{A}): Let $\xi_{\check{A}}^U = (n_{\check{A}} \rightarrow \mathcal{U}_{\check{A}})$ be a path that terminates with $\mathcal{U}_{\check{A}}$.

Lemma 3.5.4. *Let σ be a state at an error-free node $n_X = (n_{\check{A}}, n_{\check{C}}) \in \mathcal{N}_X^{\text{DW}}$ such that $\phi_{n_X}(\sigma)$ holds and σ executes $\xi_{\check{A}}^U = (n_{\check{A}} \rightarrow \mathcal{U}_{\check{A}})$ to completion. Then σ must execute some path $\xi_{\check{C}} = (n_{\check{C}} \rightarrow \mathcal{U}_{\check{C}})$ to completion in \mathbf{C} .*

Proof. Consider the execution of σ on X' starting at $n_{X'} = (n_{\check{A}'}, n_{\check{C}})$, such that $n_{X'}$ in X' is structurally identical to n_X in X . Due to (Mutex \check{A}) and (Coverage \check{A}), there can be only two cases:

1. σ executes some path $\xi_{\ddot{A}'}^x = (n_{\ddot{A}'} \rightarrow \mathcal{U}_{\ddot{A}})$ to completion in \ddot{A}' . In this case, due to (Coverage \mathbb{C}) and (Safety), some $\xi_{\mathbb{C}}^x = (n_{\mathbb{C}} \rightarrow \mathcal{U}_{\mathbb{C}})$ must be executed to completion on σ in \mathbb{C} . In this case, the lemma holds with $\xi_{\mathbb{C}} = \xi_{\mathbb{C}}^x$.
2. σ executes some path $\xi_{\ddot{A}'}^x = (n_{\ddot{A}'} \rightarrow n_{\ddot{A}'}^x)$ to completion in \ddot{A}' , where $n_{\ddot{A}'}^x \neq \mathcal{U}_{\ddot{A}'}$ and $e_{\mathbb{X}'}^{x_v} = (n_{\mathbb{X}'} \xrightarrow{\xi_{\ddot{A}'}^x; \xi_{\mathbb{C}}^{x_v}} n_{\mathbb{X}'}^{x_v}) \in \mathcal{E}_{\mathbb{X}'}$ (for $1 \leq v \leq w$) are $w \geq 1$ edges in \mathbb{X}' , where $n_{\mathbb{X}'}^{x_v} = (n_{\ddot{A}'}^x, n_{\mathbb{C}}^{x_v})$. Because \mathbb{X}' satisfies (Coverage \mathbb{C}), σ must execute a path $\xi_{\mathbb{C}}^{x_v} = (n_{\mathbb{C}} \rightarrow n_{\mathbb{C}}^{x_v})$ to completion in \mathbb{C} , for some $1 \leq v \leq w$. We show by contradiction that $\forall 1 \leq v \leq w : n_{\mathbb{C}}^{x_v} = \mathcal{U}_{\mathbb{C}}$ must hold.

Assume $n_{\mathbb{C}}^{x_v} \neq \mathcal{U}_{\mathbb{C}}$. Let memory access instructions d_1, d_2, \dots, d_k exist on path $\xi_{\ddot{A}'}^x$, such that $\xi_{\ddot{A}'}^x$ deviates from $\xi_{\ddot{A}}^U$ on one of these memory access instructions d_r ($1 \leq r \leq k$), so that $\xi_{\ddot{A}}^U$ transitions to $\mathcal{U}_{\ddot{A}}$ due to φ evaluating to **true** in a check “if φ halt(\mathcal{U})” in a (LOAD \ddot{A}) or (STORE \ddot{A}) in \ddot{A} , while $\xi_{\ddot{A}'}^x$ continues execution to reach $n_{\ddot{A}'}^x \neq \mathcal{U}_{\ddot{A}'}$ due to φ' (safety-relaxed rewrite of φ) evaluating to **false** in a corresponding check “if φ' halt(\mathcal{U})” in \ddot{A}' .

Let $[p]_w$ represent the addresses being accessed by the memory access instruction d_r . It must be true that $\exists \alpha \in [p]_w : \alpha \in \text{comp}(\Sigma_{\ddot{A}'}^{BUFUS})$ if d_r is a load instruction and $\exists \alpha \in [p]_w : \alpha \in \text{comp}(\Sigma_{\ddot{A}'}^{(B \setminus G_r)UFwUS})$ if d_r is a store instruction; this is because φ' evaluates to **false** but φ evaluates to **true** (for load and store instructions). Because \mathbb{X}' satisfies (MAC), the execution of σ starting at $n_{\mathbb{C}}$ must cause all addresses in $[p]_w$ to be accessed before execution can reach $n_{\mathbb{C}}^{x_v}$ in \mathbb{C} (and $n_{\mathbb{X}'}^{x_v}$ in \mathbb{X}'); this is because $\alpha \notin \Sigma_{\ddot{A}'}^{GUF} \cup [\text{esp}, \text{stk}_e]$. Further, because $\xi_{\ddot{A}'}^x$ contains a memory access instruction, due to lemma 3.5.1, both $\xi_{\ddot{A}'}^x$ and $\xi_{\mathbb{C}}^x$ cannot modify the address sets of common regions B . Thus, during the execution of σ starting at $n_{\mathbb{C}}$, the `accessIsSafe $\mathbb{C}_{\tau,a}$` (\cdot) check must necessarily evaluate to **false** and the execution must transition to $\mathcal{U}_{\mathbb{C}}$. This is a contradiction, and so it must be true that $n_{\mathbb{C}}^{x_v} = \mathcal{U}_{\mathbb{C}}$. Hence, the lemma holds in this case with $\xi_{\mathbb{C}} = \xi_{\mathbb{C}}^x = (n_{\mathbb{C}} \rightarrow \mathcal{U}_{\mathbb{C}})$.

□

Using lemma 3.5.4, we enumerate all such paths $\xi_{\mathbb{C}} = (n_{\mathbb{C}} \rightarrow \mathcal{U}_{\mathbb{C}})$ that can be executed in \mathbb{C} if $\xi_{\ddot{A}}^U = (n_{\ddot{A}} \rightarrow \mathcal{U}_{\ddot{A}})$ is executed in \ddot{A} starting at node $n_{\mathbb{X}} \in \mathcal{N}_{\mathbb{X}}$. As described in the proof of lemma 3.5.4, there are only a finite number of such paths. For each such path $\xi_{\mathbb{C}}$, we add an edge $e_{\mathbb{X}}^x = (n_{\mathbb{X}} \xrightarrow{\xi_{\ddot{A}}^U; \xi_{\mathbb{C}}} (\mathcal{U}_{\ddot{A}}, \mathcal{U}_{\mathbb{C}}))$ to $\mathcal{E}_{\mathbb{X}}$ if it does not exist already. (Coverage \mathbb{C}) thus follows from lemma 3.5.4. Further, (Coverage \ddot{A}) also holds for \mathbb{X} because all assembly paths that exist in \mathbb{X}' also exist in \mathbb{X} and additional paths, only potentially feasible in \ddot{A} , are added.

□

Using theorems 3.4.2 and 3.5.3, hereafter, we will use only the safety-relaxed and callers' virtual smallest semantics of the unoptimized IR and assembly procedures. We will continue to refer to the unoptimized IR with the callers' virtual smallest semantics and assembly procedure with the safety-relaxed and callers' virtual smallest semantics as C and A respectively. The corresponding annotated procedure of A will be referred as \check{A} .

Chapter 4

Automatic Construction of a Product-Program

In the previous chapter, we established that a product program $X = \ddot{A} \boxtimes C$ can be used as a witness of refinement from an unoptimized IR procedure C to an annotated assembly procedure \ddot{A} . In this chapter, we describe our algorithm, called DYNAMO, for simultaneous automatic annotation of A , to produce \ddot{A} , and construction of X . The product program X produced by our algorithm is guaranteed to satisfy the thirteen requirements that enable it to be used as a witness of refinement between an input unoptimized IR procedure C and the annotated assembly procedure \ddot{A} , also produced by our algorithm from the input assembly procedure A .

DYNAMO uses a set of internal heuristics for discovering the required annotation for A — we call this *blackbox* setting. In the other *whitebox* setting, DYNAMO is capable of utilizing external *untrusted* hints for annotating A . These hints may be *untrusted* because the annotation is validated by the algorithm. As the hints need not be trusted, they could be generated through lightweight compiler instrumentation which may be incorrect or sourced from potentially inaccurate sources such as debug headers.

4.1 The DYNAMO algorithm

The DYNAMO algorithm takes the transition graphs corresponding to the LLVM_d and assembly procedures (C and A resp.) and an *unroll factor* μ as input and returns, if successful, an annotated \ddot{A} and a product graph $X = \ddot{A} \boxtimes C = (\mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X)$ as output. In

addition, it also identifies an inductive invariant network Φ_X that maps each error-free node $n_X \in \mathcal{N}_X^{\text{DFW}}$ to its node invariant ϕ_{n_X} .

Given enough computational time, DYNAMO is guaranteed to find the required (\ddot{A}, X) if:

- (a) A is a translation of C through, potentially path-specializing, bisimilar transformations up to a maximum unrolling of μ .
- (b) For two or more allocations or procedure calls that reuse stack space in A , their relative order in C is preserved in A .
- (c) An allocation due to `alloca()` is always stack-allocated in A .
- (d) An allocation-containing path is not specialized by the compiler, such that on one specialization stack-allocation is performed and on another register-allocation is performed.
- (e) The desired annotation to \ddot{A} is identifiable either through search heuristics (blackbox mode) or through user-supplied and/or compiler hints (whitebox).
- (f) Our invariant inference procedure is able to identify the required invariant network Φ_X that captures the compiler transformations from C to A .

The restriction on non-bisimilar and de-specializing transformations stated in clause (a) stems from the limitations of the COUNTER algorithm [17] of which our algorithm is a derivative of — the former is fundamental to the algorithm. The restriction on de-specializing transformations is a performance trade-off and can potentially be lifted at the cost of runtime (see §4.4 of [17]). Clause (b) is simply a restatement of the limitation of our refinement definition (section 2.5.4), and clauses (c) and (d) are due to restrictions on a virtual-allocation annotation (section 2.6). We describe our blackbox annotation algorithm that dictates clause (e) in section 4.1.3 and our invariant inference algorithm (for clause (f)) in section 4.2.

DYNAMO constructs the solution incrementally, by relying on the property that for a non-coverage requirement to hold for a fully-annotated \ddot{A} and a fully-constructed X , it must also hold for a partially-annotated \ddot{A} and a partially-constructed subgraph of X rooted at its entry node n_X^s . A partially-constructed X (based on a partially-annotated \ddot{A}), constructed from the entry node n_X^s , that does not meet the non-coverage requirements may be safely discarded without affecting the completeness of the algorithm.

Recall that we defined callers' virtual smallest and safety-relaxed semantics for the procedures \mathbf{C} and \mathbf{A} (sections 3.4 and 3.5) — we assume that the transition graphs \mathbf{C} and \mathbf{A} passed to the algorithm have these semantics. Before beginning the construction of \mathbf{X} (and annotation of \mathbf{A}), we run an intraprocedural, flow-sensitive, field-insensitive points-to dataflow analysis [3] to compute over-approximate states of the β and β_M maps for each node $n_C \in \mathcal{N}_C$ ¹. These sound but over-approximate values, computed at each node n_C , are substituted in to replace all references to β and β_M in \mathbf{C} 's graph. After this substitution, the assignments to β and β_M (e.g., in (LOAD_C) , (STORE_C) , etc.) become vacuous.

Algorithm 1 presents the pseudo-code of the algorithm. The algorithm has two phases. In the first phase (line 4 in algorithm 1), it attempts to correlate the paths in \mathbf{A} with the paths in \mathbf{C} while simultaneously identifying the required annotation for \mathbf{A} . At the successful completion of the first phase, all paths in the original, non-annotated \mathbf{A} are correlated. However, recall that the annotation instructions, alloc_s and alloc_v , have additional paths to error nodes $\mathcal{U}_{\check{\mathbf{A}}}$ and $\mathcal{W}_{\check{\mathbf{A}}}$ (figs. 2.8, 2.10 and 2.11). These paths to error nodes are not correlated in the first phase. The second phase of the algorithm (line 6 in algorithm 1) correlates these additionally introduced (error) paths.

The sub-procedure $\text{constructX}()$, used in both phases, identifies the required correlations and annotation and builds the product program \mathbf{X} incrementally. It assumes the availability of an oracle **chooseFrom** operator, such that $\rho \leftarrow \text{chooseFrom } \vec{\rho}$ chooses a quantity ρ from a finite set $\vec{\rho}$, such that the algorithm is able to complete the refinement proof, if such a choice exists. If the search space is limited, an exhaustive search could be used to implement **chooseFrom**. For larger search spaces, a counterexample-guided best-first search procedure (described in section 4.1.9) is employed to approximate **chooseFrom**. At a high-level, $\text{constructX}()$ systematically visits each uncorrelated path $\xi_{\check{\mathbf{A}}}$ in \mathbf{A} and tries to identify a pathset $\langle \xi \rangle_C$ in \mathbf{C} that can be correlated with $\xi_{\check{\mathbf{A}}}$ such that the non-coverage requirements are not violated. If required, $\xi_{\check{\mathbf{A}}}$ is annotated on-the-fly. We discuss the algorithm in detail in the following sections.

¹Recall that $\beta(x)$ returns the set of regions that a state variable x in \mathbf{C} may point to and $\beta_M(r)$ returns the set of regions that some pointer stored in region r may point to.

Algorithm 1: Automatic construction of X

```

1 Function DYNAMO( $A, C, \mu$ )
2    $\ddot{A} \leftarrow A$ ;    $C \leftarrow \text{pointsToAnalysis}(C)$ ;
3    $N_X \leftarrow \{(n_{\ddot{A}}^s, n_C^s)\}$ ;    $\mathcal{E}_X \leftarrow \emptyset$ ;    $\mathcal{D}_X \leftarrow \emptyset$ ;    $\Phi_X \leftarrow \{(n_{\ddot{A}}^s, n_C^s) \mapsto (\Omega_{\ddot{A}} = \Omega_C)\}$ ;
4   if  $\neg \text{constructX}(\ddot{A}, C, \mu, N_X, \mathcal{E}_X, \mathcal{D}_X, \Phi_X, \text{CORRELATE\_AND\_ANNOTATE})$  then
5     return Failure
6   if  $\neg \text{constructX}(\ddot{A}, C, \mu, N_X, \mathcal{E}_X, \mathcal{D}_X, \Phi_X, \text{CORRELATE\_NEW\_ERROR\_PATHS})$  then
7     return Failure
8   if  $\neg \text{checkCoverageReqs}(N_X, \mathcal{E}_X, \mathcal{D}_X, \Phi_X, \ddot{A}, C)$  then
9     return Failure
10  return Success( $\ddot{A}, (N_X, \mathcal{E}_X, \mathcal{D}_X), \Phi_X$ )
11 end
12 Function  $\text{constructX}(\ddot{A}, C, \mu, N_X, \mathcal{E}_X, \mathcal{D}_X, \Phi_X, \text{phase})$ 
13    $Q_{\ddot{A}} \leftarrow \text{getCutPointsInRPO}(\ddot{A})$ ;
14   foreach  $q_{\ddot{A}}$  in  $Q_{\ddot{A}}$  do
15     foreach  $q_{\ddot{A}}^t$  in  $\text{cutPointSuccessorsRPO}(q_{\ddot{A}}, Q_{\ddot{A}}, \ddot{A})$  do
16       foreach  $\xi_{\ddot{A}}$  in  $\text{getAllSimplePathsBetweenCutPoints}(q_{\ddot{A}}, q_{\ddot{A}}^t, \ddot{A})$  do
17         if  $\text{pathExists}(\xi_{\ddot{A}}, \mathcal{E}_X)$  then
18           continue
19         if  $\text{pathIsInfeasible}(\xi_{\ddot{A}}, N_X, \Phi_X)$  then
20           continue
21          $\langle \xi \rangle_C \leftarrow \text{chooseFrom correlatedPathsInCOpts}(\xi_{\ddot{A}}, \mu, N_X, \mathcal{E}_X, \ddot{A}, C)$ ;
22         foreach  $\xi_C$  in  $\langle \xi \rangle_C$  do
23           if  $\text{phase} = \text{CORRELATE\_AND\_ANNOTATE}$  then
24              $(\ddot{A}, \xi_{\ddot{A}}) \leftarrow \text{chooseFrom asmAnnotOpts}(\xi_{\ddot{A}}, \xi_C, \ddot{A}, C)$ ;
25           end
26            $\vec{\xi}'_{\ddot{A}} \leftarrow \text{breakIntoSingleIOPaths}(\xi_{\ddot{A}})$ ;
27            $\vec{\xi}'_C \leftarrow \text{breakIntoSingleIOPaths}(\xi_C)$ ;
28            $\vec{\xi}^*_{\ddot{A}}, \vec{\xi}^*_C \leftarrow \text{trimToMatchPathToErrorNode}(\vec{\xi}'_{\ddot{A}}, \vec{\xi}'_C)$ ;
29           if  $\neg \text{haveSimilarStructure}(\vec{\xi}^*_{\ddot{A}}, \vec{\xi}^*_C)$  then
30             return Failure
31           foreach  $\xi'_{\ddot{A}} = (n_{\ddot{A}} \rightarrow n_{\ddot{A}}^t), \xi'_C = (n_C \rightarrow n_C^t)$  in  $\text{zip}(\vec{\xi}^*_{\ddot{A}}, \vec{\xi}^*_C)$  do
32              $e_X \leftarrow (\xi'_{\ddot{A}}; \xi'_C)$ ;    $n'_X \leftarrow (n_{\ddot{A}}^t, n_C^t)$ ;
33             if  $\text{addingEdgeWillCreateEmptyCCycle}(N_X, \mathcal{E}_X, e_X)$  then
34               return Failure
35              $\mathcal{E}_X \leftarrow \mathcal{E}_X \cup \{e_X\}$ ;    $N_X \leftarrow N_X \cup \{n'_X\}$ ;
36              $\mathcal{D}_X \leftarrow \text{addDetMappings}(e_X, \mathcal{D}_X)$ ;
37              $\Phi_X \leftarrow \text{inferInvariantsAndCounterexamples}(n'_X, N_X, \mathcal{E}_X, \mathcal{D}_X, \Phi_X, \ddot{A}, C)$ ;
38             if  $\neg \text{checkSemanticReqsExceptCoverage}(N_X, \mathcal{E}_X, \mathcal{D}_X, \Phi_X, \ddot{A}, C)$  then
39               return Failure
40           end
41         end
42       end
43     end
44   end
45   return Success
46 end

```

4.1.1 Enumerating \mathbf{A} paths

Cut-points and Simple Paths

We first define some useful predicates. Let $P \in \{\mathbf{C}, \mathbf{\ddot{A}}\}$.

Definition 4.1.1 ($\text{io}(n_P)$). $\text{io}(n_P)$ evaluates to **true** iff n_P is either a source or sink node of an I/O path.

Definition 4.1.2 ($\text{term}(n_P)$). $\text{term}(n_P)$ evaluates to **true** iff n_P is a terminating node.

We define an ordered set of nodes $Q_P \subseteq \mathcal{N}_P$, called the *cut-points* in procedure P , such that $Q_P \supseteq \{n_P : n_P \in \mathcal{N}_P \wedge (n_P = n_P^s \vee \text{io}(n_P) \vee \text{term}(n_P))\}$ and the maximum length of a path between two nodes in Q_P (not containing any other intermediate node that belongs to Q_P) is finite. The algorithm to identify Q_P for a procedure P , $\text{getCutPointsInRPO}(P)$, first initializes $Q_P := \{n_P : n_P \in \mathcal{N}_P \wedge (n_P = n_P^s \vee \text{io}(n_P) \vee \text{term}(n_P))\}$, and then identifies all cycles in the transition graph P that do not already contain a cut-point; for each such cycle, the first node belonging to that cycle in reverse postorder is added to Q_P . The cut-points Q_P returned by $\text{getCutPointsInRPO}(P)$ are arranged in a reverse postorder (RPO).

A *simple path* ($q_P \rightarrow q_P^t$) is a non-empty path connecting two cut-points $q_P, q_P^t \in Q_P$ and not containing any other cut-point as an intermediate node; q_P^t is called a *cut-point successor* of q_P . By definition, a simple path must be finite.

The $\text{cutPointSuccessorsRPO}(q_P, Q_P, P)$ function returns the cut-point successors of q_P ordered in reverse postorder. The error node successors of q_P , if any, are ordered after the error-free successors. Further, the error node successor \mathcal{W}_P (if it exists) is arranged at the very end, after all other cut-point successors. This property is expected during path enumeration.

The $\text{getAllSimplePathsBetweenCutPoints}(q_P, q_P^t, P)$ function returns *all* simple paths of the form $q_P \rightarrow q_P^t$, for $q_P, q_P^t \in Q_P$. The returned paths are mutually exclusive by construction.

Reverse postorder enumeration of \mathbf{A} paths

The procedure $\text{constructX}()$ first identifies a set of cut-points $Q_{\mathbf{\ddot{A}}}$ of procedure $\mathbf{\ddot{A}}$ using $\text{getCutPointsInRPO}(\mathbf{\ddot{A}})$ (line 13 in algorithm 1). Then, for each cut-point $q_{\mathbf{\ddot{A}}} \in Q_{\mathbf{\ddot{A}}}$,

Algorithm 2: Pseudo-code of the $pathIsInfeasible()$ procedure.

```

1 Function  $pathIsInfeasible(\xi_{\ddot{A}} = (q_{\ddot{A}} \rightarrow q_{\ddot{A}}^t), \mathcal{N}_X, \Phi_X)$ 
2    $n_X \leftarrow \text{getXNode}(q_{\ddot{A}}, \mathcal{N}_X)$  ; // unique X node corresponding to source of  $\xi_{\ddot{A}}$ 
3   if Hoare triple  $\{\phi_{n_X}\}(\xi_{\ddot{A}}; \epsilon)\{\text{false}\}$  holds then
4     return true
5   else
6     return false
7   end
8 end

```

visited in reverse postorder, it identifies the set of cut-point successors of $q_{\ddot{A}}^t$ using $\text{cutPointSuccessorsRPO}(q_{\ddot{A}}, \mathcal{Q}_{\ddot{A}}, \ddot{A})$ (line 15 in algorithm 1). These cut-point successors are visited in reverse postorder, and all simple paths $\vec{\xi}_{\ddot{A}}$ between $q_{\ddot{A}}$ and its cut-point successor $q_{\ddot{A}}^t$ are enumerated through $\text{getAllSimplePathsBetweenCutPoints}(q_{\ddot{A}}, q_{\ddot{A}}^t, \ddot{A})$ (line 16 in algorithm 1).

The paths $\vec{\xi}_{\ddot{A}}$ are enumerated in reverse postorder (i.e., the source and sink nodes are enumerated in reverse postorder) so that annotation (if any) over a path $\xi_{\ddot{A}} = (q_{\ddot{A}} \rightarrow q_{\ddot{A}}^t) \in \vec{\xi}_{\ddot{A}}$ is performed *before* the address sets (potentially modified by annotation) are used in successor instructions of $q_{\ddot{A}}^t$. This property ensures consistency in the invariant network Φ_X and enables incremental building of the product program X .

Given a simple path $\xi_{\ddot{A}} = (q_{\ddot{A}} \rightarrow q_{\ddot{A}}^t) \in \vec{\xi}_{\ddot{A}}$, $pathIsInfeasible(\xi_{\ddot{A}}, \mathcal{N}_X, \Phi_X)$ returns **true** iff $\xi_{\ddot{A}}$ is infeasible at every node $n_X = (q_{\ddot{A}}, -) \in \mathcal{N}_X$; our construction ensures there can be at most one $n_X = (q_{\ddot{A}}, -)$ for each $q_{\ddot{A}} \in \mathcal{Q}_{\ddot{A}}$. An infeasible path is not considered for correlation (line 19 in algorithm 1). Algorithm 2 shows the pseudo-code of $pathIsInfeasible()$.

For an \ddot{A} path $\xi_{\ddot{A}}$, $pathExists(\xi_{\ddot{A}}, \mathcal{E}_X)$ returns **true** iff $\xi_{\ddot{A}}$ is already correlated with some $\xi_C = (q_C \rightarrow q_C^t)$ in \mathcal{E}_X (i.e., $\exists e_X : e_X = ((q_{\ddot{A}}, q_C) \xrightarrow{\xi_{\ddot{A}}; \xi_C} (q_{\ddot{A}}^t, q_C^t)) \in \mathcal{E}_X$ holds). Because the same $constructX()$ procedure is invoked in both phases, the use of $pathExists()$ in line 17 of algorithm 1 is an optimization to avoid correlating the same paths again in the second phase. In the first call to $constructX()$, when $phase = \text{CORRELATE_AND_ANNOTATE}$, $pathExists(\xi_{\ddot{A}}, \mathcal{E}_X)$ will always return **false** as the enumeration returns fresh, uncorrelated paths. However, in the second call to $constructX()$, when $phase = \text{CORRELATE_NEW_ERROR_PATHS}$, the algorithm is only correlating the newly introduced error paths, and $pathExists(\xi_{\ddot{A}}, \mathcal{E}_X)$ will return **false** only for the uncorrelated paths to error nodes introduced due to annotation of \ddot{A} .

Trade-off between completeness and efficiency

DYNAMO attempts to correlate each cut-point in $Q_{\ddot{A}}$ with a cut point in Q_C . To allow maximum transformations, $Q_{\ddot{A}}$ should be as small as possible. On the other hand, a smaller $Q_{\ddot{A}}$ could result in potentially longer paths between cut-points; a Hoare triple over longer paths could be potentially harder for SMT solvers to reason about due to larger expression sizes. Our algorithm of $getCutPointsInRPO(P)$ returns the minimum set of cut-points for a procedure P .

4.1.2 Correlating C paths

Given an \ddot{A} path $\xi_{\ddot{A}} = (q_{\ddot{A}} \rightarrow q_{\ddot{A}}^t)$, $correlatedPathsInCOpts(\xi_{\ddot{A}}, \mu, \dots)$ identifies a non-empty set of *options* for candidate pathset² $[\langle \xi \rangle_C]$ in C that can potentially be correlated with $\xi_{\ddot{A}}$. The **chooseFrom** operator is used to choose a pathset $\langle \xi \rangle_C$ from the returned options $[\langle \xi \rangle_C]$. All paths in $\langle \xi \rangle_C$ must originate at a unique error-free cut-point q_C such that $(q_{\ddot{A}}, q_C) \in \mathcal{N}_X$. The enumeration by $correlatedPathsInCOpts()$ ensures that there will be exactly one such $(q_{\ddot{A}}, q_C)$ in \mathcal{N}_X . Paths in $\langle \xi \rangle_C$ may have different sinks, however. Unlike $\xi_{\ddot{A}}$, a path $\xi_C \in \langle \xi \rangle_C$ need not be a simple path, and can visit a node $n_C \in \mathcal{N}_C$ up to μ times — this enables DYNAMO to handle transformations such as loop peeling, unrolling, etc.

Before we describe $correlatedPathsInCOpts()$, we define a \mathcal{U} -maximal pathset.

Definition 4.1.3 (\mathcal{U} -maximal pathset). *A pathset $\langle \xi \rangle_P$ of procedure P is \mathcal{U} -maximal iff whenever there exists a path $\xi'_P = \epsilon$ or $\xi'_P = (q_P \rightarrow q_P^t)$ such that $\xi'_P \preceq \xi_P$ ³ for $\xi_P \in \langle \xi \rangle_P$ and $\xi_P^u = \xi'_P \cdot (q_P^t \rightarrow \mathcal{U}_P)$ for some edge $(q_P^t \rightarrow \mathcal{U}_P) \in \mathcal{E}_P$, $\xi_P^u \in \langle \xi \rangle_P$ holds.*

Informally, a \mathcal{U} -maximal pathset $\langle \xi \rangle_P$ is a set of paths where the \mathcal{U}_P -going path of a potentially \mathcal{U} -producing instruction is included (as part of a path) in $\langle \xi \rangle_P$, if the instruction is included (as part of a path) in $\langle \xi \rangle_P$. The union of two \mathcal{U} -maximal pathsets, if it is a pathset, is also a \mathcal{U} -maximal pathset.

The $correlatedPathsInCOpts()$ procedures enumerates \mathcal{U} -maximal pathsets as candidate options. Enumerating \mathcal{U} -maximal pathsets is essential for a fast SMT encoding (see chapter 5). The key property that \mathcal{U} -maximality enables is that the pathset condition

²Recall that a pathset is a set of pairwise mutually-exclusive paths originating at the same node (definition 3.1.2).

³Recall that $\xi'_P \preceq \xi_P$ iff ξ'_P is a prefix of ξ_P (section 3.1).

(disjunction of the path conditions of the individual paths) does not depend on the \mathcal{U} -triggering condition of an instruction in the pathset.

The DYNAMO algorithm requires *correlatedPathsInCOpts()* to behave differently based on the sink node $q_{\bar{A}}^t$ of $\xi_{\bar{A}}$. We present the associated requirements first and then later discuss our implementation.

- If $q_{\bar{A}}^t \notin \{\mathcal{U}_{\bar{A}}, \mathcal{W}_{\bar{A}}\}$, i.e. $q_{\bar{A}}^t$ is an error-free cut-point node, *correlatedPathsInCOpts()* must return candidates where each candidate $\langle \xi \rangle_{\mathcal{C}}$ is a \mathcal{U} -maximal non-empty pathset such that each path $\xi_{\mathcal{C}} \in \langle \xi \rangle_{\mathcal{C}}$ starts at $q_{\mathcal{C}}$ and either:
 - (a) ends at a unique error-free cut-point node, say $q_{\mathcal{C}}^t$, i.e., all paths $\xi_{\mathcal{C}} \in \langle \xi \rangle_{\mathcal{C}}$ ending at an error-free node end at $q_{\mathcal{C}}^t$, or
 - (b) ends at error node $\mathcal{U}_{\mathcal{C}}$.

Because we restrict all paths in $\langle \xi \rangle_{\mathcal{C}}$ with an error-free sink node to have the same sink $q_{\mathcal{C}}^t$, for a $q_{\bar{A}}^t$ node, this creates exactly one node $n_{\mathcal{X}}^t = (q_{\bar{A}}^t, q_{\mathcal{C}}^t) \in \mathcal{N}_{\mathcal{X}}$.

If $\epsilon \in \langle \xi \rangle_{\mathcal{C}}$, then no other path to an error-free node can be present in $\langle \xi \rangle_{\mathcal{C}}$. In other words, ϵ is treated as a path that starts and ends at $q_{\mathcal{C}}$.

- If $q_{\bar{A}}^t = \mathcal{U}_{\bar{A}}$, *correlatedPathsInCOpts()* must return candidates where each candidate $\langle \xi \rangle_{\mathcal{C}}$ is a \mathcal{U} -maximal non-empty pathset such that each path $\xi_{\mathcal{C}} \in \langle \xi \rangle_{\mathcal{C}}$ starts at $q_{\mathcal{C}}$ and ends at $\mathcal{U}_{\mathcal{C}}$.

correlatedPathsInCOpts() must return a pathset option containing the empty path, $\langle \xi \rangle_{\mathcal{C}} = \{\epsilon\}$, if the aforementioned enumeration is not possible, e.g., if no path from $q_{\mathcal{C}}$ to $\mathcal{U}_{\mathcal{C}}$ exists in \mathcal{C} . Notice that this choice is never a valid correlation option for a $\mathcal{U}_{\bar{A}}$ -going path because it does not satisfy the (Safety) requirement. However, because *correlatedPathsInCOpts()* must return a non-empty set of options, we simply choose $\langle \xi \rangle_{\mathcal{C}} = \{\epsilon\}$ as a sentinel value so that DYNAMO may fail as early as possible.

- If $q_{\bar{A}}^t = \mathcal{W}_{\bar{A}}$, *correlatedPathsInCOpts()* must return candidates where each candidate $\langle \xi \rangle_{\mathcal{C}}$ is a \mathcal{U} -maximal non-empty pathset such that each path $\xi_{\mathcal{C}} \in \langle \xi \rangle_{\mathcal{C}}$ either has an error-free sink or has the error node $\mathcal{U}_{\mathcal{C}}$ as sink. Unlike the $q_{\bar{A}}^t \notin \{\mathcal{U}_{\bar{A}}, \mathcal{W}_{\bar{A}}\}$ case, the paths to error-free nodes are not required to have the same, unique sink node. In this case, not restricting the enumeration to a unique error-free sink node facilitates a simultaneous proof of (MAC) and (Coverage \mathcal{C}). We demonstrate this through an example.

Consider the C and (abstracted) assembly code fragments shown below:

<pre> q_C: ... // load/store to hp/cl if (cc) { q_C^{t₁}: fcall1() ... } else { q_C^{t₂}: fcall2() ... } </pre>	<pre> q_Ä: ... // load/store to hp/cl ... q_Ä^t: esp = ... // ... halt(W) if (ca) { q_Ä^{t₁}: fcall1() ... } else { q_Ä^{t₂}: fcall2() ... } </pre>
--	---

Assume that DYNAMO has already correlated $(q_C, q_{\ddot{A}})$, $(q_C^{t_1}, q_{\ddot{A}}^{t_1})$, and $(q_C^{t_2}, q_{\ddot{A}}^{t_2})$, and is now trying to correlate $\xi_{\ddot{A}} = (q_{\ddot{A}} \rightarrow q_{\ddot{A}}^t)$ such that $q_{\ddot{A}}^t = \mathcal{W}_{\ddot{A}}$. Consider the different possibilities for a pathset $\langle \xi \rangle_C$, identified by *correlatedPathsInCOpts()*:

- If $\langle \xi \rangle_C = \{\epsilon, (q_C \rightarrow \mathcal{U}_C)\}$, such that $\xi_C = \epsilon$ is correlated with $\xi_{\ddot{A}}$ producing $e_X = ((q_{\ddot{A}}, q_C) \xrightarrow{\xi_{\ddot{A}}; \epsilon} (\mathcal{W}_{\ddot{A}}, q_C))$, then e_X would fail to uphold the (MAC) condition because the memory access to *hp* or *cl* (*load/store to hp/cl*) would remain unmatched in $\xi_{\ddot{A}}$ (recall that (MAC) requires a *hp* or *cl* access in $\xi_{\ddot{A}}$ to be *matched* in ξ_C ; see also the `checkMAC()` algorithm in section 4.1.6).
- If $\langle \xi \rangle_C = \{(q_C \rightarrow q_C^{t_1}), (q_C \rightarrow \mathcal{U}_C)\}$ or $\langle \xi \rangle_C = \{(q_C \rightarrow q_C^{t_2}), (q_C \rightarrow \mathcal{U}_C)\}$, i.e., $\langle \xi \rangle_C$ includes paths from either of the two already correlated pathsets (but not both), then (Coverage_C) would fail to hold because the correlated path $\xi_C = (q_C \rightarrow q_C^{t_1})$ (or $(q_C \rightarrow q_C^{t_2})$) may not execute to completion when $\xi_{\ddot{A}} = (q_{\ddot{A}} \rightarrow q_{\ddot{A}}^t)$ is executed to completion, i.e., `cc` (or `¬cc`) may not necessarily hold even if $\xi_{\ddot{A}}$, which terminates before reaching corresponding `if` statement, executes to completion.
- If $\langle \xi \rangle_C = \{(q_C \rightarrow q_C^{t_1}), (q_C \rightarrow q_C^{t_2}), (q_C \rightarrow \mathcal{U}_C)\}$, i.e., $\langle \xi \rangle_C$ includes paths with different sink nodes ($q_C^{t_1}$ and $q_C^{t_2}$), then it is easy to see that both (MAC) and (Coverage_C) can simultaneously hold.

Note that both (MAC) and (Coverage_C) can simultaneously hold when $\xi_{\ddot{A}} = (q_{\ddot{A}} \rightarrow q_{\ddot{A}}^{t_1})$ is correlated with $\xi_C = (q_C \rightarrow q_C^{t_1})$ or $\xi_{\ddot{A}} = (q_{\ddot{A}} \rightarrow q_{\ddot{A}}^{t_2})$ is correlated with $\xi_C = (q_C \rightarrow q_C^{t_2})$. In both cases, (Coverage_C) can hold if `cc` (`¬cc`) and `ca` (`¬ca`) evaluate to identical value in a lockstep execution of ξ_C and $\xi_{\ddot{A}}$. Thus, not restricting the candidates in $\langle \xi \rangle_C$ to a single sink node is required for satisfying (MAC) in the particular case of $q_{\ddot{A}}^t = \mathcal{W}_{\ddot{A}}$.

As *correlatedPathsInCOpts()* identifies options of a pathset, which, by definition, is a set of mutually-exclusive paths, the (Mutex_C) requirement is satisfied by construction.

Further, because a path $\xi_C = (q_C \rightarrow \mathcal{W}_C)$ is never returned as part of any candidate, (Well-formedness) is also satisfied.

With an appropriate unroll-factor μ , this path enumeration strategy supports path specializing compiler transformations like loop peeling, unrolling, splitting, unswitching, etc., but does not support a path de-specializing transformation like loop re-rolling. A path de-specialization transformation requires a single $\xi_{\ddot{A}} = (q_{\ddot{A}} \rightarrow q_{\ddot{A}}^t)$ to be correlated with a set of paths $\{\xi_C^1, \xi_C^2, \dots, \xi_C^m\}$ with potentially different end-points such that $\forall_{1 \leq j \leq m} : \xi_C^j = (q_C \rightarrow q_C^j)$ and $\forall_{1 \leq j_1 < j_2 \leq m} : q_C^{j_1} \neq q_C^{j_2}$. Such a construction would require an error-free cut-point in \ddot{A} to be correlated with multiple error-free cut-points in C — recall that we correlate an error-free cut-point in \ddot{A} with exactly one error-free cut-point in C . For comparison, a path specialization transformation requires a cut-point q_C in C to be correlated with multiple cut-points in \ddot{A} (which our algorithm supports).

correlatedPathsInCOpts() algorithm

Algorithm 3 shows the pseudo-code of the *correlatedPathsInCOpts()* procedure. The procedure accepts the path $\xi_{\ddot{A}} = (q_{\ddot{A}} \rightarrow q_{\ddot{A}}^t)$, unroll-factor μ , nodes \mathcal{N}_X and edges \mathcal{E}_X of X , and the two procedures \ddot{A} and C as input parameters and produces a set of pathsets (pathset options) as output.

correlatedPathsInCOpts() presumes the existence of a pathset enumeration sub-procedure *enumPathsetsTillUnroll()*. The *enumPathsetsTillUnroll*(q_P, q_P^t, μ, P) procedure returns a set of pathsets $\overrightarrow{\langle \xi \rangle}_P$ such that for each pathset $\langle \xi \rangle_P \in \overrightarrow{\langle \xi \rangle}_P$ and each path $\xi_P \in \langle \xi \rangle_P$, ξ_P has q_P as source node and q_P^t as sink node, i.e., $\xi_P = (q_P \rightarrow q_P^t)$ and maximum unrolling of a node in ξ_P is bounded by μ . Thus, *enumPathsetsTillUnroll()* returns options for a mutually exclusive set of paths such that each path in the set starts at q_P and ends at q_P^t and a node appears at most μ times in a path. The *getFullPathsetAtAllDelta*(q_P, q^t, P, μ) procedure from COUNTER (§3.11 in [17]) satisfies the requirements for *enumPathsetsTillUnroll()* and is used in our implementation. *getFullPathsetAtAllDelta*(q_P, q^t, P, μ) returns a set of (μ, δ) -unrolled full pathsets that are constructed such that paths in a returned pathset are disjoint from paths in rest of the pathsets. We will elaborate on (μ, δ) -unrolled full pathset when we talk about representation of a pathset later in this section.

correlatedPathsInCOpts() begins by identifying the unique node q_C such that $n_X = (q_{\ddot{A}}, q_C) \in \mathcal{N}_X$ (recall that q_C is guaranteed to exist and be unique). It then constructs

Algorithm 3: Pseudo-code of the *correlatedPathsInCOpts()* algorithm.

```

1 Function correlatedPathsInCOpts( $\xi_{\bar{A}} = (q_{\bar{A}} \rightarrow q_{\bar{A}}^t), \mu, \mathcal{N}_X, \mathcal{E}_X, \bar{A}, C$ )
2    $n_X = (q_{\bar{A}}, q_C) \leftarrow \text{getXNode}(q_{\bar{A}}, \mathcal{N}_X)$ ; // unique X node corresponding to  $q_{\bar{A}}$ 
3    $ret \leftarrow \{\text{mkPathset}(\{\epsilon\})\}$ ; // pathset with empty path is always an option
4   if  $q_{\bar{A}}^t = \mathcal{U}_{\bar{A}}$  then // enumerate options for  $\mathcal{U}_C$ -going paths
5      $\langle \xi \rangle_C \leftarrow \text{enumPathsetsTillUnroll}(q_C, \mathcal{U}_C, \mu, C)$ ;
6     foreach  $\langle \xi \rangle_C \in \overrightarrow{\langle \xi \rangle_C}$  do
7        $ret \leftarrow ret \cup \text{getAllUMaximalSubsets}(\langle \xi \rangle_C, C)$ ;
8     end
9   else if  $q_{\bar{A}}^t = \mathcal{W}_{\bar{A}}$  then // enumerate combinations of already correlated paths
10     $\vec{e}_X^o \leftarrow \text{getOutgoingEdges}(n_X, \mathcal{E}_X)$ ;
11     $\langle \xi \rangle_C \leftarrow \text{prjCPathsets}(\vec{e}_X^o)$ ; // project C pathsets out of outgoing edges
12     $ret \leftarrow ret \cup \overrightarrow{\langle \xi \rangle_C}$ ;
13    /* include the maximal combinations of the pathsets as well */
14    foreach  $\langle \xi \rangle_C \in \overrightarrow{\langle \xi \rangle_C}$  do
15       $r \leftarrow \langle \xi \rangle_C$ ;
16      foreach  $\langle \xi \rangle'_C \in \overrightarrow{\langle \xi \rangle_C}$  do
17        if  $\text{unionRemainsMutex}(r, \langle \xi \rangle'_C)$  then
18           $r \leftarrow \text{unionPathsets}(r, \langle \xi \rangle'_C)$ ;
19        end
20       $ret \leftarrow ret \cup r$ ;
21    end
22  else //  $q_{\bar{A}}^t \notin \{\mathcal{U}_{\bar{A}}, \mathcal{W}_{\bar{A}}\}$ : enumerate  $\mu$ -unrolled paths to reachable cut-points
23    foreach  $q_C^t \in \text{getAllNonErrorReachableCutPoints}(q_C, C)$  do
24       $\langle \xi \rangle_C \leftarrow \text{enumPathsetsTillUnroll}(q_C, q_C^t, \mu, C)$ ;
25      foreach  $\langle \xi \rangle_C^t \in \overrightarrow{\langle \xi \rangle_C}$  do
26         $ret \leftarrow ret \cup \text{getAllUMaximalSubsets}(\langle \xi \rangle_C^t, C)$ ;
27      end
28    end
29  return  $ret$ 
30 Function getAllUMaximalSubsets( $\langle \xi \rangle_P, P$ )
31    $ret \leftarrow \emptyset$ ;
32   foreach  $\langle \xi \rangle'_P \in \text{powerset}(\langle \xi \rangle_P)$  do
33      $ret \leftarrow ret \cup \text{mkUMaximal}(\langle \xi \rangle'_P, P)$ ;
34   end
35   return  $ret$ 
36 end

```

(using $\text{mkPathset}(\{\epsilon\})$) the pathset containing the empty path ϵ (line 3 in algorithm 3) as an output option. This ensures that the procedure returns at least one candidate option — even though it may not be the correct candidate in all cases, e.g., when $q_{\bar{A}}^t = \mathcal{U}_{\bar{A}}$. The other output options returned by *correlatedPathsInCOpts()* are determined by $q_{\bar{A}}^t$.

When $q_{\bar{A}}^t = \mathcal{U}_{\bar{A}}$ (line 4 in algorithm 3), the procedure uses $\text{enumPathsetsTillUnroll}()$ to enumerate a set of pathsets $\overrightarrow{\langle \xi \rangle_C}$ containing μ -unrolled paths to \mathcal{U}_C . For a pathset

$\langle \xi \rangle_C \in \overrightarrow{\langle \xi \rangle_C}$, the *getAllUMaximalSubsets*($\langle \xi \rangle_C, C$) sub-procedure (defined in lines 30 to 36 of algorithm 3) computes a \mathcal{U} -maximal pathset for each subset of $\langle \xi \rangle_C$ — computing over each subset of $\langle \xi \rangle_C$ instead of just $\langle \xi \rangle_C$ increases generality at the cost of exponential increase in the number of returned candidates. The ranking and pruning strategies proposed by COUNTER [17] have been demonstrated to help in effectively navigating such large search-spaces. The *mkUMaximal*($\langle \xi \rangle_P, P$) function, used in line 33 of algorithm 3, returns the smallest \mathcal{U} -maximal pathset $\langle \xi \rangle_P^U$ such that $\langle \xi \rangle_P \subseteq \langle \xi \rangle_P^U$. The computed \mathcal{U} -maximal pathsets are added as candidate output options.

When $q_{\bar{A}}^t = \mathcal{W}_{\bar{A}}$ (line 9 in algorithm 3), the procedure returns a set of pathset options derived from the already correlated paths in \mathcal{E}_X . Recall that *cutPointSuccessorsRPO*(q_P, \dots) orders the cut-point successor $\mathcal{W}_{\bar{A}}$ of q_P after all other cut-point successors of q_P . Due to this, $\xi_{\bar{A}} = (q_{\bar{A}} \rightarrow \mathcal{W}_{\bar{A}})$ is considered for correlation only after all paths of the form $\xi_{\bar{A}}^{\mathcal{X}} = (q_{\bar{A}} \rightarrow q_{\bar{A}}^{\mathcal{X}})$ (for $q_{\bar{A}}^{\mathcal{X}} \neq \mathcal{W}_{\bar{A}}$) have been correlated. A pathset $\langle \xi \rangle_C^{\mathcal{X}}$ already correlated with $\xi_{\bar{A}}^{\mathcal{X}}$ ($(\xi_{\bar{A}}^{\mathcal{X}}; \langle \xi \rangle_C^{\mathcal{X}}) \in \mathcal{E}_X$) is used as a candidate option.

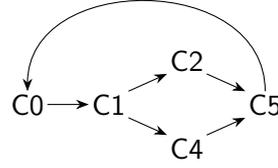
To increase coverage, the procedure also includes the maximal mutually-exclusive combinations of the pathsets $\overrightarrow{\langle \xi \rangle_C^{\mathcal{X}}}$ — computed using iterative merging in lines 13 to 20 of algorithm 3 — as candidate options. For example, if there are two pathsets $\overrightarrow{\langle \xi \rangle_C^{\mathcal{X}^1}}$ and $\overrightarrow{\langle \xi \rangle_C^{\mathcal{X}^2}}$ correlated with paths of the form $\xi_{\bar{A}}^{\mathcal{X}}$ (as described above), such that a union of $\overrightarrow{\langle \xi \rangle_C^{\mathcal{X}^1}}$ and $\overrightarrow{\langle \xi \rangle_C^{\mathcal{X}^2}}$ is also a pathset $\overrightarrow{\langle \xi \rangle_C^{\mathcal{X}^u}}$, then all three pathsets $\overrightarrow{\langle \xi \rangle_C^{\mathcal{X}^1}}$, $\overrightarrow{\langle \xi \rangle_C^{\mathcal{X}^2}}$, and $\overrightarrow{\langle \xi \rangle_C^{\mathcal{X}^u}}$ will be included in the candidate options. An example where this is useful was presented earlier in this section when requirements associated with *correlatedPathsInCOpts*() were described. Note that each of the returned candidate pathset will be \mathcal{U} -maximal by construction.

When $q^t \notin \{\mathcal{U}_{\bar{A}}, \mathcal{W}_{\bar{A}}\}$ (line 21 in algorithm 3), the procedure first identifies all error-free cut-points Q_C^{UW} ($Q_C^{UW} \cap \{\mathcal{U}_C, \mathcal{W}_C\} = \emptyset$) reachable from q_C using the function *getAllNonErrorReachableCutPoints*() (line 22 in algorithm 3). For each cut-point $q_C^t \in Q_C^{UW}$, pathsets $\overrightarrow{\langle \xi \rangle_C}$ with paths from q_C to q_C^t are enumerated using *enumPathsetsTillUnroll*(). For each enumerated pathset $\langle \xi \rangle_C^t \in \overrightarrow{\langle \xi \rangle_C}$, the procedure identifies a \mathcal{U} -maximal pathset for each subset of $\langle \xi \rangle_C^t$ and adds it to the set of candidate options (lines 24 to 26 in algorithm 3).

```

C0:  for (i = 0; i < n; ++i) {
C1:    if (c[i]) {
C2:      x += a[i];
C3:    } else {
C4:      x += b[i];
C5:    }
C6:  }

```



(a) C source fragment.

(b) Abbreviated graph for fig. 4.1a.

Figure 4.1: C source fragment and its abbreviated control-flow graph.

Representation of a pathset

The number of enumerated paths can be exponential in the number of procedure nodes and unroll factor μ . COUNTER [17] suggests use of a compact series-parallel digraph representation, called *SP-graph* representation in [42], for efficiently representing a (μ, δ) -unrolled full pathset — a (μ, δ) -unrolled full pathset is a maximal set of pairwise mutually-exclusive paths such that all paths have same source node and same sink node and no node in a path is repeated more than μ times in a path and the sink node repeats exactly δ times in each path ($\delta \leq \mu$). The SP-graph representation of a pathset is a structured serial (\cdot) and parallel ($+$) combination defined by the grammar $SP ::= \epsilon \mid e \mid SP \cdot SP \mid SP + SP$, where e is an edge and ϵ represents the empty path. The paths represented by an SP-graph can be enumerated by a recursive traversal of its structure. For example, consider the C procedure fragment and its abbreviated control-flow graph in fig. 4.1. A $(1, 1)$ -unrolled full pathset of loop path from C0, say $FP_{C0 \rightsquigarrow C0}^{1,1}$, has paths corresponding to both branches of if at C1, i.e., $FP_{C0 \rightsquigarrow C0}^{1,1} = \{C0 \rightarrow C1 \rightarrow C2 \rightarrow C5 \rightarrow C0, C0 \rightarrow C1 \rightarrow C4 \rightarrow C5 \rightarrow C0\}$. The SP-graph of $FP_{C0 \rightsquigarrow C0}^{1,1}$ is $(C0 \rightarrow C1) \cdot (((C1 \rightarrow C2) \cdot (C2 \rightarrow C5)) + ((C1 \rightarrow C4) \cdot (C4 \rightarrow C5))) \cdot (C5 \rightarrow C0)$, or, after abbreviating serial combinations, $FP_{C0 \rightsquigarrow C0}^{1,1} = (C0 \rightarrow C1) \cdot (C1 \rightarrow C2 \rightarrow C5 + C1 \rightarrow C4 \rightarrow C5) \cdot (C5 \rightarrow C0)$. An another example, the $(4, 4)$ -unrolled full pathset of loop path from C0, say $FP_{C0 \rightsquigarrow C0}^{4,4}$, contains 16 paths and is represented by a serial concatenation of four $FP_{C0 \rightsquigarrow C0}^{0,0}$: $FP_{C0 \rightsquigarrow C0}^{4,4} = ((C0 \rightarrow C1) \cdot (C1 \rightarrow C2 \rightarrow C5 + C1 \rightarrow C4 \rightarrow C5) \cdot (C5 \rightarrow C0))^4$, where the repetition is indicated by the exponent.

COUNTER correlates a (μ, δ) -unrolled full pathset in a single step (such that a product graph edge contains a (μ, δ) -unrolled full pathset). The primary advantage of using (μ, δ) -unrolled full pathset is that its SP-graph representation enables an efficient SMT

encoding of a Hoare triple over such a pathset. Using such SP-graph representation, it is possible to have linear-sized (in number of nodes in the SP-graph) SMT proof obligations even if the SP-graph represents an exponential number of paths. SMT solvers are able to discharge these linear-sized SMT obligations faster than time taken to discharge the exponential number of proof obligations for individual paths. For example, the *weakest precondition*[14] of predicate $x = 0$ over $\text{FP}_{\text{C0} \rightsquigarrow \text{C0}}^{4,4}$ is a linear sized expression

$$\begin{aligned} & x + (\text{c}[i] \quad ? \text{a}[i] \quad : \text{b}[i]) \\ & + (\text{c}[i + 1] ? \text{a}[i + 1] : \text{b}[i + 1]) \\ & + (\text{c}[i + 2] ? \text{a}[i + 2] : \text{b}[i + 2]) \\ & + (\text{c}[i + 3] ? \text{a}[i + 3] : \text{b}[i + 3]) = 0 \end{aligned}$$

where the C-like ternary operator ‘?:’ is used as a shorthand for the ‘if-then-else’ operator of SMT. SMT solvers are able to discharge one such proof obligation faster than 16 obligations (for each path in $\text{FP}_{\text{C0} \rightsquigarrow \text{C0}}^{4,4}$) performed over individual paths.

Recall that the set of paths in a pathset $\langle \xi \rangle_{\text{C}}$ must be mutually exclusive by definition. Further, a candidate pathset $\langle \xi \rangle_{\text{C}}$ enumerated for a path $\xi_{\bar{\text{A}}} = (q_{\bar{\text{A}}} \rightarrow q_{\bar{\text{A}}}^{\text{UW}})$ ($q_{\bar{\text{A}}}^{\text{UW}} \notin \{\mathcal{U}_{\bar{\text{A}}}, \mathcal{W}_{\bar{\text{A}}}\}$) can be partitioned into two pathsets $\langle \xi \rangle_{\text{C}}^t$ and $\langle \xi \rangle_{\text{C}}^u$ where the pathset $\langle \xi \rangle_{\text{C}}^t$ contains paths with an error-free sink node q_{C}^t and the pathset $\langle \xi \rangle_{\text{C}}^u$ contains paths with sink node \mathcal{U}_{C} . If the $\langle \xi \rangle_{\text{C}}^t$ pathset is a (μ, δ) -unrolled full pathset, i.e., the `enumPathsetsTillUnroll()` function uses the `getFullPathsetAtAllDelta()` from [17] (as we do in our implementation) then the SP-graph representation can be used for representing $\langle \xi \rangle_{\text{C}}^t$.

4.1.3 Identifying A annotation

For each feasible simple path $\xi_{\bar{\text{A}}}$ and each (potentially non-simple) path $\xi_{\text{C}} \in \langle \xi \rangle_{\text{C}}$, the `asmAnnotOpts()` procedure enumerates the options for annotating $\xi_{\bar{\text{A}}}$ with `allocs,v`, `deallocs,v` instructions and operands for assembly `call` instructions.

An annotation option includes the positions and the operands of the (de)allocation instructions:

- For an `allocs` instruction ‘ $p_{\bar{\text{A}}}^j : \text{alloc}_s \ e_v, e_w, a, z$ ’, an annotation option would have $p_{\bar{\text{A}}}^j$ as the position, e_v as the start address, e_w as the allocation size, a as the

alignment, and $z \in Z$ as the allocation site.

- For an `allocv` instruction ‘ $p_{\check{A}}^j : \text{alloc}_v e_w, a, zl$ ’, an annotation option would have $p_{\check{A}}^j$ as the position, e_w as the allocation size, a as the alignment, and $zl \in Z_l$ as the allocation site.

Recall that we restrict `allocv` annotation to allocation sites due to a declaration of a local variable $zl \in Z_l$, excluding allocation sites due to `alloca()` (section 2.6). *asmAnnotOpts()* thus limits enumeration of annotation options for `allocv` to allocation sites Z_l .

For an assembly call instruction ‘ $p_{\check{A}}^j : \text{call } \rho$ ’, an annotation option would have the return type γ , the types and values of arguments $(\vec{\tau} \vec{x})$, and the set of callee-observable regions β^* . The annotated call instruction would be ‘ $p_{\check{A}}^j : \text{call } \gamma \rho(\vec{\tau} \vec{x}) \beta^*$ ’.

asmAnnotOpts() algorithm

Algorithm 4 shows the pseudo-code for the *asmAnnotOpts()* procedure. The procedure takes two correlated paths $\xi_{\check{A}}$ and $\xi_{\check{C}}$ and procedures \check{A} and \check{C} as input and returns the *options* for (potentially) annotated \check{A} and modified path $\xi_{\check{A}}^*$ as output.

(SingleIO) requirement enforces a lockstep correlation between (de)allocation and procedure-call instructions in $\xi_{\check{A}}$ and $\xi_{\check{C}}$. *asmAnnotOpts()* thus annotates a (de)`allocs,v`, `call` instruction in $\xi_{\check{A}}$ only if a corresponding (de)`alloc`, `call` instruction is present in $\xi_{\check{C}}$. The call to `collectAllocDeallocCallInsns()` at line 2 in algorithm 4 returns, in $\vec{\clubsuit}_{\check{C}}$, the list of (de)`alloc` and `call` instructions in input $\xi_{\check{C}}$. *asmAnnotOpts()* attempts to generate annotation options, represented using $\vec{\pi}$, for each returned instruction $\clubsuit_{\check{C}} \in \vec{\clubsuit}_{\check{C}}$ (lines 4 to 24 in algorithm 4). The Cartesian product of the generated annotation options $\vec{\pi}$ (for each instruction $\clubsuit_{\check{C}} \in \vec{\clubsuit}_{\check{C}}$) forms the annotation candidates $\vec{\ominus}$ for $\xi_{\check{A}}$ and each annotation candidate $\ominus \in \vec{\ominus}$, which is a sequence of annotation options (one for each $\clubsuit_{\check{C}} \in \vec{\clubsuit}_{\check{C}}$), is separately applied to input \check{A} to generate an output annotated \check{A} (and the updated $\xi_{\check{A}}, \xi_{\check{A}}^*$) option.

In the whitebox setting (lines 6 to 9 in algorithm 4), annotation hints are available that can be used for precisely identifying the required annotation π . These annotation hints are derived from multiple sources, including hints from instrumented compiler (see section 6.1.4 for an example), debug headers of the executable, or manually provided by the user. Further, they annotation hints are of varied quality, for example, hints harvested from debug headers of an highly optimized executable are not very reliable

Algorithm 4: Pseudo-code of the *asmAnnotOpts()* procedure.

```

1 Function asmAnnotOpts( $\xi_{\bar{A}}, \xi_C, \bar{A}, C$ )
2    $\vec{\clubsuit}_C \leftarrow \text{collectAllocDeallocCallInsns}(\xi_C, C)$ ;
3    $\vec{\ominus} \leftarrow \{\emptyset\}$ ;
4   foreach  $\clubsuit_C$  in  $\vec{\clubsuit}_C$  do
5      $\vec{\pi} \leftarrow \emptyset$ ;
6     if hints are available for  $\clubsuit_C$  then // whitebox setting
7       foreach hint  $h$  do
8          $\vec{\pi} \leftarrow \vec{\pi} \cup \text{genAnnotOptUsingHint}(\xi_{\bar{A}}, \clubsuit_C, \xi_C, h, \bar{A}, C)$ ;
9       end
10      // blackbox enumeration
11      if  $\clubsuit_C$  is alloc then
12        if alloc insn not already present for  $\clubsuit_C$  in  $\bar{A}$  then
13           $\vec{\pi} \leftarrow \text{genAnnotOptsForAlloc}(\xi_{\bar{A}}, \clubsuit_C, \xi_C, \vec{\ominus}, \bar{A}, C)$ ;
14        else if  $\clubsuit_C$  is dealloc then
15           $\vec{\pi} \leftarrow \text{genAnnotOptsForDealloc}(\xi_{\bar{A}}, \clubsuit_C, \xi_C, \vec{\ominus}, \bar{A}, C)$ ;
16        else // procedure call
17           $\vec{\pi} \leftarrow \text{genAnnotOptsForFcallUsingCallConvAndCPath}(\xi_{\bar{A}}, \clubsuit_C, \xi_C, \vec{\ominus}, \bar{A}, C)$ ;
18         $\vec{\ominus}' = \emptyset$ ;
19        // cross-product all possibilities with existing options
20        foreach  $\ominus$  in  $\vec{\ominus}$  do
21          foreach  $\pi$  in  $\vec{\pi}$  do
22             $\vec{\ominus}' \leftarrow \vec{\ominus}' \cup \{\ominus \cdot \pi\}$ ;
23          end
24        end
25         $\vec{\ominus} \leftarrow \vec{\ominus}'$ ;
26      end
27       $ret \leftarrow \emptyset$ ;
28      // generate an annotated  $\bar{A}$  for each option
29      foreach  $\ominus$  in  $\vec{\ominus}$  do
30         $\bar{A}', \xi_{\bar{A}}^* \leftarrow \text{applyAnnots}(\bar{A}, \xi_{\bar{A}}, \ominus)$ ;
31         $ret \leftarrow ret \cup (\bar{A}', \xi_{\bar{A}}^*)$ ;
32      end
33      return  $ret$ 
34 end

```

[26]. We use all available hints for generating (multiple) annotation options. A best-first search implementation (see section 4.1.9) may choose to rank options from certain sources (e.g., hints from instrumented compiler) ahead of others.

If no annotation hints are available, the correlated (de)alloc/call instruction \clubsuit_C in ξ_C can be utilized for deducing partial annotation. The allocation size, alignment, and allocation site operands of an (de)alloc_{s,v} instruction are uniquely identified from the correlated (de)alloc instruction \clubsuit_C . Similarly, the arguments' count and types, the return type, and the callee-observable regions for a call instruction in $\xi_{\bar{A}}$ are identified using a correlated call instruction \clubsuit_C ; the deduction of arguments'

count further enables identification of the addresses for the arguments of `call` using calling conventions (`genAnnotOptsForFcallUsingCallConvAndCPath()` at line 16 in algorithm 4).

After the deduction of other parameters, only the position and start address⁴ of an `(de)allocs,v` instruction need to be determined (or guessed) for a complete annotation. For a `(de)allocs` due to a procedure call argument, the position and start address are determined based on argument order and calling conventions. For rest of the `(de)allocs,v` instructions, we reduce the search space for the position and start address (at the cost of reduced generality) using the following three restrictions:

1. An `allocs,v` (`deallocs,v`) annotation is added only after (before) an instruction that updates the stackpointer `esp`.
2. For an `allocs` instruction, the stackpointer value `esp` after the update is used as the annotation for the start address expression.
3. For a single allocation site in `C`, at most one `allocs,v` instruction (but potentially multiple `deallocs,v` instructions) is added to \ddot{A} (line 11 in algorithm 4).

Thus, in a blackbox setting, due to the third restriction, a refinement proof may fail if the compiler specializes a path containing a local variable allocation. Due to the first and second restriction, a refinement proof may fail for certain (arguably rare) types of (de)allocation order preserving stack reallocation and stack merging performed by the compiler. An example is discussed in section 6.2.4. Note that these limitations hold only for the blackbox setting.

An annotation option π , obtained either from whitebox hints or blackbox enumeration, is accumulated in each sequence of annotation options $\Theta \in \vec{\Theta}$ enumerated so far (lines 18 to 22 in algorithm 4). Once all instructions in \clubsuit_C have been considered, $\vec{\Theta}$ contains the Cartesian product of annotation options enumerated for each $\clubsuit_C \in \vec{\clubsuit}_C$ and each $\Theta \in \vec{\Theta}$ forms an annotation candidate for $\xi_{\ddot{A}}$.

Each enumerated annotation candidate $\Theta \in \vec{\Theta}$ is applied separately to A to incrementally construct \ddot{A} (`applyAnnots()` at line 27 in algorithm 4). The application of an annotation may potentially update the path $\xi_{\ddot{A}}$ to $\xi_{\ddot{A}}^*$ due to addition of edges for `(de)allocs,v` and annotated `call`. `asmAnnotOpts()` returns the options for the updated \ddot{A} and $\xi_{\ddot{A}}^*$.

⁴start address is only required for the `allocs` instruction

Algorithm 5: Pseudo-code of the *trimToMatchPathToErrorNode()* procedure.

```

1 Function trimToMatchPathToErrorNode( $\vec{\xi}_{\ddot{A}}, \vec{\xi}_{\ddot{C}}$ )
2    $l \leftarrow \min(|\vec{\xi}_{\ddot{A}}|, |\vec{\xi}_{\ddot{C}}|)$ ;
3    $\vec{\xi}_{\ddot{A}}^*, \vec{\xi}_{\ddot{C}}^* \leftarrow \text{take}(l, \vec{\xi}_{\ddot{A}}), \text{take}(l, \vec{\xi}_{\ddot{C}})$ ; // take first  $l$  elements from  $\vec{\xi}_{\ddot{A}}$  and  $\vec{\xi}_{\ddot{C}}$ 
4   if  $\text{sink}(\text{last}(\vec{\xi}_{\ddot{A}}^*)) \in \{\mathcal{U}_{\ddot{A}}, \mathcal{W}_{\ddot{A}}\}$  or  $\text{sink}(\text{last}(\vec{\xi}_{\ddot{C}}^*)) = \mathcal{U}_{\ddot{C}}$  then // return trimmed
      sequences if either ends with error node
5     | return  $(\vec{\xi}_{\ddot{A}}^*, \vec{\xi}_{\ddot{C}}^*)$ 
6   else // otherwise, return the original sequences
7     | return  $\vec{\xi}_{\ddot{A}}, \vec{\xi}_{\ddot{C}}$ 
8 end

```

4.1.4 Validating structure of identified paths

asmAnnotOpts() produces a set of options for annotation and the **chooseFrom** operator chooses one such that (if possible) the annotated instructions generate identical traces in \ddot{A} and \ddot{C} .

It must be emphasized here that due to these annotated instructions, extra paths to error nodes $\mathcal{U}_{\ddot{A}}$ and $\mathcal{W}_{\ddot{A}}$ are added to \ddot{A} — recall the multi-line graph instructions translations of (de)alloc_{s,v} instructions, with checks for overlap and alignment, presented in (ALLOCS'), (DEALLOCS'), (ALLOCV), and (DEALLOCV). These paths are *not* a part of the annotated $\xi_{\ddot{A}}^*$ that contains the error-free sub-paths of the annotated (de)alloc_{s,v} instructions. Instead, the correlation of these extra paths to error nodes happens in second phase of the algorithm (when *phase* = CORRELATE_NEW_ERROR_PATHS; recall that the algorithm operates in two phases).

After annotation, the annotated path $\xi_{\ddot{A}}^*$ may become a non-simple path⁵ due to the extra I/O instructions introduced by the annotation. The (potentially non-simple) path $\xi_{\ddot{A}}^*$ is therefore broken into a sequence of constituent paths $\vec{\xi}'_{\ddot{A}}$ using *breakIntoSingleIOPaths()* (line 26 of algorithm 1) so that each I/O path appears by itself (and not as a sub-path of a longer constituent path) — this caters to the (SingleIO) requirement (section 3.3.1). *breakIntoSingleIOPaths()* is similarly used on $\xi_{\ddot{C}}$ to obtain a sequence of simple paths $\vec{\xi}'_{\ddot{C}}$.

The (SingleIO) requirement requires that each I/O path $\xi'_{\ddot{A}} \in \vec{\xi}'_{\ddot{A}}$ is correlated separately with an I/O path $\xi'_{\ddot{C}} \in \vec{\xi}'_{\ddot{C}}$ of similar kind. However, the sequences of simple paths $\vec{\xi}'_{\ddot{A}}$ and $\vec{\xi}'_{\ddot{C}}$ obtained after *breakIntoSingleIOPaths()* may not have identical lengths. For

⁵Recall that a simple path cannot have a cut-point as an intermediate node.

Algorithm 6: Pseudo-code of the *haveSimilarStructure()* procedure.

```

1 Function haveSimilarStructure( $\vec{\xi}_{\ddot{A}}, \vec{\xi}_{\mathcal{C}}$ )
2   if  $|\vec{\xi}_{\ddot{A}}| \neq |\vec{\xi}_{\mathcal{C}}|$  then
3     return false
4    $(n_{\ddot{A}}, n_{\ddot{A}}^t) \leftarrow \text{src}(\text{first}(\vec{\xi}_{\ddot{A}})), \text{sink}(\text{last}(\vec{\xi}_{\ddot{A}}));$ 
5    $(n_{\mathcal{C}}, n_{\mathcal{C}}^t) \leftarrow \text{src}(\text{first}(\vec{\xi}_{\mathcal{C}})), \text{sink}(\text{last}(\vec{\xi}_{\mathcal{C}}));$ 
6   if  $n_{\ddot{A}}^t = \mathcal{U}_{\ddot{A}} \wedge n_{\mathcal{C}}^t \neq \mathcal{U}_{\mathcal{C}}$  then // (Safety)
7     return false
8   if  $n_{\ddot{A}}^t \notin \{\mathcal{U}_{\ddot{A}}, \mathcal{W}_{\ddot{A}}\} \wedge n_{\mathcal{C}}^t \notin \{\mathcal{U}_{\mathcal{C}}, \mathcal{W}_{\mathcal{C}}\} \wedge \text{term}(n_{\ddot{A}}^t) \neq \text{term}(n_{\mathcal{C}}^t)$  then // (Termination)
9     return false
10  // (SingleIO) --- each I/O path-pair is of same kind
11  foreach  $(\xi_{\ddot{A}}, \xi_{\mathcal{C}})$  in  $\text{zip}(\vec{\xi}_{\ddot{A}}, \vec{\xi}_{\mathcal{C}})$  do
12    if  $\text{isIOPath}(\xi_{\ddot{A}}) \neq \text{isIOPath}(\xi_{\mathcal{C}})$  then
13      return false
14    if  $\text{isIOPath}(\xi_{\ddot{A}}) \wedge \neg \text{IOPathsOfSameKind}(\xi_{\ddot{A}}, \xi_{\mathcal{C}})$  then
15      return false
16  end
17  return true
18 end

```

example, if $\xi_{\ddot{A}}$ is a path to an error-free node $n_{\ddot{A}}^t$ and $\xi_{\mathcal{C}}$ is a path to the error node $\mathcal{U}_{\mathcal{C}}$, then the sequence $\vec{\xi}'_{\ddot{A}}$ may be larger than $\vec{\xi}'_{\mathcal{C}}$. In such a scenario, because $\mathcal{U}_{\mathcal{C}}$ is a terminating node, the path sequence $\vec{\xi}'_{\ddot{A}}$ can be *trimmed* to make the sequence lengths identical. This trimming is permissible because under the refinement definition the generated traces are required to be identical only till \mathcal{C} halts with error \mathcal{U} or \ddot{A} halts with error \mathcal{W} (section 2.4). The procedure *trimToMatchPathToErrorNode()* (line 28 of algorithm 1) attempts to make the lengths of $\vec{\xi}'_{\ddot{A}}$ and $\vec{\xi}'_{\mathcal{C}}$ identical by trimming $\vec{\xi}'_{\ddot{A}}$ and $\vec{\xi}'_{\mathcal{C}}$ in this fashion. Algorithm 5 shows the pseudo-code of *trimToMatchPathToErrorNode()*.

Next, the *haveSimilarStructure()* procedure validates the structure of $\vec{\xi}_{\ddot{A}}$ and $\vec{\xi}_{\mathcal{C}}$ obtained after *trimToMatchPathToErrorNode()* (line 29 in algorithm 1). *haveSimilarStructure*($\vec{\xi}_{\ddot{A}}, \vec{\xi}_{\mathcal{C}}$) returns true iff the sequence of paths $\vec{\xi}_{\ddot{A}}$ and $\vec{\xi}_{\mathcal{C}}$ have identical lengths and are *similarly structured*, where structural similarity is defined with respect to the structural requirements of X . Algorithm 6 shows the pseudo-code of *haveSimilarStructure()*. Let $\text{pos}(\xi, \vec{\xi})$ represent the position of path ξ in a sequence of paths $\vec{\xi}$. Let $\xi_{\ddot{A}}^j$ and $\xi_{\mathcal{C}}^j$ denote paths such that $\xi_{\ddot{A}}^j \in \vec{\xi}_{\ddot{A}}$, $\xi_{\mathcal{C}}^j \in \vec{\xi}_{\mathcal{C}}$ and $\text{pos}(\xi_{\mathcal{C}}^j, \vec{\xi}_{\mathcal{C}}) = \text{pos}(\xi_{\ddot{A}}^j, \vec{\xi}_{\ddot{A}}) = j$; we will refer to the pair $(\xi_{\ddot{A}}^j, \xi_{\mathcal{C}}^j)$ as a “coupled path-pair”. *haveSimilarStructure()* ensures that, i.e., it returns true if, if the partially-constructed X satisfies the structural requirements of section 3.3.1, then it will continue to satisfy the (SingleIO), (Safety), and (Termination) requirements after adding edges corresponding to coupled path-pairs $e_{\mathsf{X}}^j = (\xi_{\ddot{A}}^j, \xi_{\mathcal{C}}^j)$ (for all j) to X . Let $\xi_{\ddot{A}}$ and $\xi_{\mathcal{C}}$ be the last paths in the path sequences

$\vec{\xi}_{\bar{A}}$ and $\vec{\xi}_{\bar{C}}$ respectively such that $\text{pos}(\xi_{\bar{A}}, \vec{\xi}_{\bar{A}}) = \text{pos}(\xi_{\bar{C}}, \vec{\xi}_{\bar{C}}) = |\vec{\xi}_{\bar{A}}|$. (Safety) requires that if the sink node $n_{\bar{A}}^t$ of $\xi_{\bar{A}}$ is $\mathcal{U}_{\bar{A}}$, then the sink node $n_{\bar{C}}^t$ of $\xi_{\bar{C}}$ must be $\mathcal{U}_{\bar{C}}$ and (Termination) requires that if both $n_{\bar{A}}^t$ and $n_{\bar{C}}^t$ are error-free, then they must agree on terminating status, i.e., $\text{term}(n_{\bar{A}}^t) = \text{term}(n_{\bar{C}}^t)$. The (SingleIO) check obligations include ensuring that a coupled I/O path-pair $(\xi_{\bar{A}}^j, \xi_{\bar{C}}^j)$ is an I/O path-pair of *same kind*. This is expressed through `IOPathsOfSameKind()` in algorithm 6 wherein `IOPathsOfSameKind($\xi_{\bar{A}}^j, \xi_{\bar{C}}^j$)` returns true iff $\xi_{\bar{A}}^j$ and $\xi_{\bar{C}}^j$ are either both reads or both writes for the same type of value (implemented as syntactic checks on the read/written value⁶).

4.1.5 Incremental construction of (\bar{A}, X)

At this point (line 31 in algorithm 1), we have two sequences of paths $\vec{\xi}_{\bar{A}}^*$ and $\vec{\xi}_{\bar{C}}^*$ that have identical number of elements. In its next step, the DYNAMO algorithm constructs an X edge $e_X = (\xi'_{\bar{A}} = (n_{\bar{A}} \rightarrow n_{\bar{A}}^t), \xi'_{\bar{C}} = (n_{\bar{A}} \rightarrow n_{\bar{A}}^t))$ from a coupled path-pair $(\xi'_{\bar{A}}, \xi'_{\bar{C}})$ (through `zip()` in line 31 of algorithm 1). To ensure (Similar-Speed) requirement (section 3.3.1), the algorithm checks that the addition of e_X does not create an empty cycle of \bar{C} paths in \mathcal{E}_X (`addingEdgeWillCreateEmptyCCycle()` in algorithm 1). We only need to perform the check on \bar{C} paths because our enumeration guarantees that $\xi'_{\bar{A}}$ is never empty — recall that `correlatedPathsInCOpts()` allows a path $\xi_{\bar{C}}$ in a pathset $\langle \xi \rangle_{\bar{C}}$ to be empty (so that $\xi'_{\bar{C}}$ can be empty) but $\xi_{\bar{A}}$, as returned by `getAllSimplePathsBetweenCutPoints()`, can never be empty. The constructed edge $e_X = (\xi'_{\bar{A}}; \xi'_{\bar{C}})$ is added to \mathcal{E}_X and its destination node $(n_{\bar{A}}^t, n_{\bar{C}}^t)$ is added to \mathcal{N}_X , if not already present.

If $\xi'_{\bar{C}}$ contains an edge labeled with a choose instruction ($'\vec{v} := \theta(\vec{\tau})'$), then we update the deterministic choice map \mathcal{D}_X to include determinized mappings for each choose instruction in $\xi'_{\bar{C}}$ through `addDetMappings()`; algorithm 7 shows the pseudo-code of `addDetMappings()`. For example, if $\xi'_{\bar{C}}$ represents a path between `wr(allocBegin(...))` and `wr(allocEnd(...))` for an `alloc` instruction in \bar{C} ((`ALLOC`) in fig. 2.5), and $\xi'_{\bar{A}}$ is a corresponding path due to an `allocs,v` instruction, and edges $e_{\bar{C}}^{\theta_a}$ and $e_{\bar{C}}^{\theta_m}$ in $\xi'_{\bar{C}}$ are labeled with instructions $'a_b := \theta(i_{32})'$ and $'\theta(i_{32} \rightarrow i_8)'$ respectively due to (`ALLOC`), we add mappings $\mathcal{D}_X(e_X, e_{\bar{C}}^{\theta_a}, 1) = v$ and $\mathcal{D}_X(e_X, e_{\bar{C}}^{\theta_m}, 1) = M_{\bar{A}}$, where v is the address defined in $\xi'_{\bar{A}}$ due to either `allocs` ((`ALLOCS`)) or `allocv` ((`ALLOCV`)). In algorithm 7, the $e_{\bar{C}}^{\theta_a}$ edge is identified using `identifyAllocAddrEdge()` and the

⁶Recall that we use different value constructors for different instructions; see item 4 in section 2.2.7

Algorithm 7: Pseudo-code of the *addDetMappings()* and *inferInvariantsAndCounterexamples()* procedures.

```

1 Function addDetMappings( $e_X, \mathcal{D}_X$ )
2    $\mathcal{D}_X' \leftarrow \mathcal{D}_X$ ;
3    $(\xi_{\bar{A}}, \xi_C) \leftarrow e_X$ ;
4   if  $e_C^{\theta_a} \leftarrow \text{identifyAllocAddr}\theta\text{Edge}(\xi_C)$  then
5      $e_C^{\theta_m} \leftarrow \text{identifyAllocMem}\theta\text{Edge}(\xi_C)$ ;
6      $\mathcal{D}_X'(e_X, e_C^{\theta_a}, 1) \leftarrow \text{identifyAllocAddr}(\xi_{\bar{A}})$ ; //  $v$  in (ALLOCS), (ALLOCV)
7      $\mathcal{D}_X'(e_X, e_C^{\theta_m}, 1) \leftarrow M_{\bar{A}}$ ;
8   else if  $e_C' \leftarrow \text{identifyEntryMem}\theta\text{Edge}(\xi_C)$  then
9      $\mathcal{D}_X'(e_X, e_C', 1) \leftarrow M_{\bar{A}}$ ;
10  return  $\mathcal{D}_X'$ 
11 end
12 Function inferInvariantsAndCounterexamples( $n_X, \mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X, \Phi_X, \bar{A}, C$ )
13  if  $n_X$  is an error node then
14    return
15   $\star_{\bar{A}}(n_X) \leftarrow \text{computeInterestingExprsAtNodeForA}(\bar{A}, \mathcal{N}_X, \mathcal{E}_X, \Phi_X)$ ;
16   $\star_C(n_X) \leftarrow \text{computeInterestingExprsAtNodeForC}(C, \mathcal{N}_X, \mathcal{E}_X, \Phi_X)$ ;
17   $\Phi_X \leftarrow \text{inferInvariantsForExprs}(\star_{\bar{A}}, \star_C, \mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X, \Phi_X, \bar{A}, C)$ ;
18 end

```

$e_C^{\theta_m}$ edge is identified using `identifyAllocMem θ Edge()`. Notice that our algorithm only populates $\mathcal{D}_X(e_X, e_C^\theta, n)$ for $n = 1$, even though section 3.2 defines \mathcal{D}_X more generally. Algorithm 7 also populates \mathcal{D}_X for an edge e_C' that corresponds to the ‘ $\theta(i_{32} \rightarrow i_8)$ ’ edge due to procedure-entry ((ENTRY_C) in fig. 2.5 and (ENTRY_A) in fig. 2.7), identified using `identifyEntryMem θ Edge()` in algorithm 7.

Recall that DYNAMO maintains the invariant that all non-coverage requirements hold over the partially-constructed X . So far, the algorithm has ensured that each structural requirement is satisfied by the newly added edge $e_X = (n_X \rightarrow (n_{\bar{A}}^t, n_C^t))$. If the destination node $(n_{\bar{A}}^t, n_C^t)$ is an error-free node, the invariant network Φ_X may no longer be inductive due to the addition of the new edge, thereby violating the (Inductive) requirement (see section 3.3.1). The *inferInvariantsAndCounterexamples()* procedure, shown in algorithm 7, updates Φ_X to ensure its inductivity. To keep invariant inference tractable, the state elements and expressions participating in invariant inference at a node $n_X = (n_{\bar{A}}, n_C)$ are restricted to *interesting expressions* at n_X — which include at least the live registers, ghost variables, and stack slots at $n_{\bar{A}}$ in \bar{A} (shown as $\star_{\bar{A}}(n_X)$ in algorithm 7) and a subset of all defined variables (including ghost variables) for C (shown as $\star_C(n_X)$ in algorithm 7). These choices are similar to the ones considered in [42]. We defer a detailed discussion of the invariant inference, including a description of the candidate invariant grammar, to section 4.2.

Algorithm 8: Pseudo-code of the *checkSemanticReqsExceptCoverage()* procedure.

```

1 Function checkSemanticReqsExceptCoverage( $\mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X, \Phi_X, \check{A}, C$ )
2   if  $\neg$ invariantsAreInductive( $\mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X, \Phi_X, \check{A}, C$ ) then
3     return false
4   if  $\neg$ checkEquivalence( $\mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X, \Phi_X, \check{A}, C$ ) then
5     return false
6   if  $\neg$ checkMAC( $\mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X, \Phi_X, \check{A}, C$ ) then
7     return false
8   if  $\neg$ checkMemEq( $\mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X, \Phi_X, \check{A}, C$ ) then
9     return false
10  return true
11 end

```

4.1.6 Checking requirements on partial X

Once the invariant network has been updated, the non-coverage semantic requirements are checked through *checkSemanticReqsExceptCoverage()* procedure. A partially constructed X that fails to satisfy the requirements check is discarded⁷. The pseudo-code of *checkSemanticReqsExceptCoverage()*, shown in algorithm 8, performs checks for the following semantic requirements on the partial X: (Inductive), (Equivalence), (Memory Access Correspondence) or (MAC), and (MemEq). The (Inductive) check is realized through *invariantsAreInductive()*, which fails if the invariant network Φ_X is not inductive. Similarly, the (Equivalence), (MemEq), and (Memory Access Correspondence) or (MAC) checks are realized through *checkEquivalence()*, *checkMemEq()*, and *checkMAC()* respectively. The former two, *checkEquivalence()* and *checkMemEq()*, simply involve checking if the set of invariants ϕ_{n_X} , inferred at each error-free node $n_X \in \mathcal{N}_X$, include the required invariants, $\Omega_C = \Omega_{\check{A}}$ and $M_{\check{A}} =_{\Sigma_{\check{A}}^B \setminus (\Sigma_{\check{A}}^{Z_I} |^v)} M_C$ respectively⁸.

(Memory Access Correspondence) or (MAC) check

The pseudo-code for the *checkMAC()* sub-procedure is shown in algorithm 8. Recall that the (MAC) requirement check (section 3.3.1) entails ensuring that for each edge $e_X = (n_X \xrightarrow{\xi_{\check{A}}; \xi_C} n'_X) \in \mathcal{E}_X$, such that $n'_X \neq (-, \mathcal{U}_C)$ and for each memory access to interval $[\alpha]_w$ in $\xi_{\check{A}}$, either:

⁷Recall that the **chooseFrom** operator chooses option such that this check does not fail if the required X can be constructed using DYNAMO.

⁸As we will see later in section 4.2, these are instantiations of the invariant shapes WEq and MemEq respectively.

Algorithm 9: Pseudo-code of the *checkMAC()* procedure.

```

1 Function checkMAC( $\mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X, \Phi_X, \ddot{A}, C$ )
2   foreach  $\xi_{\ddot{A}}, \xi_C$  such that  $e_X = (n_X \xrightarrow{\xi_{\ddot{A}}; \xi_C} n'_X) \in \mathcal{E}_X$  and  $n'_X \neq (-, \mathcal{U}_C)$  do
3      $memAcc_C \leftarrow collectMemAccessesInPath(\xi_C)$ ;
4      $isSafe_C(rd) \leftarrow \bigwedge_{(\xi_C^j, \alpha_j, w_j, rd) \in memAcc_C} WP_{[\xi_C^j]_{\mathcal{D}_X}^{ex}}(accessIsSafe(\alpha_j, w_j, rd, C))$ ;
5      $isSafe_C(wr) \leftarrow \bigwedge_{(\xi_C^j, \alpha_j, w_j, wr) \in memAcc_C} WP_{[\xi_C^j]_{\mathcal{D}_X}^{ex}}(accessIsSafe(\alpha_j, w_j, wr, C))$ ;
6     foreach  $(\xi_{\ddot{A}}^j, \alpha_j, w_j, k_j) \in collectMemAccessesInPath(\xi_{\ddot{A}})$  do
7        $isSafe_{\ddot{A}} \leftarrow accessIsSafe(\alpha_j, w_j, k_j, \ddot{A})$ ;
8       if Hoare triple  $\{\phi_{n_X} \wedge pathcond([\xi_C]_{\mathcal{D}_X}^{ex})\}(\xi_{\ddot{A}}^j; \epsilon)\{isSafe_C(k_j) \implies isSafe_{\ddot{A}}\}$  does
9         not hold then
10        | return false
11      end
12    end
13  return true
14 Function accessIsSafe( $\alpha, w, k, P$ )
15   // page size parametric access-is-safe condition generator
16    $pgMask \leftarrow \sim(PAGE\_SIZE - 1)$ ;
17    $un \leftarrow$  unique variable different from all state variables in  $C$  and  $\ddot{A}$ ;
18    $unL \leftarrow un \ \& \ pgMask$ ;     $unH \leftarrow unL + PAGE\_SIZE - 1$ ;
19    $l, h \leftarrow \alpha, \alpha + w - 1$ ;
20    $inUnalloc \leftarrow ov([l, h], [unL, unH])$ ;
21   if  $P = \ddot{A}$  then
22     | if  $k = rd$  then
23       |  $inGFS \leftarrow [l, h] \subseteq (\Sigma_{\ddot{A}}^{GUF} \cup [esp, \boxed{stk_e}])$ ;
24     | else //  $k = wr$ 
25       |  $inGFS \leftarrow [l, h] \subseteq (\Sigma_{\ddot{A}}^{G_wUF_w} \cup [esp, \boxed{stk_e}])$ ;
26     | return  $inUnalloc \implies inGFS$ 
27   else //  $P = C$ 
28     | return  $\neg inUnalloc$ 

```

- A corresponding memory access (a corresponding read or a corresponding write) to identical interval is present in ξ_C . In other words, the access is *matched* in ξ_C .
- Or, for a read access, $[\alpha]_w \subseteq (\Sigma_{\ddot{A}}^{GUF} \cup [esp, \boxed{stk_e}])$, and for a write access, $[\alpha]_w \subseteq (\Sigma_{\ddot{A}}^{G_wUF_w} \cup [esp, \boxed{stk_e}])$.

In the *checkMAC()* procedure, the first clause is generalized to *page-granularity* where an access in $\xi_{\ddot{A}}$ is deemed to be matched if it remains within the same memory page as an access in ξ_C , i.e., an access interval $[\alpha]_w$ ($[\alpha]_w \subseteq \boxed{\Sigma_{\ddot{A}}^{rd}}$ or $[\alpha]_w \subseteq \boxed{\Sigma_{\ddot{A}}^{wr}}$) is matched

iff the following holds:

$$\forall_p : \left(\begin{array}{l} \text{aligned}_{\text{PAGE_SIZE}}(p) \\ \wedge \text{ov}([\alpha]_w, [p]_{\text{PAGE_SIZE}}) \end{array} \right) \implies \exists_{\alpha', w'} : \left(\begin{array}{l} [\alpha']_{w'} \subseteq \Sigma_{\mathbb{C}}^{\text{rd/wr}} \\ \wedge \text{ov}([\alpha']_{w'}, [p]_{\text{PAGE_SIZE}}) \end{array} \right)$$

where $\Sigma_{\mathbb{C}}^{\text{rd/wr}} = \Sigma_{\mathbb{C}}^{\text{rd}}$ if $[\alpha]_w \subseteq \Sigma_{\mathbb{A}}^{\text{rd}}$ and $\Sigma_{\mathbb{C}}^{\text{rd/wr}} = \Sigma_{\mathbb{C}}^{\text{wr}}$ if $[\alpha]_w \subseteq \Sigma_{\mathbb{A}}^{\text{wr}}$, and PAGE_SIZE is the size of the page (e.g., $\text{PAGE_SIZE} = 4096$). This page granular version is equivalent to the original requirement when $\text{PAGE_SIZE} = 1$.

The `checkMAC()` implementation in algorithm 8 formulates the (MAC) check as a search for an arbitrary *unallocated* page of size PAGE_SIZE starting at an (page-aligned) address *un* such that the access interval $[\alpha]_w$ (in \mathbb{A}) overlaps with it but none of the memory accesses in $\xi_{\mathbb{C}}$ overlap with it.

Let $e_X = ((n_{\mathbb{A}}, n_{\mathbb{C}}) \xrightarrow{\xi_{\mathbb{A}}; \xi_{\mathbb{C}}} n_X^t)$ such that $n_X^t \neq (-, \mathcal{U}_{\mathbb{C}})$. A memory access on $\xi_{\mathbb{A}}$ to address $\alpha_{\mathbb{A}}$, of size $w_{\mathbb{A}}$, and kind $k_{\mathbb{A}} \in \{\text{rd}, \text{wr}\}$ ($k_{\mathbb{A}} = \text{rd}$ for memory read and $k_{\mathbb{A}} = \text{wr}$ for memory write) reachable after traversing path $\xi'_{\mathbb{A}}$ starting from $n_{\mathbb{A}}$ is represented by tuple $(\xi'_{\mathbb{A}}, \alpha_{\mathbb{A}}, w_{\mathbb{A}}, k_{\mathbb{A}})$. Similarly, $(\xi_{\mathbb{C}}^j, \alpha_{\mathbb{C}}^j, w_{\mathbb{C}}^j, k_{\mathbb{C}}^j)$ for $1 \leq j \leq m$ represents j th memory access (out of m) on $\xi_{\mathbb{C}}$ of same kind as $k_{\mathbb{A}}$, i.e., $\forall_{1 \leq j \leq m} : k_{\mathbb{C}}^j = k_{\mathbb{A}}$. Let $\text{WP}_{\xi_P}(\alpha)$ represent *weakest precondition*[14] of α after executing ξ_P . An access $(\xi'_{\mathbb{A}}, \alpha_{\mathbb{A}}, w_{\mathbb{A}}, k_{\mathbb{A}})$ is *safe* iff the following holds:

$$\left(\begin{array}{l} \phi(n_{\mathbb{A}}, n_{\mathbb{C}}) \\ \wedge \text{pathcond}(\xi'_{\mathbb{A}}) \wedge \text{pathcond}([\xi_{\mathbb{C}}]_{\mathcal{D}_X}^{e_X}) \\ \wedge \bigwedge_{j=1}^m \text{WP}_{\xi_{\mathbb{C}}^j}(\neg \text{ov}([un]_{\text{PAGE_SIZE}}, [\alpha_{\mathbb{C}}^j]_{w_{\mathbb{C}}^j})) \end{array} \right) \implies \text{WP}_{\xi'_{\mathbb{A}}} \left(\begin{array}{l} \text{ov}([un]_{\text{PAGE_SIZE}}, [\alpha_{\mathbb{A}}]_{w_{\mathbb{A}}}) \\ \implies [\alpha_{\mathbb{A}}]_{w_{\mathbb{A}}} \subseteq \Sigma^{GFS} \end{array} \right) \quad (4.1)$$

where $\Sigma^{GFS} = \Sigma_{\mathbb{A}}^{GUF} \cup [\text{esp}, \text{stk}_e]$ if $k_{\mathbb{A}} = \text{rd}$, and $\Sigma^{GFS} = \Sigma_{\mathbb{A}}^{G_w \cup F_w} \cup [\text{esp}, \text{stk}_e]$ if $k_{\mathbb{A}} = \text{wr}$ (see procedure `accessIsSafe()` in algorithm 9).

It is easy to see that eq. (4.1) will not hold for an access $[\alpha_{\mathbb{A}}]_{w_{\mathbb{A}}}$ if it is neither matched by any of the accesses $[\alpha_{\mathbb{C}}^j]_{w_{\mathbb{C}}^j}$ for any j in $\xi_{\mathbb{C}}$ nor does it belong to the address set Σ^{GFS} .

4.1.7 Correlating paths to error nodes due to annotated instructions

Recall that the DYNAMO operates in two phases: in the first phase all paths in the original, non-annotated \mathbb{A} are annotated and correlated, and in the second phase the extra error-going paths introduced due to annotation are correlated.

The correlation of these extra paths follow the same procedure, shown as call to the same phase-parametric *constructX()* procedure at line 6 in algorithm 1. As the required annotation is assumed to have been performed in the first phase, in the second call to *constructX()* with *phase* = CORRELATE_NEW_ERROR_PATHS, the procedure skips the call to the *asmAnnotOpts()* procedure (line 23 in algorithm 1). Similarly, as an optimization, the procedure avoids correlating an already correlated $\xi_{\mathbb{A}}$ (line 17 in algorithm 1).

4.1.8 Soundness of DYNAMO algorithm

When all feasible simple paths between the cut points of \mathbb{A} are exhaustively correlated (including the paths introduced due to annotation), the (Coverage \mathbb{A}) requirement must be satisfied. The *checkCoverageReqs()* procedure further checks the satisfaction of (Coverage \mathbb{C}) before returning **Success** (line 8 in algorithm 1).

Recall that the path enumeration ensures (Mutex \mathbb{A}), (Mutex \mathbb{C}) and (Well-formedness) are satisfied by construction, the *addingEdgeWillCreateEmptyCCycle()* procedure ensures that (Similar-Speed) holds, and the *haveSimilarStructure()* procedure ensures that the remaining three structural requirements are also satisfied. The four non-coverage semantic requirements are checked in *checkSemanticReqsExceptCoverage()*. DYNAMO is sound because it returns **Success** only if all the thirteen search-algorithm requirements are satisfied.

4.1.9 Counterexample Guided Best-First Search

The **chooseFrom** operator must attempt to maximize the chances of returning **Success**, even if only a fraction of the search space has been explored. DYNAMO uses the counterexamples generated when a proof obligation is falsified (e.g., during invariant inference) to guide the search towards the more promising options. A counterexample is a proxy for the machine states of \mathbb{C} and \mathbb{A} that may appear at a node n_X during the

lockstep execution encoded by \mathbf{X} . Thus, if at any step during the construction of \mathbf{X} , the execution of a counterexample for a candidate partial solution $(\ddot{\mathbf{A}}, \mathbf{X})$ results in the violation of a non-coverage requirement (e.g., (MemEq)) that candidate is discarded and the search backtracks (instead of failing). The counterexample-guided best-first search exhausts the entire search space of correlations before returning **Failure**.

Our (MemEq) requirement is generalization of the heap relation⁹ pruning criteria of COUNTER [17]. In COUNTER, a candidate that does not have matching heap states at the end of counterexample execution is discarded. DYNAMO, through (MemEq), extends this to matching of memory states of stack-allocated locals as well. If the assembly procedure $\ddot{\mathbf{A}}$ does not have store-sinking like memory optimizations that cause divergence in the *non-stack* memory state of the two procedures, (MemEq) effectively prunes the search space without any loss in completeness guarantee.

The execution of counterexamples opportunistically weakens the node invariants in \mathbf{X} . Like COUNTER [17], we use the number of live registers in $\ddot{\mathbf{A}}$ related through the current invariants in $\Phi_{\mathbf{X}}$ to rank the enumerated partial candidate solutions to implement a best-first search. This ranking criteria is key to (relative) scaling of COUNTER and DYNAMO algorithms.

4.2 Invariant Inference

Our invariant inference procedure for identifying node invariants is a counterexample-guided algorithm, similar to COUNTER [17]. Candidate invariants of a partial product graph \mathbf{X} are formed by conjoining the predicates drawn from the predicate grammar shown in fig. 4.2. A candidate invariant $\boxed{\phi}$ at node $n_{\mathbf{X}}$ is validated using Hoare triples of the form: $\{\phi_{n_{\mathbf{X}}}\}(\xi_{\ddot{\mathbf{A}}}^{jk}; \xi_{\mathbf{C}}}^{jk})\{\boxed{\phi}\}$ for all $n_{\mathbf{X}}^j \in \mathcal{N}_{\mathbf{X}}$ and $(\xi_{\ddot{\mathbf{A}}}^{jk}, \xi_{\mathbf{C}}}^{jk})$ such that $e_{\mathbf{X}}^{jk} = (n_{\mathbf{X}}^j \xrightarrow{\xi_{\ddot{\mathbf{A}}}^{jk}; \xi_{\mathbf{C}}}^{jk}} n_{\mathbf{X}}) \in \mathcal{E}_{\mathbf{X}}$.

The predicate grammar includes shape for inferring affine relations ($\boxed{\text{affine}}$) between *interesting* bitvector expressions of $\ddot{\mathbf{A}}$ and \mathbf{C} — recall that the set of expressions participating in invariant inference are drawn from a set of *interesting* expressions in $\ddot{\mathbf{A}}$ and \mathbf{C} (*inferInvariantsAndCounterexamples()* in algorithm 7). Apart from $\boxed{\text{affine}}$, we also infer inequality relations through $\boxed{\text{ineq}}$ and $\boxed{\text{ineqC}}$ over these interesting

⁹Unlike DYNAMO, COUNTER does not distinguish between different global variables in the program and treat all memory as a single "heap".

$$\begin{array}{lll}
\boxed{\text{affine}} \quad \sum_i c_i v_i = c & \boxed{\text{ineqC}} \quad v \odot c & \boxed{\text{ineq}} \quad v_1 \odot v_2 \\
\boxed{\text{MemEq}} \quad M_C =_{\Sigma_{\ddot{A}}^B \setminus (\Sigma_{\ddot{A}}^{z_l} |^v)} M_{\ddot{A}} & \boxed{\text{AllocEq}} \quad \forall_{r \in B} : \Sigma_C^r = \Sigma_{\ddot{A}}^r & \boxed{\text{WEq}} \quad \Omega_{\ddot{A}} = \Omega_C \\
\boxed{\text{sp0rd}} \quad \boxed{\text{sp.p}_{\ddot{A}}^{j_1}} \leq_u (\boxed{\text{sp.p}_{\ddot{A}}^{j_2}} - v^*) & \boxed{\text{zEmpty}} \quad \{\Sigma_C^z, \Sigma_{\ddot{A}}^{z_l} |^s, \Sigma_{\ddot{A}}^{z_l} |^v\} \{=, \neq\} \emptyset & \\
\boxed{\text{spzBd}} \quad \boxed{\text{em.z}} \vee (\boxed{\text{sp.p}_{\ddot{A}}^j} \odot \{\boxed{\text{lb.z}}, \boxed{\text{ub.z}}\}) & \boxed{\text{spzBd}'} \quad \boxed{\text{em.z}} \vee (\boxed{\text{sp.p}_{\ddot{A}}^j} \leq_u (\boxed{\text{lb.z}} - v^*)) &
\end{array}$$

Figure 4.2: Predicate grammar for constructing candidate invariants. v represents a bitvector variable (registers, stack slots, and ghost variables), c represents a bitvector constant. $\odot \in \{\leq_{s,u}, <_{s,u}, >_{s,u}, \geq_{s,u}\}$. v^* represents a bitvector value drawn from a restricted grammar (explained in text).

expressions. An efficient counterexample-guided algorithm for computing affine relations over bitvectors in an incremental setting such as ours is described in [42].

The candidate invariants include shapes $\boxed{\text{MemEq}}$ and $\boxed{\text{AllocEq}}$ for equality of memory and allocation state of common regions across \ddot{A} and C and the $\boxed{\text{WEq}}$ shape for capturing the equality of outside world states. $\boxed{\text{MemEq}}$ and $\boxed{\text{WEq}}$ cater to the (MemEq) and (Equivalence) requirements respectively. Note that $\boxed{\text{AllocEq}}$ follows from DYNAMO’s \mathcal{D}_X construction (*addDetMappings()* in algorithm 1) and the execution semantics that observe each (de)allocation event. If the (de)allocations in both \ddot{A} and C were unobservable (and thus allowed to potentially differ indefinitely), invariant inference would become significantly harder, especially when equating memory regions of locals.

Recall that in our graph representations we save stackpointer value at the boundary of a stackpointer updating instruction at PC $p_{\ddot{A}}^j$ in ghost variable $\boxed{\text{sp.p}_{\ddot{A}}^j}$ ((OP-ESP) in fig. 2.6). These ghost variables make it convenient to express relationships between stack-allocated local regions. To prove separation between different local variables (allocated by different stackpointer decrements), we require invariants that lower-bound the gap between two ghost variables, say $\boxed{\text{sp.p}_{\ddot{A}}^{j_1}}$ and $\boxed{\text{sp.p}_{\ddot{A}}^{j_2}}$, by some value v^* that depends on the allocation size operand of an `allocs` instruction ($\boxed{\text{sp0rd}}$).

To capture the various relations between lower bounds, upper bounds, region sizes, and $\boxed{\text{sp.p}_{\ddot{A}}^j}$, the guessing grammar includes shapes $\boxed{\text{spzBd}}$ and $\boxed{\text{spzBd}'}$ that are of the form: “either a local variable region is empty or its bounds are related to $\boxed{\text{sp.p}_{\ddot{A}}^j}$ in these possible ways”. The “a local variable region is empty” part caters to the case of a conditional `alloca()` where n_X may have some incoming paths where allocation did not happen. Similar to $\boxed{\text{sp0rd}}$, the value v^* in $\boxed{\text{spzBd}'}$ is derived from the allocation

$$\begin{array}{l}
\boxed{\text{Empty}} \quad \forall_{r \in G \cup F \cup Y \cup Z} : (\Sigma_{\ddot{A}}^r = \emptyset \Leftrightarrow \boxed{\text{em.r}}) \\
\boxed{\text{gfySz}} \quad \forall_{r \in G \cup F \cup Y \setminus \{\text{vrdc}\}} : (\boxed{\text{sz.r}} = \text{sz}(\text{T}(r))) \quad \boxed{\text{vrdcSz}} \quad (\boxed{\text{em.vrdc}} \Leftrightarrow \boxed{\text{sz.vrdc}} = 0) \\
\boxed{\text{gfyIntvl}} \quad \forall_{r \in G \cup F \cup Y} : \left[\boxed{\text{em.r}} \vee \left(\begin{array}{l} (\boxed{\text{lb.r}} \leq_u \boxed{\text{ub.r}}) \wedge (\boxed{\text{lb.r}} + \boxed{\text{sz.r}} - 1_{i_{32}} = \boxed{\text{ub.r}}) \\ \wedge ([\boxed{\text{lb.r}}, \boxed{\text{lb.r}}] = \Sigma_{\ddot{A}}^r) \end{array} \right) \right] \\
\boxed{\text{zlIntvl}} \quad \boxed{\text{em.zl}} \vee \left(\begin{array}{l} (\boxed{\text{lb.zl}} \leq_u \boxed{\text{ub.zl}}) \wedge (\boxed{\text{lb.zl}} + \boxed{\text{lstSz.zl}} - 1_{i_{32}} = \boxed{\text{ub.zl}}) \\ \wedge ([\boxed{\text{lb.zl}}, \boxed{\text{lb.zl}}] = \Sigma_{\ddot{A}}^{zl}) \end{array} \right) \\
\boxed{\text{zaBd}} \quad \boxed{\text{em.za}} \vee \left(\begin{array}{l} (\boxed{\text{lb.za}} \leq_u \boxed{\text{ub.za}}) \wedge (\boxed{\text{lb.za}} + \boxed{\text{lstSz.za}} - 1_{i_{32}} \leq_u \boxed{\text{ub.za}}) \\ \wedge (\boxed{\text{lb.za}} = \text{lb}(\Sigma_{\ddot{A}}^{za}) \wedge \boxed{\text{ub.za}} = \text{ub}(\Sigma_{\ddot{A}}^{za})) \end{array} \right) \\
\boxed{\text{StkBd}} \quad \Sigma_{\ddot{A}}^{\{\text{stk}\} \cup Y} \cup (\Sigma_{\ddot{A}}^Z \setminus (\Sigma_{\ddot{A}}^{Z_l |^v})) = [\text{esp}, \boxed{\text{stk}_e}] \quad \boxed{\text{csBd}} \quad \Sigma_{\ddot{A}}^{\{\text{cs}, \text{cl}\}} = [\boxed{\text{stk}_e} + 1, \boxed{\text{cs}_e}] \\
\boxed{\text{NoOverlapC}} \quad \neg \text{ov}(\Sigma_{\ddot{A}}^{hp}, \Sigma_{\ddot{A}}^{cl}, \Sigma_{\ddot{A}}^{\text{vrdc}}, \dots, i_{\ddot{A}}^g, \dots, i_{\ddot{A}}^y, \dots, \Sigma_{\ddot{A}}^z) \\
\boxed{\text{NoOverlapA}} \quad \neg \text{ov}(\Sigma_{\ddot{A}}^{\{\text{hp}, \text{cl}\} \cup G \cup Y}, \dots, \Sigma_{\ddot{A}}^z |^s, \dots, i_{\ddot{A}}^f, \dots, \Sigma_{\ddot{A}}^{\text{stk}}, \Sigma_{\ddot{A}}^{\text{cs}}) \\
\boxed{\text{ROMA}} \quad \forall_{r \in F_r} : (M_{\ddot{A}} =_{i^r} \text{ROM}_{\ddot{A}}^r(i^r)) \quad \boxed{\text{ROMC}} \quad \forall_{r \in G_r} : (M_C =_{i^r} \text{ROM}_C^r(i^r))
\end{array}$$

Figure 4.3: Global invariants that hold at each non-entry, error-free node $n_X \in \mathcal{N}_X^{\text{DW}}$.

size operand of an `allocs` instruction. $\boxed{\text{zEmpty}}$ tracks the emptiness of the address set of a local region z in \mathbb{C} and the address sets $\Sigma_{\ddot{A}}^{zl |^s}$ and $\Sigma_{\ddot{A}}^{zl |^v}$ for $zl \in Z_l$ in \ddot{A} . We need to track the emptiness of the latter address sets in \ddot{A} to prove the infeasibility of the \mathcal{U} -going paths due to (DEALLOCV) and (DEALLOC'S') (sections 2.6.1 and 2.6.2).

Together, the predicate shapes $\boxed{\text{spOrd}}$, $\boxed{\text{spzBd}}$, $\boxed{\text{spzBd}'}$, $\boxed{\text{zEmpty}}$ (and $\boxed{\text{affine}}$ and $\boxed{\text{ineq}}$ for relations between $\boxed{\text{sp.p}_{\ddot{A}}^j}$), enable disambiguation between stack writes involving spilled pseudo-registers and stack-allocated locals.

4.2.1 Global Invariants

Recall that due to our execution semantics certain *global invariants* hold by construction at each non-entry, error-free node $n_X \in \mathcal{N}_X^{\text{DW}}$ in X (section 3.3.3). We add these global invariants to the set of node invariants ϕ_{n_X} at a node n_X , along with the inferred invariants described in previous section. We list the predicate expressions for these global invariants in fig. 4.3 and discuss each below.

- $\boxed{\text{Empty}}$ asserts that the ghost variable $\boxed{\text{em.r}}$ (for tracking the emptiness of region) for $r \in G \cup F \cup Y \cup Z$ tracks the emptiness of its address set.

- $\boxed{\text{gfySz}}$ equates the ghost variable $\boxed{\text{sz}.r}$ (for tracking size of region $r \in G \cup F \cup (Y \setminus \{\text{vrdc}\})$) to the size of the variable named r (recall that a region identifier is also the name of the variable)¹⁰.

$\boxed{\text{vrdcSz}}$ encodes that the $\boxed{\text{sz.vrdc}}$ is zero iff $\boxed{\text{em.vrdc}}$ holds for the variadic parameter $\text{vrdc} \in Y$.

- $\boxed{\text{gfyIntvl}}$ encodes that the address set of region $r \in G \cup F \cup Y$ is an interval of size $\boxed{\text{sz}.r}$ bounded by ghost variables $\boxed{\text{lb}.r}$ and $\boxed{\text{ub}.r}$. Note that only the vrdc region can potentially be empty such that $\boxed{\text{em}.r}$ holds for $r = \text{vrdc}$.

- $\boxed{\text{zlIntvl}}$ captures the property that a local variable region $zl \in Z_l$, if non-empty, must be an interval of size $\boxed{\text{lstSz}.zl}$ bounded by ghost variables $\boxed{\text{lb}.zl}$ and $\boxed{\text{ub}.zl}$.

$\boxed{\text{zlIntvl}}$ is encoding the same “ r is an interval” property presented in $\boxed{\text{gfyIntvl}}$ for $r \in Z_l$.

- $\boxed{\text{zaBd}}$ captures a weaker property (than $\boxed{\text{zlIntvl}}$) for a local region $za \in Z_a$ (recall that za is a local allocated using $\text{alloca}()$): if non-empty, this region must be bounded by its ghost variables ($\boxed{\text{lb}.zl}$ and $\boxed{\text{ub}.zl}$) and must be at least $\boxed{\text{lstSz}.za}$ large.

The difference from $\boxed{\text{zlIntvl}}$ is that the region za need not be an interval. This matches the expectation from intuition that multiple stack decrements in \ddot{A} corresponding to executions of an $\text{alloca}()$ in C need not be contiguous.

- $\boxed{\text{StkBd}}$ encodes the invariant that the interval $[\text{esp}, \boxed{\text{stk}_e}]$ represents the union of the address sets of stk , regions in Y , and stack-allocated local regions ($\Sigma_{\ddot{A}}^Z \setminus (\Sigma_{\ddot{A}}^{Z_l |^v})$).
- $\boxed{\text{csBd}}$ is similarly shaped as $\boxed{\text{StkBd}}$ and encodes that the interval $[\boxed{\text{stk}_e} + 1, \boxed{\text{cs}_e}]$ represents the union of the address sets of regions cs and cl .
- $\boxed{\text{NoOverlapC}}$ encodes the disjointedness of all regions $r \in B$ (recall that B denote the common regions present in both C and A).
- $\boxed{\text{NoOverlapA}}$ encodes the disjointedness of all regions in \ddot{A} except virtually-allocated regions. Note that $\boxed{\text{NoOverlapA}}$ does not encode disjointedness of regions $\{hp, cl\} \cup G \cup Y$ — we rely on $\boxed{\text{AllocEq}}$ and $\boxed{\text{NoOverlapC}}$ for this.

¹⁰Note that $\text{sz}(T(r)) > 0$ for all r because a variable can never have zero size in the C programming language.

- $\boxed{\text{ROMC}}$ and $\boxed{\text{ROMA}}$ encode the preservation of memory contents of read-only regions in \mathbb{C} and \mathbb{A} .

Let $\boxed{\phi_X}$ represent the conjunction of the global invariants listed in fig. 4.3. Let $\boxed{\phi_{n_X}}$ represent the conjunction of the inductively provable invariants drawn from grammar shown in fig. 4.2. The invariants ϕ_{n_X} at an error-free node $n_X \in \mathcal{N}_X$ are obtained through conjunction of $\boxed{\phi_X}$ and $\boxed{\phi_{n_X}}$, i.e., $\phi_{n_X} = \boxed{\phi_X} \wedge \boxed{\phi_{n_X}}$.

4.3 Running Example of the Algorithm

Figure 4.5 shows the abbreviated Transition Graphs of the unoptimized IR and fully annotated assembly procedure of the `fib` procedure from fig. 2.1, reproduced for convenience in fig. 4.4. A node in figs. 4.5a and 4.5b is identified by its PC and we use a subscript notation for the PCs due to constituent edges of an unoptimized IR/assembly instruction, e.g., the procedure call at I13 corresponds to graph edges $I13 \rightarrow I13_1$, $I13_1 \rightarrow I13_2$, $I13_2 \rightarrow I13_3$, and $I13_3 \rightarrow I14$. The abbreviated graphs retain only a subset of the nodes — roughly, we retain nodes corresponding to the PCs of the statements in the listings in fig. 4.4. We omit almost all edge labels except for the `rd` and `wr` instructions. We show transitions to error nodes in an abbreviated edge using edge labels, e.g., in fig. 4.5a, the transition to error node \mathcal{W}_C due to (ENTRY_C) is shown as the label $\rightarrow \mathcal{W}_C$ on the edge $I0_1 \rightarrow I0_2$. For an abbreviated edge due to $(\text{de})\text{alloc}$ and $(\text{de})\text{alloc}_{s,v}$ instruction, we show the constituent (still abbreviated) edges using an exploded view, e.g., the abbreviated edge $I1 \rightarrow I2$ corresponding to `alloc` instruction at I1 in fig. 4.4 is exploded into (still abbreviated) edges $I1 \rightarrow I1_1$, $I1_1 \rightarrow I1_2$, $I1_2 \rightarrow I1_3$, $I1_3 \rightarrow I1_4$, and $I1_4 \rightarrow I2$ in fig. 4.5a. In the exploded view for edges due to `alloc` at I9 and I10 in fig. 4.4b and due to `alloc_s` at A17¹ and A17², we omit the nodes as well and just retain the labels. Because the graphs are abbreviated we will refer both the unoptimized IR and assembly procedures in fig. 4.4 and the graphs in fig. 4.5 in our discussion below.

We show the execution of the DYNAMO algorithm on the two graphs in fig. 4.5 at unroll-factor $\mu = 2$. Note that the \mathbb{A} graph in fig. 4.5b shows the fully annotated graph. During our exposition, we will treat the graph as if it were not annotated till the point it becomes annotated in our discussion i.e., we will ignore the annotated edges until they are *inserted* by our execution of the algorithm. The algorithm begins with the points-to analysis on \mathbb{C} that over-approximates the β and β_M sets at each PC of \mathbb{C} .

```

    int printf(const char*, ...);

C0: int fib(int n, int m) {
C1:   int v[n+2];
C2:   v[0]=0; v[1]=1;
C3:   for(int i=2; i<=m; ++i)
C4:     v[i]=v[i-1]+v[i-2];
C5:   printf("fib(%d) = %d", m, v[m]);
C6:   return v[m];
C7: }

    (a) C program with VLA.

I0: int fib(int* n, int* m):
I1:   i=alloc 1, int, 4;
I2:   v=alloc *n+2, int, 4;
I3:   v[0]=0; v[1]=1;
I4:   *i=2;
I5:   if(*i >_s *m) goto I9;
I6:     v[*i]=v[*i-1]+v[*i-2];
I7:     ++(*i);
I8:     goto I5;
I9:   pI9=alloc 1, char*, 4;
I10:  pI10=alloc 1, struct{int;int;}, 4;
I11:  *pI9=__S__;
I12:  *pI10=*m; *(pI10+4)=v[*m];
I13:  t=call int printf(pI9, pI10);
I14:  dealloc I10;
I15:  dealloc I9;
I16:  r=v[*m];
I17:  dealloc I2;
I18:  dealloc I1;
I19:  ret r;

A0: fib:
A1:   push ebp; ebp = esp;
A2:   push {edi, esi, ebx};
A3:   esp -= 12;
A31: vI1 = allocv 4, 4, I1;
A4:   eax = mem4[ebp+8]; ebx = mem4[ebp+12];
A5:   esp -= 0xFFFFFFFF & (4*(eax+2)+15));
A51: allocs esp, 4*(eax+2), 4, I2;
A6:   esi = ((esp+3)>>2)*4;
A7:   mem4[esi] = 0; mem4[esi+4] = 1;
A8:   if(ebx ≤s 1) jmp A15;
A9:   edi = 0; edx = 1; eax = 2;
A10:  ecx = edx+edi;
A11:  edi = edx; edx = ecx;
A12:  mem4[esi+4*eax] = ecx;
A13:  ++eax;
A14:  if(eax ≤s ebx) jmp A10;
A15:  edi = mem4[esi+4*ebx];
A16:  esp -= 4;
A17:  push {edi, ebx, __S__};
A171: allocs esp, 4, 4, I9;
A172: allocs esp+4, 8, 4, I10;
A18:  call int printf
      (<char*> esp,
       <struct{int; int;}> esp+4)
      {hp, cl, I9, I10};
A181: deallocs I10;
A182: deallocs I9;
A19:  eax = edi;
A191: deallocs I2;
A192: deallocv I1;
A20:  esp = ebp-12;
A21:  pop {ebx, esi, edi, ebp};
A22:  ret;

    (b) (Abstracted) Unoptimized IR.

    (c) (Abstracted) 32-bit x86 assembly code.

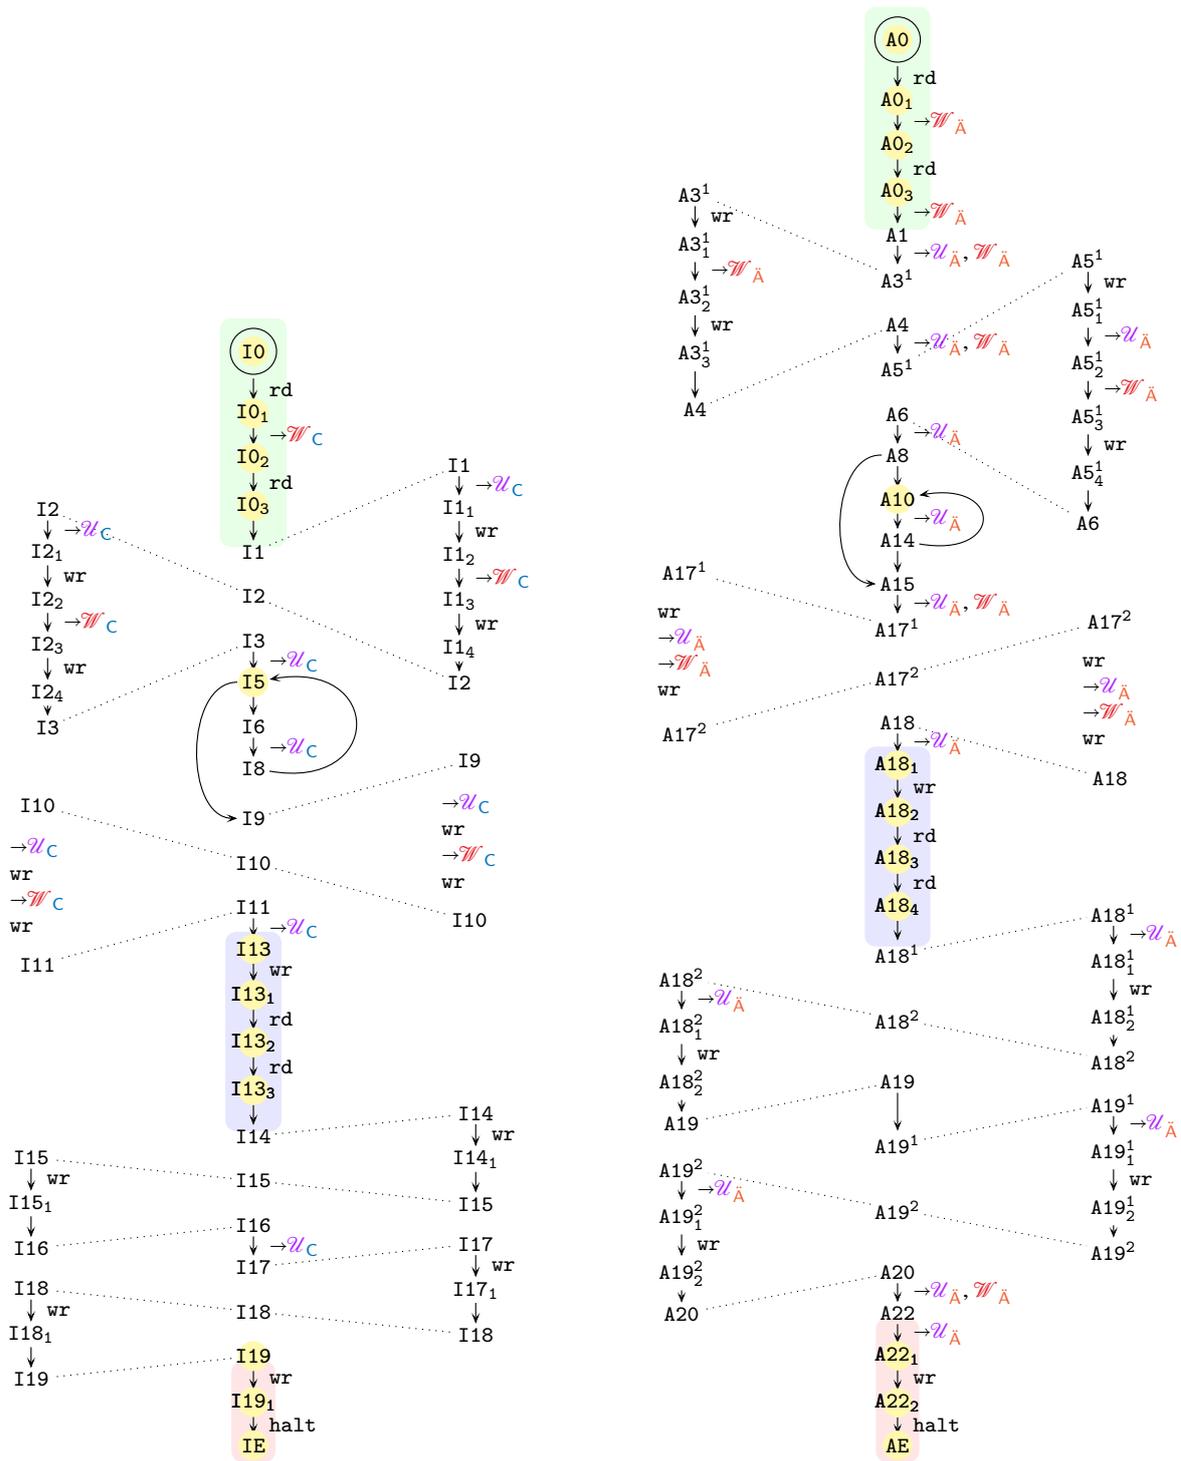
```

Figure 4.4: Reproduced C program and its unoptimized IR and assembly from fig. 2.1.

For example, at node I3 (or PC I3 in fig. 4.4b), $\beta(n) = \{n\}$, $\beta(m) = \{m\}$, $\beta(i) = \{I1\}$, $\beta(v) = \{I2\}$ and $\beta_M(n) = \beta_M(m) = \{hp, cl\}$, $\beta_M(I1) = \beta_M(I2) = \emptyset$, where n and m refer to both the state variables and regions $n, m \in Y$ and $I1, I2 \in Z_l$.

Enumerating A paths

After the points-to analysis, DYNAMO begins its first phase with the call to the *constructX()* procedure with parameter *phase* = CORRELATE_AND_ANNOTATE. *constructX()*



(a) Abbreviated Transition Graph for the unoptimized IR of the `fib` procedure from fig. 4.4b.

(b) Abbreviated Transition Graph for the compiler-generated assembly of the `fib` procedure from fig. 4.4c.

Figure 4.5: Abbreviated Transition Graphs for the unoptimized IR and assembly of the `fib` procedure from fig. 4.4.

identifies the set of cut-points in fig. 4.5b. A subset of cut-points nodes in figs. 4.5a and 4.5b are highlighted with  background¹¹. The cut-points nodes identified by *getCutPointsInRPO()* are: $A0, A0_1, A0_2, A0_3, A10, A18_1, A18_2, A18_3, A18_4, A22_1, A22_2$, and AE ¹². The correlation of (abbreviated) edges, $A0 \rightarrow A0_1, A0_1 \rightarrow A0_2$, and $A0_2 \rightarrow A0_3$, due to ($\text{ENTRY}_{\check{A}}$), is rather straightforward and we will instead consider the case when the cut-point node is $A0_3$, i.e., when $\mathcal{N}_X = \{(A0, I0), (A0_1, I0_1), (A0_2, I0_2), (A0_3, I0_3)\}$ and $\mathcal{E}_X = \{((A0, I0) \xrightarrow{(A0 \rightarrow A0_1); (I0 \rightarrow I0_1)} (A0_1, I0_1)), ((A0_1, I0_1) \xrightarrow{(A0_1 \rightarrow A0_2); (I0_1 \rightarrow I0_2)} (A0_2, I0_2)), ((A0_2, I0_2) \xrightarrow{(A0_2 \rightarrow A0_3); (I0_2 \rightarrow I0_3)} (A0_3, I0_3))\}$.

The (ordered) cut-point successors of $A0_3$, as identified by *cutPointSuccessorsRPO()*, are: $A10, A18_1, \mathcal{U}_{\check{A}}$, and $\mathcal{W}_{\check{A}}$ ¹³. The cut-point $A10$, due to the cycle $A10 \rightarrow A14 \rightarrow A10$, is considered first as it precedes others in the order. The only simple path between $A0_3$ and $A10$ is the path $(A0_3 \rightarrow A10) = A0_3 \rightarrow A1 \rightarrow A4 \rightarrow A6 \rightarrow A8 \rightarrow A10$. As $(A0_3 \rightarrow A10)$ is not demonstrably infeasible, it is considered for the next step where its \mathbf{C} pathset options are enumerated.

Correlating \mathbf{C} paths

The *correlatedPathsInCOpts()* procedure (algorithm 3) is called with arguments $\xi_{\check{A}} = (A0_3 \rightarrow A10)$ and $\mu = 2$. The unique X node for $A0_3$ in \mathcal{N}_X is $(A0_3, I0_3)$. As the sink $A10$ of $(A0_3 \rightarrow A10)$ is an error-free node, the third case in *correlatedPathsInCOpts()* is triggered (line 21 in algorithm 3):

1. First, the set of *all* error-free cut-points reachable from $I0_3$ are identified (*getAllNonErrorReachableCutPoints()*). In fig. 4.5a, these are: $I1_1, I1_2, I1_3, I1_4, I2_1, I2_2, I2_3, I2_4, I5, I13, I13_1, I13_2, I13_3, I19, I19_1, IE, \dots$ (omitting the ones due to $(\text{de})\text{alloc}$ instructions at $I9, I10, I14, I15, I17$, and $I18$).
2. For each cut-point $q_{\mathbf{C}}^t$ enumerated in previous step, the set of pathsets from $I0_3$ to $q_{\mathbf{C}}^t$ with unrolling up to $\mu = 2$ are enumerated using *enumPathsetsTillUnroll()*. For $q_{\mathbf{C}}^t = I5$, the possible pathsets returned by *enumPathsetsTillUnroll()* are the singleton sets $\{I0_3 \rightarrow I1 \rightarrow I2 \rightarrow I3 \rightarrow I5\}$ and $\{I0_3 \rightarrow I1 \rightarrow I2 \rightarrow I3 \rightarrow I5 \rightarrow I6 \rightarrow I8 \rightarrow I5\}$ ¹⁴ — notice that the latter has two unrolling of node $I5$.

¹¹All cut-point nodes except those due to $(\text{de})\text{alloc}$ and $(\text{de})\text{alloc}_{s,v}$ instructions are highlighted in figs. 4.5a and 4.5b.

¹²As established earlier, we will ignore the inserted edges at this point.

¹³Recall that the nodes between $A3^1$ and $A4$ and between $A5^1$ and $A6$ will not be considered because they are inserted later due to annotation.

¹⁴Here and henceforth, we will omit the edges due to $(\text{de})\text{alloc}$ for brevity.

3. For each pathset $\langle \xi \rangle_{\mathcal{C}}^t$ returned by `enumPathsetsTillUnroll()`, set of its all \mathcal{U} -maximal subsets are added as candidate options. A candidate pathset added due to the singleton pathset $\langle \xi \rangle_{\mathcal{C}}^t = \{I0_3 \rightarrow I1 \rightarrow I2 \rightarrow I3 \rightarrow I5\}$ includes the paths: $I0_3 \rightarrow I1 \rightarrow \mathcal{U}_{\mathcal{C}}$, $I0_3 \rightarrow I1 \rightarrow I2 \rightarrow \mathcal{U}_{\mathcal{C}}$, $I0_3 \rightarrow I1 \rightarrow I2 \rightarrow I3 \rightarrow \mathcal{U}_{\mathcal{C}}$, and $I0_3 \rightarrow I1 \rightarrow I2 \rightarrow I3 \rightarrow I5$, where the $\mathcal{U}_{\mathcal{C}}$ going paths due to the statement I3 in fig. 4.4b have been abbreviated to $I0_3 \rightarrow I1 \rightarrow I2 \rightarrow I3 \rightarrow \mathcal{U}_{\mathcal{C}}$ for brevity.

In `constructX()`, let the **chooseFrom** operator choose the pathset containing the paths $\{I0_3 \rightarrow I1 \rightarrow \mathcal{U}_{\mathcal{C}}, I0_3 \rightarrow I1 \rightarrow I2 \rightarrow \mathcal{U}_{\mathcal{C}}, I0_3 \rightarrow I1 \rightarrow I2 \rightarrow I3 \rightarrow \mathcal{U}_{\mathcal{C}}, I0_3 \rightarrow I1 \rightarrow I2 \rightarrow I3 \rightarrow I5\}$ as the candidate pathset $\langle \xi \rangle_{\mathcal{C}}$ for correlation with $\xi_{\check{A}} = (A0_3 \rightarrow A10)$. Each path $\xi_{\mathcal{C}} \in \langle \xi \rangle_{\mathcal{C}}$ is considered separately; we demonstrate the correlation of $\xi_{\mathcal{C}} = I0_3 \rightarrow I1 \rightarrow I2 \rightarrow I3 \rightarrow I5$. As `phase = CORRELATE_AND_ANNOTATE`, `constructX()` calls `asmAnnotOpts()` for potentially annotating $\xi_{\check{A}} = (A0_3 \rightarrow A10)$ (line 23 in algorithm 1).

Identifying A annotation

The `asmAnnotOpts()` procedure begins by collecting the sequence of (edges due to) `alloc`, `dealloc`, and `call` instructions in $\xi_{\mathcal{C}} = I0_3 \rightarrow I1 \rightarrow I2 \rightarrow I3 \rightarrow I5$. The (abbreviated) edges $I1 \rightarrow I2$ and $I2 \rightarrow I3$ correspond to the `alloc` instructions at I1 and I2 respectively in fig. 4.4b. For each `alloc` instruction, `asmAnnotOpts()` enumerates `allocs,v` annotation options in the path $\xi_{\check{A}} = A0_3 \rightarrow A1 \rightarrow A4 \rightarrow A6 \rightarrow A8 \rightarrow A10$.

In the blackbox mode, the enumeration is dictated by the restrictions described in section 4.1.3. The first restriction confines the position of an `allocs,v` annotation to just after a stackpointer updating instruction — in our case (fig. 4.4c), the candidates are the PCs just after the statements A1, A2, A3, and A5. The second restriction confines the options for start address in an `allocs` annotation to the stackpointer value `esp` at the respective position — for each of the positions enumerated above, the corresponding stackpointer values are represented by the ghost variables `sp.A1`, `sp.A2`, `sp.A3`, and `sp.A5`¹⁵. The other parameters for the `allocs,v` instructions are inferred from the corresponding `alloc` instruction, e.g. for the `alloc` instruction at I2 (fig. 4.4b), the size of allocation is inferred to be $4*(\text{eax}+2)$ (inferred using the relation $\text{eax} = (*\text{n}) = \text{sel}_4(\text{mem}, \text{ebp} + 8)$ ¹⁶), the required alignment to be 4 (from alignment of

¹⁵Recall that the translation rule (OP-ESP) for a stackpointer updating instruction at PC $p_{\check{A}}^j$ stores the stackpointer value at the end of the update to a ghost variable `sp.pcheck{A}j` identified uniquely by $p_{\check{A}}^j$.

¹⁶A more general method is to build a \mathcal{D}_X -like map for enabling propagation of information from \mathcal{C} to \check{A} . Such a construction is enabled due to lockstep correlation of `alloc` and `allocs,v` in X .

int), and the region identifier to be I2. Thus, the annotation ‘`allocv 4, 4, I1`’ after A3 for `alloc` at I1 and ‘`allocs [sp.A5], 4 * (eax + 2), 4, I2`’ after A5 for `alloc` at I2 will be enumerated by the blackbox enumeration procedure (`genAnnotOptsForAlloc()` in algorithm 4). The `asmAnnotOpts()` procedure considers all combinations of individual annotation options, one of which will include the combined annotation ‘`allocv 4, 4, I1`’ after A3 and ‘`allocs [sp.A5], 4 * (eax + 2), 4, I2`’ after A5 as an annotation candidate. The application of the above annotation candidate inserts the edges $A3^1 \rightarrow A3_1^1 \rightarrow A3_2^1 \rightarrow A3_3^1 \rightarrow A4$ and $A5^1 \rightarrow A5_1^1 \rightarrow A5_2^1 \rightarrow A5_3^1 \rightarrow A5_4^1 \rightarrow A6$ to \ddot{A} as shown in fig. 4.5b. Note that the updated $\xi_{\ddot{A}}$, $\xi_{\ddot{A}}^* = A0_3 \rightarrow A1 \rightarrow A3^1 \rightarrow A4 \rightarrow A5^1 \rightarrow A6 \rightarrow A8 \rightarrow A10$ ¹⁷, obtained after including the annotated edges, does not include the $\mathcal{U}_{\ddot{A}}$ and $\mathcal{W}_{\ddot{A}}$ going paths added due to annotation — these paths are correlated in the second phase (when `phase = CORRELATE_NEW_ERROR_PATHS`).

Transforming identified paths to correlation paths

After annotation, `constructX()` breaks $\xi_{\ddot{A}}$ and ξ_C into constituent paths using `breakIntoSingleIOPaths()`, so that each I/O path (section 3.1) appears by itself and not as sub-path of some larger path. The path sequences obtained after both calls to `breakIntoSingleIOPaths()` are $\vec{\xi}'_{\ddot{A}} = ((A0_3 \rightarrow A3^1), (A3^1 \rightarrow A3_1^1), (A3_1^1 \rightarrow A3_2^1), (A3_2^1 \rightarrow A3_3^1), (A3_3^1 \rightarrow A5^1), (A5^1 \rightarrow A5_1^1), (A5_1^1 \rightarrow A5_2^1), (A5_2^1 \rightarrow A5_3^1), (A5_3^1 \rightarrow A5_4^1), (A5_4^1 \rightarrow A10))$ and $\vec{\xi}'_C = ((I0_3 \rightarrow I1_1), (I1_1 \rightarrow I1_2), (I1_2 \rightarrow I1_3), (I1_3 \rightarrow I1_4), (I1_4 \rightarrow I2_1), (I2_1 \rightarrow I2_2), (I2_2 \rightarrow I2_3), (I2_3 \rightarrow I2_4), (I2_4 \rightarrow I5))$.

As neither ξ_C nor $\xi_{\ddot{A}}$ ends at an error node, the call to `trimToMatchPathToErrorNode()` (algorithm 5) will not perform any trimming. For demonstrating `trimToMatchPathToErrorNode()`, we consider the case when for the same $\xi_{\ddot{A}}$, $\xi_C = I0_3 \rightarrow I1 \rightarrow I2 \rightarrow \mathcal{U}_C$ is correlated so that $\vec{\xi}'_C$ (obtained after `breakIntoSingleIOPaths()`) comes out to be $((I0_3 \rightarrow I1_1), (I1_1 \rightarrow I1_2), (I1_2 \rightarrow I1_3), (I1_3 \rightarrow I1_4), (I1_4 \rightarrow \mathcal{U}_C))$. In this case, $\vec{\xi}'_C$ ends at error node \mathcal{U}_C and the two paths sequences have difference lengths ($|\vec{\xi}'_{\ddot{A}}| = 9$ and $|\vec{\xi}'_C| = 5$). `trimToMatchPathToErrorNode()` trims $\vec{\xi}'_{\ddot{A}}$ to match the length of the error going path sequence $\vec{\xi}'_C$, i.e., we obtain $\vec{\xi}_{\ddot{A}}^* = ((A0_3 \rightarrow A3^1), (A3^1 \rightarrow A3_1^1), (A3_1^1 \rightarrow A3_2^1), (A3_2^1 \rightarrow A3_3^1), (A3_3^1 \rightarrow A5^1))$ (dropping the last four path components from $\vec{\xi}'_{\ddot{A}}$) and $\vec{\xi}'_C$ remains identical to $\vec{\xi}'_C$. Effectively, with `trimToMatchPathToErrorNode()`, the lockstep execution of either of the error-free path in an X edge is restricted to the point where the other path encounters error.

¹⁷Omitting the constituent edges between $A3^1$ and $A4$, and $A5^1$ and $A6$.

Validating the structure of the paths

$constructX()$ validates the structure of $\vec{\xi}_{\bar{A}}^*$ and $\vec{\xi}_{\bar{C}}^*$ using $haveSimilarStructure()$. The $haveSimilarStructure()$ predicate verifies that both path sequences $\vec{\xi}_{\bar{A}}^*$ and $\vec{\xi}_{\bar{C}}^*$ have identical number of paths and each coupled path-pair satisfies the (SingleIO), (Safety), and (Termination) structural requirements. With $\vec{\xi}_{\bar{A}}^* = ((A0_3 \rightarrow A3^1), (A3^1 \rightarrow A3_1^1), (A3_1^1 \rightarrow A3_2^1), (A3_2^1 \rightarrow A3_3^1), (A3_3^1 \rightarrow A5^1), (A5^1 \rightarrow A5_1^1), (A5_1^1 \rightarrow A5_3^1), (A5_3^1 \rightarrow A5_4^1), (A5_4^1 \rightarrow A10))$ and $\vec{\xi}_{\bar{C}}^* = ((I0_3 \rightarrow I1_1), (I1_1 \rightarrow I1_2), (I1_2 \rightarrow I1_3), (I1_3 \rightarrow I1_4), (I1_4 \rightarrow I2_1), (I2_1 \rightarrow I2_2), (I2_2 \rightarrow I3), (I2_3 \rightarrow I2_4), (I2_4 \rightarrow I5))$, the (Safety) and (Termination) requirements hold trivially and it is easy to see that the I/O paths of same kind are coupled together for the (SingleIO) requirement, e.g., the I/O path $(A3^1 \rightarrow A3_1^1)$ that corresponds to the $wr(allocBegin())$ instruction due to $alloc_v$ (see (ALLOCV)) is paired with the path $(I1_1 \rightarrow I1_2)$ that also corresponds to a $wr(allocBegin())$ instruction due to $alloc$ (see (ALLOC)).

Adding X edges and updating \mathcal{D}_X

After validating the structural requirements, the identically-sized path sequences $\vec{\xi}_{\bar{A}}^*$ and $\vec{\xi}_{\bar{C}}^*$ are deconstructed into path pairs that form an X edge and added to \mathcal{E}_X . The $addingEdgeWillCreateEmptyCCycle()$ check, performed for each path pair, holds for each path pair in our example as none of the paths are empty. For the path pairs, $((A3_1^1 \rightarrow A3_2^1), (I1_2 \rightarrow I1_3))$ and $((A5_1^1 \rightarrow A5_3^1), (I2_2 \rightarrow I2_3))$ that contain the *choose* instruction (θ), the $addDetMappings()$ procedure updates \mathcal{D}_X to include the mappings for the *choose* instructions in $(I1_2 \rightarrow I1_3)$ and $(I2_2 \rightarrow I2_3)$. For the \bar{C} path $(I1_2 \rightarrow I1_3)$, two mappings are added, one for the address and another for memory, $\mathcal{D}_X(((A3_1^1 \rightarrow A3_2^1), (I1_2 \rightarrow I1_3)), (I1_2 \rightarrow I1_3^a), 1) = (v_{I1})$ and $\mathcal{D}_X(((A3_1^1 \rightarrow A3_2^1), (I1_2 \rightarrow I1_3)), (I1_2 \rightarrow I1_3^m), 1) = (M_{\bar{A}})$ where $(I1_2 \rightarrow I1_3^a)$ is the edge for $\theta(i_{32})$ instruction (in the abbreviated edge $(I1_2 \rightarrow I1_3)$) and $(I1_2 \rightarrow I1_3^m)$ is the edge for $\theta(i_{32} \rightarrow i_8)$ instruction. Similarly, for $(I2_2 \rightarrow I2_3)$, the following mappings are added: $\mathcal{D}_X(((A5_1^1 \rightarrow A5_3^1), (I2_2 \rightarrow I2_3)), (I2_2 \rightarrow I2_3^a), 1) = (sp.A5)$ and $\mathcal{D}_X(((A5_1^1 \rightarrow A5_3^1), (I2_2 \rightarrow I2_3)), (I2_2 \rightarrow I2_3^m), 1) = (M_{\bar{A}})$.

Invariant inference and checking semantic requirements

Next, DYNAMO performs invariant inference to update the invariant network Φ_X . Table 4.1 shows some of the inferred invariants for the fully-constructed X. We discuss some of the key invariants below.

n_X	ϕ_{n_X}
(A3 ₃ ¹ , I1 ₄)	$\textcircled{1} n = \text{lb.n} = \text{sp.entry} + 4$ $\textcircled{2} \text{sp.entry} = \text{ebp} + 4$ $\textcircled{3} \text{sp.A3} = \text{esp}$ $\textcircled{4} m = \text{lb.m} = \text{sp.entry} + 8$ $\textcircled{5} \text{esp} * 2^{28} = 0$ $\textcircled{6} \text{sp.A3} = \text{ebp} - 24$ $\textcircled{7} i = \text{lb.I1} = v_{I1}$ $\textcircled{8} \text{lstSz.I1} = 4$ $\textcircled{9} \text{em.I1} = \text{false}$ $\textcircled{10} \Sigma_{\bar{A}}^{I1 s} = \Sigma_{\bar{C}}^{I2} = \Sigma_{\bar{C}}^{I9} = \Sigma_{\bar{C}}^{I10} = \emptyset$ $\textcircled{11} \text{stk}_e = \text{ebp} + 15$ $\textcircled{12} \text{esp} \leq_u \text{stk}_e$ $\textcircled{13} M_{\bar{C}} =_{\Sigma_{\bar{A}}^B \setminus (\Sigma_{\bar{A}}^{Z1 v})} M_{\bar{A}}$ $\textcircled{14} \forall r \in B : \Sigma_{\bar{C}}^r = \Sigma_{\bar{A}}^r$ $\textcircled{15} \Omega_{\bar{A}} = \Omega_{\bar{C}}$
(A5 ¹ , I2 ₁)	$\textcircled{1} n = \text{lb.n} = \text{sp.entry} + 4$ $\textcircled{2} (\text{sp.entry} + 4) * 2^{28} = 0$ $\textcircled{3} \text{eax} = *n$ $\textcircled{4} m = \text{lb.m} = \text{sp.entry} + 8$ $\textcircled{5} \text{esp} * 2^{28} = 0$ $\textcircled{6} \text{ebx} = *m$ $\textcircled{7} \text{sp.entry} = \text{ebp} + 4$ $\textcircled{8} \text{sp.A3} = \text{ebp} - 24$ $\textcircled{9} \text{esp} = \text{sp.A5}$ $\textcircled{10} \Sigma_{\bar{A}}^{I1 s} = \Sigma_{\bar{C}}^{I2} = \Sigma_{\bar{C}}^{I9} = \Sigma_{\bar{C}}^{I10} = \emptyset$ $\textcircled{11} \text{sp.A5} \leq_u \text{sp.A3} - 4 * (\text{eax} + 2)$ $\textcircled{12} M_{\bar{C}} =_{\Sigma_{\bar{A}}^B \setminus (\Sigma_{\bar{A}}^{Z1 v})} M_{\bar{A}}$ $\textcircled{13} \forall r \in B : \Sigma_{\bar{C}}^r = \Sigma_{\bar{A}}^r$ $\textcircled{14} \Omega_{\bar{A}} = \Omega_{\bar{C}}$
(A10, I5)	$\textcircled{1} v = \text{lb.I2} = \text{esi}$ $\textcircled{2} v[*i - 1] = \text{edi}$ $\textcircled{3} v[*i - 2] = \text{edx}$ $\textcircled{4} m = \text{lb.m}$ $\textcircled{5} \text{ebx} = *m$ $\textcircled{6} \text{eax} = *i$ $\textcircled{7} \Sigma_{\bar{A}}^{I1 s} = \Sigma_{\bar{A}}^{I2 v} = \Sigma_{\bar{C}}^{I9} = \Sigma_{\bar{C}}^{I10} = \emptyset$ $\textcircled{8} \text{sp.entry} = \text{ebp} + 4$ $\textcircled{9} \text{esp} * 2^{28} = 0$ $\textcircled{10} M_{\bar{C}} =_{\Sigma_{\bar{A}}^B \setminus (\Sigma_{\bar{A}}^{Z1 v})} M_{\bar{A}}$ $\textcircled{11} \forall r \in B : \Sigma_{\bar{C}}^r = \Sigma_{\bar{A}}^r$ $\textcircled{12} \Omega_{\bar{A}} = \Omega_{\bar{C}}$
(A18 ₂ ² , I15 ₁)	$\textcircled{1} \Sigma_{\bar{A}}^{I1 s} = \Sigma_{\bar{A}}^{I2 v} = \Sigma_{\bar{C}}^{I9} = \Sigma_{\bar{C}}^{I10} = \emptyset$ $\textcircled{2} \text{sp.entry} = \text{ebp} + 4$ $\textcircled{3} v[*m] = \text{edi}$ $\textcircled{4} M_{\bar{C}} =_{\Sigma_{\bar{A}}^B \setminus (\Sigma_{\bar{A}}^{Z1 v})} M_{\bar{A}}$ $\textcircled{5} \forall r \in B : \Sigma_{\bar{C}}^r = \Sigma_{\bar{A}}^r$ $\textcircled{6} \Omega_{\bar{A}} = \Omega_{\bar{C}}$

Table 4.1: Some of the inferred inductive node invariants for the product graph X of the two procedures shown in fig. 4.5. $*v$ is short for $\text{sel}_4(M_{\bar{C}}, v)$, e.g., $*i = \text{sel}_4(M_{\bar{C}}, i)$, $*m = \text{sel}_4(M_{\bar{C}}, m)$ and so on.

- At node (A5¹, I2₁) (that corresponds to start of `alloc` and `allocv` in unoptimized IR and assembly respectively), the invariants, $\textcircled{9} \text{esp} = \text{sp.A5}$, $\textcircled{11} \text{sp.A5} \leq_u \text{sp.A3} - 4 * (\text{eax} + 2)$, and $\textcircled{3} \text{eax} = *n$ are consequential in proof of the clause $[v, v + w - 1]_{i_{32}} \in \text{intrvlInSet}_a(v, v + w - 1]_{i_{32}}, \Sigma_{\bar{A}}^{stk})$ check due to (ALLOCS). The proof of $\phi = \text{intrvlInSet}_a(v, v + w - 1]_{i_{32}}, \Sigma_{\bar{A}}^{stk})$ makes the path (A5₁¹ \rightarrow $\mathcal{U}_{\bar{A}}$) with path condition $\neg\phi$ infeasible.
- At node (A10, I5) (that corresponds to loop head in both unoptimized IR and assembly), the invariant $\textcircled{9} \text{esp} * 2^{28} = 0$, which implies that the stackpointer `esp` is aligned by 16, helps in falsifying the check `-aligned16(esp)` due to procedure call at A18 in fig. 4.4c (see (CALL_A) in fig. 2.8), thereby making the $\mathcal{U}_{\bar{A}}$ -going path (A18 \rightarrow $\mathcal{U}_{\bar{A}}$) infeasible.
- At node (A18₂², I15₁) (that corresponds to deallocation of first argument to `printf`), the invariant $\textcircled{1} \Sigma_{\bar{A}}^{I2|v} = \emptyset$ helps in proving the infeasibility of the $\mathcal{U}_{\bar{A}}$ -going path

due to (ALLOCS').

After invariant inference, the non-coverage semantics requirements are checked through *checkSemanticReqsExceptCoverage()*. The checks for (Equivalence) and (MemEq) involve ensuring that the predicates corresponding to instantiations of $\boxed{\text{WEq}}$ and $\boxed{\text{MemEq}}$ are present in the inferred invariants ϕ_{n_X} . For example, table 4.1 contains instantiations of $\boxed{\text{WEq}}$ and $\boxed{\text{MemEq}}$ at each node; the presence of these invariant shapes is sufficient to ensure (Equivalence) and (MemEq). For the (Memory Address Correspondence) or (MAC) check, we take the example of the path pair $((A3_3^1 \rightarrow A5^1), (I1_4 \rightarrow I2_1))$ that has two memory accesses in \ddot{A} (at A4 in fig. 4.4c): $\text{mem}_4(\text{ebp} + 8)$ and $\text{mem}_4(\text{ebp} + 12)$. There are no memory accesses on the path $(I1_4 \rightarrow I2_1)$, thus, for (MAC) it must be established that these memory accesses belong to the address set $\Sigma_{\ddot{A}}^{GUF} \cup [\text{esp}, \text{stk}_e]$. This is easily provable because $[\text{ebp} + 8, \text{ebp} + 15] \subseteq \Sigma_{\ddot{A}}^{\{n,m\}} \subseteq [\text{esp}, \text{stk}_e]$ is provable over the path $(A3_3^1 \rightarrow A5^1)$ due to invariants $\text{esp} = \boxed{\text{sp.A3}} = \text{ebp} - 24$, $\boxed{\text{stk}_e} = \text{ebp} + 15$, and $\text{esp} \leq_u \boxed{\text{stk}_e}$ at $(A3_3^1, I1_4)$.

Once the first phase has finished, DYNAMO invokes the second phase where the error-going paths that were added due to annotation are correlated. In fig. 4.5b, these are the $\mathcal{U}_{\ddot{A}}$ and $\mathcal{W}_{\ddot{A}}$ -going paths originating from nodes $A3_1^1, A5_1^1, A5_2^1, A17^1, A17^2, A18^1, A18^2, A19^1$, and $A19^2$.

In the last step, DYNAMO checks the (CoverageC) requirement for the added edges. We take the example of the loop path correlation here. The \ddot{A} loop path $\xi_{\ddot{A}} = A10 \rightarrow A14 \rightarrow A10$ is correlated with the paths $\xi_C^1 = I5 \rightarrow I6 \rightarrow I8 \rightarrow I5$ and $\xi_C^2 = I5 \rightarrow I6 \rightarrow \mathcal{U}_C^{18}$, to produce the two edges, $e_X^1 = ((A10, I5) \xrightarrow{A10 \rightarrow A14 \rightarrow A10; I5 \rightarrow I6 \rightarrow I8 \rightarrow I5} (A10, I5))$ and $e_X^2 = ((A10, I5) \xrightarrow{A10 \rightarrow A14 \rightarrow A10; I5 \rightarrow I6 \rightarrow \mathcal{U}_C} (A10, \mathcal{U}_C))$. The (CoverageC) obligation involves proving that the path cover $\{e_X^1, e_X^2\} \langle \mathcal{D}_X, \xi_{\ddot{A}} \rangle$, which is equivalent to the Hoare triple $\{\phi_{(A10, I5)}\}(\xi_{\ddot{A}}; \epsilon) \{ \text{pathcond}([\xi_C^1]_{\mathcal{D}_X}^{e_X^1}) \vee \text{pathcond}([\xi_C^2]_{\mathcal{D}_X}^{e_X^2}) \}$, holds. The invariants, ⑤ $\text{ebx} = *m$ and ⑥ $\text{eax} = *i$ at $(A10, I5)$ in table 4.1 are adequate to successfully discharge this Hoare triple.

Table 4.2 shows the edges \mathcal{E}_X in the final product graph X for the graphs in fig. 4.5 — $\text{alloc}_{s,v}$ paths not shown in fig. 4.5b have been omitted and the \ddot{A} and C paths have been abbreviated for brevity as done in fig. 4.5b. Notice that certain paths to $\mathcal{U}_{\ddot{A}}$ in \ddot{A} are not in \mathcal{E}_X because they were proven infeasible, e.g. $(A1 \rightarrow \mathcal{U}_A)$, $(A5_1^1 \rightarrow \mathcal{U}_A)$, $(A18 \rightarrow \mathcal{U}_A)$, $(A20 \rightarrow \mathcal{U}_A)$, $(A22 \rightarrow \mathcal{U}_A)$ and so on.

¹⁸Representing all \mathcal{U}_C -going paths using a single $I5 \rightarrow I6 \rightarrow \mathcal{U}_C$ path.

n_X	n_X^t	$\xi_{\bar{A}}$	ξ_C
(A0, I0)	(A0 ₁ , I0 ₁)	A0 → A0 ₁	I0 → I0 ₁
(A0 ₁ , I0 ₁)	(A0 ₂ , I0 ₂)	A0 ₁ → A0 ₂	I0 ₁ → I0 ₂
(A0 ₁ , I0 ₁)	(\mathcal{W}_A , I0 ₁)	A0 ₁ → \mathcal{W}_A	ϵ
(A0 ₂ , I0 ₂)	(A0 ₃ , I0 ₃)	A0 ₂ → A0 ₃	I0 ₂ → I0 ₃
(A0 ₃ , I0 ₃)	(A3 ¹ , I1 ₁) (A3 ¹ , \mathcal{U}_C)	A0 ₃ → A1 → A3 ¹	I0 ₃ → I1 → I1 ₁ I0 ₃ → I1 → \mathcal{U}_C
(A0 ₃ , I0 ₃)	(\mathcal{W}_A , I0 ₃)	A0 ₃ → \mathcal{W}_A A0 ₃ → A1 → \mathcal{W}_A	ϵ
(A3 ¹ , I1 ₁)	(A3 ₁ ¹ , I1 ₂)	A3 ¹ → A3 ₁ ¹	I1 ₁ → I1 ₂
(A3 ₁ ¹ , I1 ₂)	(A3 ₂ ¹ , I1 ₃)	A3 ₁ ¹ → A3 ₂ ¹	I1 ₂ → I1 ₃
(A3 ₁ ¹ , I1 ₂)	(\mathcal{W}_A , I1 ₂)	A3 ₁ ¹ → \mathcal{W}_A	ϵ
(A3 ₂ ¹ , I1 ₃)	(A3 ₃ ¹ , I1 ₄)	A3 ₂ ¹ → A3 ₃ ¹	I1 ₃ → I1 ₄
(A3 ₃ ¹ , I1 ₄)	(A5 ¹ , I2 ₁) (A5 ¹ , \mathcal{U}_C)	A3 ₃ ¹ → A4 → A5 ¹	I1 ₄ → I2 → I2 ₁ I1 ₄ → I2 → \mathcal{U}_C
(A3 ₃ ¹ , I1 ₄)	(\mathcal{W}_A , I2 ₁) (\mathcal{W}_A , \mathcal{U}_C)	A3 ₃ ¹ → A4 → \mathcal{W}_A	I1 ₄ → I2 → I2 ₁ I1 ₄ → I2 → \mathcal{U}_C
(A3 ₃ ¹ , I1 ₄)	(\mathcal{U}_A , \mathcal{U}_C)	A3 ₃ ¹ → A4 → \mathcal{U}_A	I1 ₄ → I2 → \mathcal{U}_C
(A5 ¹ , I2 ₁)	(A5 ₁ ¹ , I2 ₂)	A5 ¹ → A5 ₁ ¹	I2 ₁ → I2 ₂
(A5 ₁ ¹ , I2 ₂)	(A5 ₂ ¹ , I2 ₃)	A5 ₁ ¹ → A5 ₂ ¹	I2 ₂ → I2 ₃
(A5 ₁ ¹ , I2 ₂)	(\mathcal{W}_A , I2 ₂)	A5 ₁ ¹ → \mathcal{W}_A	ϵ
(A5 ₂ ¹ , I2 ₃)	(A5 ₃ ¹ , I2 ₄)	A5 ₂ ¹ → A5 ₃ ¹	I2 ₃ → I2 ₄
(A5 ₃ ¹ , I2 ₄)	(A10, I5) (A10, \mathcal{U}_C)	A5 ₃ ¹ → A6 → A8 → A10	I2 ₄ → I3 → I5 I2 ₄ → I3 → \mathcal{U}_C
(A10, I5)	(A10, I5) (A10, \mathcal{U}_C)	A10 → A14 → A10	I5 → I6 → I8 → I5 I5 → I6 → \mathcal{U}_C
(A10, I5)	(\mathcal{U}_A , \mathcal{U}_A)	A10 → \mathcal{U}_A	I5 → I6 → \mathcal{U}_A
(A10, I5)	(A17 ¹ , I9) (A17 ¹ , \mathcal{U}_C)	A10 → A14 → A15 → A17 ¹	I5 → I6 → I8 → I5 → I9 I5 → I6 → \mathcal{U}_C
(A10, I5)	(\mathcal{W}_A , I9) (\mathcal{W}_A , \mathcal{U}_C)	A10 → A14 → A15 → \mathcal{W}_A	I5 → I6 → I8 → I5 → I9 I5 → I6 → \mathcal{U}_C
(A18, I11)	(A18 ₁ , I13) (A18 ₁ , \mathcal{U}_C)	A18 → A18 ₁	I11 → I13 I11 → \mathcal{U}_C
(A18 ₁ , I13)	(A18 ₂ , I13 ₁)	A18 ₁ → A18 ₂	I13 → I13 ₁
(A18 ₂ , I13 ₁)	(A18 ₃ , I13 ₂)	A18 ₂ → A18 ₃	I13 ₁ → I13 ₂
(A18 ₃ , I13 ₂)	(A18 ₄ , I13 ₃)	A18 ₃ → A18 ₄	I13 ₂ → I13 ₃
(A18 ₄ , I13 ₃)	(A18 ₁ ¹ , I14)	A18 ₄ → A18 ₁ ¹	I13 ₃ → I14
(A18 ₁ ¹ , I14)	(A18 ₂ ¹ , I14 ₁)	A18 ₁ ¹ → A18 ₂ ¹	I14 → I14 ₁
(A18 ₂ ¹ , I14 ₁)	(A18 ₁ ² , I15 ₁)	A18 ₂ ¹ → A18 ₁ ²	I14 ₁ → I15
(A18 ₁ ² , I15)	(A18 ₂ ² , I15 ₁)	A18 ₁ ² → A18 ₂ ²	I15 → I15 ₁
(A18 ₂ ² , I15 ₁)	(A19 ₁ ¹ , I17) (A19 ₁ ¹ , \mathcal{U}_C)	A18 ₂ ² → A19 ₁ ¹	I15 ₁ → I17 I15 ₁ → \mathcal{U}_C
(A19 ₁ ¹ , I17)	(A19 ₂ ¹ , I17 ₁)	A19 ₁ ¹ → A19 ₂ ¹	I17 → I17 ₁
(A19 ₂ ¹ , I17 ₁)	(A19 ₁ ² , I18)	A19 ₂ ¹ → A19 ₁ ²	I17 ₁ → I18
(A19 ₁ ² , I18)	(A19 ₂ ² , I18 ₁)	A19 ₁ ² → A19 ₂ ²	I18 → I18 ₁
(A19 ₂ ² , I18 ₁)	(A22 ₁ , I19)	A19 ₂ ² → A20 → A22 → A22 ₁	I18 ₁ → I19
(A19 ₂ ² , I18 ₁)	(\mathcal{W}_A , I19)	A19 ₂ ² → A20 → \mathcal{W}_A	I18 ₁ → I19
(A22 ₁ , I19)	(A22 ₂ , I19 ₁)	A22 ₁ → A22 ₂	I19 → I19 ₁
(A22 ₂ , I19 ₁)	(AE, IE)	A22 ₂ → AE	I19 ₁ → IE

Table 4.2: \mathcal{E}_X of X for the procedures shown in fig. 4.5. Each row represents an X edge $e_X = (n_X \xrightarrow{\xi_{\bar{A}}, \xi_C} n_X^t) \in \mathcal{E}_X$.

Chapter 5

SMT Encoding

In the previous chapter, we described our algorithm DYNAMO for simultaneous automatic construction of a product graph X and automatic inference of an annotation \ddot{A} of procedure A . The DYNAMO algorithm generates verification conditions or proof obligations in the form of Hoare triples during its execution. These generated Hoare triples are discharged using off-the-shelf SMT solvers. In this chapter, we describe our encoding for translating a Hoare triple over X into an SMT format compatible with solvers. Our primary contribution here is an efficient SMT representation of an address set and encoding of various relations over an address set.

We organize the chapter as follows. In section 5.1, we discuss some preliminary steps we take before the SMT encoding. In section 5.2, we present an *allocation state array* representation of address sets and encoding of address set relations for this representation. In section 5.3, we describe a faster *interval encoding* for address set relations and present its proof of soundness. We conclude in section 5.4 with description of an alternate semantics for \ddot{A} that are amenable to a simpler SMT encoding.

5.1 Preliminary Steps

At an error-free node $n_X = (n_{\ddot{A}}, n_C) \in \mathcal{N}_X^{\text{bw}}$ of $X = (\mathcal{N}_X, \mathcal{E}_X, \mathcal{D}_X)$, DYNAMO may generate a proof obligation O in the form of a Hoare triple $\{\phi_{n_X}\}(\xi_{\ddot{A}}; [\xi_C]_{\mathcal{D}_X}^{\epsilon_X})\{post\}$ — recall that both path cover (used in (Coverage C)) and path infeasibility (used in (Coverage \ddot{A})) conditions are also represented as Hoare triples with $\xi_C = \epsilon$. A Hoare triple O is encoded as a first-order logic predicate in theories of bitvector, arrays,

Predicate	First-order logic encoding using $\alpha \in \Sigma_P^r$
$\alpha \in \Sigma_P^{\vec{r}}$ $\vec{r} \subseteq R$	$\bigvee_{r \in \vec{r}} \alpha \in \Sigma_P^r$
$\forall_{r \in B} : \Sigma_C^r = \Sigma_{\ddot{A}}^r$	$\forall_{\alpha} : (\alpha \in \Sigma_C^B \Leftrightarrow \alpha \in \Sigma_{\ddot{A}}^B)$
$\Sigma_P^r = \emptyset$	$\forall_{\alpha} : \neg(\alpha \in \Sigma_P^r)$
$(\llbracket \text{lb.z} \rrbracket = \text{lb}(\Sigma_C^z) \wedge \llbracket \text{ub.z} \rrbracket = \text{ub}(\Sigma_C^z))$	$(\forall_{\alpha} : \alpha \in \Sigma_C^z \Rightarrow (\llbracket \text{lb.z} \rrbracket \leq_u \alpha \leq_u \llbracket \text{ub.z} \rrbracket)) \wedge \llbracket \text{lb.z} \rrbracket \in \Sigma_C^z \wedge \llbracket \text{ub.z} \rrbracket \in \Sigma_C^z$
$\text{ov}([\alpha_b, \alpha_e], \Sigma_P^{\vec{r}})$ $\vec{r} \subseteq R$	$\exists_{\alpha} : (\alpha_b \leq_u \alpha \leq_u \alpha_e) \wedge \alpha \in \Sigma_P^{\vec{r}}$
$[\alpha_b, \alpha_e] \subseteq \Sigma_P^{\vec{r}}$ $\vec{r} \subseteq R$	$\forall_{\alpha} : (\alpha_b \leq_u \alpha \leq_u \alpha_e) \Rightarrow \alpha \in \Sigma_P^{\vec{r}}$
$[\alpha_b, \alpha_e] = \Sigma_P^r$	$\forall_{\alpha} : (\alpha_b \leq_u \alpha \leq_u \alpha_e) \Leftrightarrow \alpha \in \Sigma_P^r$
$\Sigma_{\ddot{A}}^{\{\text{stk}\} \cup Y} \cup (\Sigma_{\ddot{A}}^Z \setminus (\Sigma_{\ddot{A}}^{Z_l} ^v)) = [\text{esp}, \text{stk}_e]$	$\forall_{\alpha} : (\alpha \in \Sigma_{\ddot{A}}^{\{\text{stk}\} \cup Y} \vee (\alpha \in \Sigma_{\ddot{A}}^Z \wedge \neg(\alpha \in \Sigma_{\ddot{A}}^{Z_l} ^v))) \Leftrightarrow (\text{esp} \leq_u \alpha \leq_u \text{stk}_e)$
$\Sigma_{\ddot{A}}^{\{\text{cs}, \text{cl}\}} = [\text{stk}_e + 1, \text{cs}_e]$	$\forall_{\alpha} : (\alpha \in \Sigma_{\ddot{A}}^{\{\text{cs}, \text{cl}\}}) \Leftrightarrow (\llbracket \text{stk}_e \rrbracket + 1 \leq_u \alpha \leq_u \llbracket \text{cs}_e \rrbracket)$

Table 5.1: Encodings of address set predicates using the address set membership predicate. R is the set of all region identifiers.

and uninterpreted functions over the state elements at n_X and discharged using an off-the-shelf SMT solver. The state elements of a procedure $P \in \{C, \ddot{A}\}$ are represented using bitvectors for a register/variable/ghost variable, arrays for memory M_P , and uninterpreted functions for $\text{read}_{\vec{r}}(\Omega_P)$ and $\text{io}(\Omega_P, \text{rw}, \vec{v})^1$.

For address sets, we describe encoding for the set-membership predicate $\alpha \in \Sigma_P^r$ for an arbitrary address α and region identifier r . All other address set predicates can be expressed in terms of the set-membership predicate; see table 5.1 for a construction. To simplify the encodings, we rely on the correct-by-construction global invariants $\llbracket \phi_X \rrbracket \subseteq \phi_{n_X}$ (section 4.2.1) and assume that the partially-constructed X (and, therefore, node n_X) satisfies the (Equivalence), (MAC), and (MemEq) requirements. The (Equivalence) requirement implies the predicate $\llbracket \text{AllocEq} \rrbracket$ and (MemEq) implies the predicate $\llbracket \text{MemEq} \rrbracket$ (both defined in section 4.2) so that $\llbracket \text{AllocEq} \rrbracket \in \phi_{n_X}$ and $\llbracket \text{MemEq} \rrbracket \in \phi_{n_X}$.

Sets of stack-allocated and virtually-allocated locals at a node

Recall that for an allocation site $z_l \in Z_l$ due to a variable declaration (or procedure parameter), at a node $n_X \in \mathcal{N}_X$, we use $\Sigma_{\ddot{A}}^{z_l |^s}$ and $\Sigma_{\ddot{A}}^{z_l |^v}$ to represent the address sets corresponding to the stack and virtual allocations performed in \ddot{A} for z_l (section 2.6).

¹Recall that read and io are the uninterpreted functions used for representing interaction with the outside world Ω_P (section 2.2.7).

Let $Zls = \{zl \mid zl \in Z_l \wedge \Sigma_{\ddot{A}}^{zl|s} \neq \emptyset\}$ and $Zlv = \{zl \mid zl \in Z_l \wedge \Sigma_{\ddot{A}}^{zl|v} \neq \emptyset\}$ represent the set of non-empty stack-allocated and virtually-allocated locals at n_X respectively. Recall that we restrict ourselves to only those compiler transformations that ensure the validity of $Zls \cap Zlv = \emptyset$ at each n_X (section 2.6). This enables us to work with Zls and Zlv instead of Z_l , which simplifies the encoding of $\alpha \in \Sigma_{\ddot{A}}^{zl}$ to either $\alpha \in \Sigma_{\ddot{A}}^{zl|s}$ (for $zl \in Zls$) or $\alpha \in \Sigma_{\ddot{A}}^{zl|v}$ (for $zl \in Zlv$) or **false** (if $zl \notin Zls \cup Zlv$). Further, because relations $\Sigma_{\ddot{A}}^{zl|s} \{=, \neq\} \emptyset$ and $\Sigma_{\ddot{A}}^{zl|v} \{=, \neq\} \emptyset$ are tracked through invariant shape `zEmpty` (fig. 4.2), Zls and Zlv can be identified through syntactic pattern matching over ϕ_{n_X} .

Using invariants for simplifying encoding

We use `AllocEq` (assumed to be in ϕ_{n_X} due to (Equivalence)) to replace $\alpha \in \Sigma_{\ddot{A}}^r$ with $\alpha \in \Sigma_C^r$ for $r \in B$. For example, the `MemEq` predicate $M_C =_{\Sigma_{\ddot{A}}^B \setminus (\Sigma_{\ddot{A}}^{Zl|v})} M_{\ddot{A}}$ is encoded as:

$$\forall \alpha : \alpha \in \Sigma_C^{GUYUZlsUZa \cup \{hp, cl\}} \Rightarrow \text{sel}_1(M_C, \alpha) = \text{sel}_1(M_{\ddot{A}}, \alpha)$$

5.2 Representing address sets using allocation state array

Let $\mathcal{L}_P : i_{32} \rightarrow R$ be an *allocation state array* that maps an address to a region identifier in procedure $P \in \{\ddot{A}, C\}$. We add a state element \mathcal{L}_P to procedure P for tracking the allocation state, address sets Σ_P^r for all $r \in R \setminus \{\text{free}\}$, in procedure P . The set-membership predicate $\alpha \in \Sigma_P^r$ for $r \in R$, which can be used for encoding all other predicates over address set Σ_P^r , (table 5.1), is encoded using \mathcal{L}_P as follows.

- For $r \notin Zlv$, $\alpha \in \Sigma_P^r$ is encoded as $\text{sel}_1(\mathcal{L}_P, \alpha) = r$.
- For $r \in Zlv$, both $\alpha \in \Sigma_C^r$ and $\alpha \in \Sigma_{\ddot{A}}^r$ are encoded as $\text{sel}_1(\mathcal{L}_C, \alpha) = r$, i.e., the encodings for both \ddot{A} and C use \mathcal{L}_C for virtually-allocated locals (by relying on the `AllocEq` invariant that is guaranteed to hold at n_X).

Thus, $\mathcal{L}_{\ddot{A}}$ is not used to track the virtually-allocated locals; instead, an address belonging to a virtually allocated-region maps to one of $\{\text{free}, \text{stk}, \text{cs}\} \cup F$ regions in $\mathcal{L}_{\ddot{A}}$.

An array-based encoding that maps a unique region to an address is possible in C because of global invariant `NoOverlapC` that forbids any overlap among all regions in

#	Instruction	SMT Encoding using \mathcal{L}_P
1	$\Sigma_P^r := \Sigma_P^r \cup [\alpha_b, \alpha_e]; \quad r \in \{stk\} \cup Z$	$\mathcal{L}_{P'} = \text{cwrite}(\mathcal{L}_P, \lambda x. x \in [\alpha_b, \alpha_e], r)$
2	$\Sigma_{\ddot{A}}^{stk} := \Sigma_{\ddot{A}}^{stk} \setminus [\alpha_b, \alpha_e];$	$\mathcal{L}_{\ddot{A}'} = \text{cwrite}(\mathcal{L}_{\ddot{A}}, \lambda x. x \in [\alpha_b, \alpha_e], \text{free})$
3	$\Sigma_P^z := \emptyset;$	$\mathcal{L}_{P'} = \text{cwrite}(\mathcal{L}_P, \lambda x. \text{sel}_1(\mathcal{L}_P, x) = z, \text{free})$
4	$\Sigma_{\ddot{A}}^{stk} := \{[\text{esp}, \boxed{\text{stk}_e}]\} \setminus \Sigma_{\ddot{A}}^Y;$	$\mathcal{L}_{\ddot{A}'} = \text{cwrite}(\mathcal{L}_{\ddot{A}}, \lambda x. x \in [\text{esp}, \boxed{\text{stk}_e}] \wedge \bigwedge_{y \in Y} x \notin \Sigma_{\ddot{A}}^y, \text{stk})$

Table 5.2: SMT encoding of address set updating instructions using allocation state array \mathcal{L}_P . $P \in \{\mathbf{C}, \ddot{A}\}$. $\mathcal{L}_{P'}$ is the allocation state array obtained after executing the instruction.

\mathbf{C} at an error-free node $n_{\mathbf{C}} \in \mathcal{N}_{\mathbf{C}}^{\text{err-free}}$. In \ddot{A} , $\boxed{\text{NoOverlap}\ddot{A}}$ permits overlap of $\Sigma_{\ddot{A}}^{Zlv}$ with $\Sigma_{\ddot{A}}^{\{stk, cs\} \cup F}$ and, consequently, $\mathcal{L}_{\ddot{A}}$ is only used for tracking regions other than Zlv .

As \mathcal{L}_P is an array state element, similar to M_P , it is directly encode-able in SMT using the theory of arrays.

5.2.1 Encoding of address set updating instructions

An address set updating instruction of the form ‘ $\Sigma_P^r := e(\dots)$ ’ updates the address set of region $r \in R$. In an allocation state array \mathcal{L}_P representation, an update ‘ $\Sigma_P^r := e(\dots)$ ’ produces a new allocation state array $\mathcal{L}_{P'}$. Table 5.2 shows \mathcal{L}_P based SMT encoding of the graph instructions that update address sets. We list the encoding for the four kinds of address set updating instructions that appear in our translations (figs. 2.5 to 2.8, 2.10 and 2.11).

Table 5.2 uses an auxiliary “conditional write” operator cwrite to encode the update of \mathcal{L}_P . If $\mathcal{L}_{P'} = \text{cwrite}(\mathcal{L}_P, \lambda x. c, v)$, then the following holds:

$$\begin{aligned} \forall \alpha : \quad & (\lambda x. c)(\alpha) \Rightarrow \text{sel}_1(\mathcal{L}_{P'}, \alpha) = v \\ & \wedge \neg(\lambda x. c)(\alpha) \Rightarrow \text{sel}_1(\mathcal{L}_{P'}, \alpha) = \text{sel}_1(\mathcal{L}_P, \alpha) \end{aligned}$$

Here, $(\lambda x. c)$ represents a function that takes an address x and returns a boolean evaluated through expression c , and $(\lambda x. c)(\alpha)$ represents the application of this function to address α . Thus, $\text{cwrite}(\mathcal{L}_P, \lambda x. c, v)$ represents the modification of allocation state array \mathcal{L}_P to value v for all addresses α that satisfy the boolean condition c . In other words, $\text{cwrite}(\mathcal{L}_P, \lambda x. c, v)$ is equivalent to $\text{st}_1(\dots \text{st}_1(\dots \text{st}_1(\mathcal{L}_P, \alpha_1, v), \dots, \alpha_i, v), \dots, \alpha_n, v)$

for all $\alpha_1, \dots, \alpha_i, \dots, \alpha_n$ where the predicate c holds ².

We discuss the specific cases of allocation and deallocation below:

- Allocation of an interval $[\alpha_b, \alpha_e]$ to region r through instruction ‘ $\Sigma_p^r := \Sigma_p^r \cup [\alpha_b, \alpha_e]$ ’ is encoded as $\mathcal{L}_{P'} = \text{cwrite}(\mathcal{L}_P, \lambda x. x \in [\alpha_b, \alpha_e], r)$ which translates to “mark the addresses in interval $[\alpha_b, \alpha_e]$ as belonging to region r in $\mathcal{L}_{P'}$ ” (shown in row 1 of table 5.2 with $r \in \{stk\} \cup Z$).
- Similarly, deallocation of an interval $[\alpha_b, \alpha_e]$ from region r through instruction ‘ $\Sigma_p^r := \Sigma_p^r \setminus [\alpha_b, \alpha_e]$ ’ is encoded as $\mathcal{L}_{P'} = \text{cwrite}(\mathcal{L}_P, \lambda x. x \in [\alpha_b, \alpha_e], \text{free})$ which translates to “mark the addresses in interval $[\alpha_b, \alpha_e]$ as belonging to region **free** in $\mathcal{L}_{P'}$ ” (shown in row 2 of table 5.2 with $r = stk$).

The deallocation of a region r through instruction ‘ $\Sigma_p^r := \emptyset$ ’ is encoded as $\mathcal{L}_{P'} = \text{cwrite}(\mathcal{L}_P, \lambda x. \text{sel}_1(\mathcal{L}_P, x) = r, \text{free})$ where the addresses to be set to **free** are identified using the predicate $\text{sel}_1(\mathcal{L}_P, x)$ (instead of a range check $x \in [\alpha_b, \alpha_e]$ as done in case of deallocation of an interval from r).

As $\mathcal{L}_{\check{A}}$ does not track virtually-allocated locals, the (de)allocation instructions $\Sigma_{\check{A}}^{zlv}|^v := \Sigma_{\check{A}}^{zlv}|^v \cup [v]_w$ and $\Sigma_{\check{A}}^{zlv}|^v := \emptyset$ ((**ALLOCV**) and (**DEALLOCV**) in fig. 2.10) for $zlv \in Zlv$ become vacuous in \check{A} , i.e., they do not change any state element in \check{A} .

5.2.2 Full-array encoding

We call this allocation state arrays \mathcal{L}_C and $\mathcal{L}_{\check{A}}$ based representation of address sets and corresponding encoding of address set relations a *full-array encoding*. In a full-array encoding, the address sets of **C** are tracked using \mathcal{L}_C and the address set of \check{A} are tracked using a combination of \mathcal{L}_C (for a region $r \in Zlv$) and $\mathcal{L}_{\check{A}}$ (otherwise). A predicate $\alpha \in \Sigma_p^r$ for a region r is encoded as an SMT *array select* operation $\text{select}(\mathcal{L}_P, \alpha)$ over the respective array \mathcal{L}_P .

A proof obligation encoded using full-array encoding contains constructs with quantifiers over SMT arrays (tables 5.1 and 5.2). Such a proof obligation, with quantifiers over arrays, can be (relatively) slow to discharge using SMT solvers. In the subsequent sections, we describe an *interval encoding* that makes use of the global invariants for a more performant encoding attainable under certain conditions. We confirm the relative better performance of our interval encoding in our experiments (section 6.2).

²Recall that $\text{st}_1(\mathcal{L}_P, \alpha_1, v)$ is our size-associated *store* operation which is equivalent to SMT expression $\text{store}(\mathcal{L}_P, \alpha_1, v)$ in this particular case.

5.3 Interval Encoding

5.3.1 Interval encoding for $r \in G \cup F \cup Y \cup Z_l \cup \{stk\}$

Recall that the address set Σ_p^r of a region $r \in G \cup F \cup Y \setminus \{\text{vrdc}\}$ is an interval and the address set Σ_p^r of a region $r \in Z_l \cup \{\text{vrdc}\}$ is either empty or an interval (section 4.2.1) — the global invariants $\boxed{\text{gfyIntvl}}$ and $\boxed{\text{zlIntvl}}$ encode this in ϕ_{n_X} . We use these invariants (and $\boxed{\text{AllocEq}}$) for a more performant *interval encoding*. In the interval encoding, we encode $\alpha \in \Sigma_p^r$ for $r \in G \cup F \cup Y \cup Z_l$ as³

$$\neg \text{em}.r \wedge (\text{lb}.r \leq_u \alpha \leq_u \text{ub}.r)$$

Moreover, if there are no local variables allocated due to the `alloca()` operator (i.e., $\Sigma_p^{Z_a} = \emptyset$), then all local variables are contiguous and thus amenable for an interval encoding. In this case, due to $\boxed{\text{StkBd}}$, the `stk` region can be identified as $[\text{esp}, \text{stk}_e] \setminus \Sigma_{\bar{A}}^{Y \cup Z_l}$.

As the interval encoding utilizes ghost variables ($\text{em}.z$, $\text{lb}.z$, $\text{ub}.z$), that are updated during an (de)allocation, the SMT encoding of an address set mutating instruction becomes vacuous.

5.3.2 Interval encoding for $r \in \{hp, cl, cs\}$

Even though the regions `hp` (heap), `cl` (callers' locals), and `cl` (callers' stack) can be discontinuous in general, we over-approximate these regions to their contiguous covers to be able to soundly encode them using intervals. Recall that our proof obligation O is a Hoare triple of the form $\{\phi_{n_X}\}(\xi_{\bar{A}}; [\xi_C]_{\mathcal{D}_X}^{ex})\{post\}$ such that $n_X = (n_{\bar{A}}, n_C) \in \mathcal{N}_X^{UW}$. Here, if $\xi_{\bar{A}}$ is an I/O path, its execution interacts with the outside world, and so an over-approximation of an externally-visible address set is unsound. We thus restrict ourselves to an I/O-free $\xi_{\bar{A}}$ (consequently, I/O-free ξ_C due to (SingleIO)) for our interval encoding.

Let $n_{\bar{A}}^1, n_{\bar{A}}^2, \dots, n_{\bar{A}}^m$ be the nodes on path $\xi_{\bar{A}} = (n_{\bar{A}} \rightarrow n_{\bar{A}}^t)$, such that $n_{\bar{A}}^1 = n_{\bar{A}}$ and $n_{\bar{A}}^m = n_{\bar{A}}^t$. Let $SP_{min}(\xi_{\bar{A}})$ represent the minimum value of stackpointer `esp` observed at any node $n_{\bar{A}}^j$ ($1 \leq j \leq m$) visited during the execution of path $\xi_{\bar{A}}$. Similarly, let $Zlv_U(\xi_{\bar{A}})$ be the union of the values of address set $\Sigma_{\bar{A}}^{Zlv}$ observed at any $n_{\bar{A}}^j$ ($1 \leq j \leq m$) visited during

³Recall that $\text{em}.r$, $\text{lb}.r$, and $\text{ub}.r$ are ghost variables that capture the emptiness of region r , lower bound of region r , and upper bound of region r respectively.

$\xi_{\check{A}}$'s execution. Intuitively, $SP_{min}(\xi_{\check{A}})$ is the lowest possible value of \mathbf{esp} ⁴ and $Zlv_U(\xi_{\check{A}})$ is the largest possible value of set $\Sigma_{\check{A}}^{Zlv}$ seen during $\xi_{\check{A}}$'s execution starting at n_X .

Let $HP(\xi_{\check{A}})$, $CL(\xi_{\check{A}})$, and $CS(\xi_{\check{A}})$ be defined as shown below.

$$\begin{aligned} HP(\xi_{\check{A}}) &= \text{comp}(\Sigma_{\check{A}}^{GUF} \cup Zlv_U(\xi_{\check{A}}) \cup [SP_{min}(\xi_{\check{A}}), \boxed{\mathbf{cs}_e}]) \\ CL(\xi_{\check{A}}) &= [\boxed{\mathbf{stk}_e} + 1_{i_{32}}, \boxed{\mathbf{cs}_e}] \setminus Zlv_U(\xi_{\check{A}}) \\ CS(\xi_{\check{A}}) &= [\boxed{\mathbf{stk}_e} + 1_{i_{32}}, \boxed{\mathbf{cs}_e}] \cap Zlv_U(\xi_{\check{A}}) \end{aligned}$$

Intuitively, $HP(\xi_{\check{A}})$ and $CL(\xi_{\check{A}})$ are the largest possible values for $\Sigma_{\check{A}}^{hp}$ and $\Sigma_{\check{A}}^{cl}$ respectively such that ϕ_{n_X} is satisfied and $\xi_{\check{A}}$ is executed to completion; as cs and cl are complementary (due to $\boxed{\mathbf{csBd}}$), this makes $CS(\xi_{\check{A}})$ the smallest possible value under the same conditions.

Theorem 5.3.1. *Let $O = \{pre\}(\xi_{\check{A}}; [\xi_C]_{D_X}^{ex})\{post\}$ be a proof obligation generated by DYNAMO. Let O' be obtained from O by strengthening precondition pre to $pre' = pre \wedge (\Sigma_{\check{A}}^{hp} = HP(\xi_{\check{A}})) \wedge (\Sigma_{\check{A}}^{cl} = CL(\xi_{\check{A}})) \wedge (\Sigma_{\check{A}}^{cs} = CS(\xi_{\check{A}}))$. If $\xi_{\check{A}}$ is I/O-free, $O \Leftrightarrow O'$ holds.*

Proof sketch. $O \Rightarrow O'$ is trivial. The proof for $O' \Rightarrow O$ relies on the limited shapes of predicates that may appear in pre and $post$, the enumeration of \mathcal{U} -maximal pathsets for C (section 4.1.2), and use of safety-relaxed semantics for \check{A} (section 3.5). For I/O-free $\xi_{\check{A}}$, pre and $post$ shapes are limited by our invariant grammar (fig. 4.2), and the edge conditions appearing in our execution semantics (figs. 2.4 to 2.8, 2.10 and 2.11). The full proof is available in section 5.3.3. \square

Using theorem 5.3.1, we rewrite $\alpha \in \Sigma_p^{hp}$ to $\alpha \in HP(\xi_{\check{A}})$, $\alpha \in \Sigma_p^{cl}$ to $\alpha \in CL(\xi_{\check{A}})$, and $\alpha \in \Sigma_p^{cs}$ to $\alpha \in CS(\xi_{\check{A}})$ in proof obligation O . If $\Sigma_p^{Za} = \emptyset$ holds at n_X , we encode all non-free regions using intervals — we call this a *full-interval encoding*; else, we encode regions in $Y \cup Z_a \cup Zls \cup \{stk\}$ using an allocation state array, and $GUF \cup Zlv \cup \{hp, cl, cs\}$ using intervals — we call this a *partial-interval encoding*.

Table 5.3 shows the SMT encoding of $\alpha \in \Sigma_p^r$ in *full-array*, *partial-interval*, and *full-interval* encoding. The column selects the encoding and the row selects the region. We group regions with common encoding into a single row, e.g., the third row gives the encoding for a region $r \in G \cup Zlv$. We merge cells with common entries for

⁴Recall that stack is allocated by decrementing \mathbf{esp} .

$\alpha \in \Sigma_P^r$	Full-array encoding		Partial-interval encoding	Full-interval encoding		
	$P = C$	$P = A$	$(\Sigma_P^{Z_a} \neq \emptyset)$	$(\Sigma_P^{Z_a} = \emptyset)$		
$r = hp$	$\text{sel}_1(\mathcal{L}_C, \alpha) = r$		$\alpha \notin (\Sigma_{\check{A}}^{G \cup F} \cup \text{Zlv}_U(\xi_{\check{A}}) \cup [SP_{\min}(\xi_{\check{A}}), \text{cs}_e])$			
$r = cl$			$\alpha \in [\text{stk}_e + 1, \text{cs}_e] \wedge \alpha \notin \text{Zlv}_U(\xi_{\check{A}})$			
$r \in G \cup \text{Zlv}$			$\neg \text{em}.r \wedge (\text{lb}.r \leq_u \alpha \leq_u \text{ub}.r)$			
$r \in Y \cup \text{Z}_a \cup \text{Zls}$						
$r \in F$	false	$\text{sel}_1(\mathcal{L}_{\check{A}}, \alpha) = r$		$\alpha \in [\text{esp}, \text{stk}_e] \wedge \bigwedge_{r \in Y \cup \text{Zls}} (\alpha \notin \Sigma_{\check{A}}^r)$		
$r = cs$					$\alpha \in [\text{stk}_e + 1, \text{cs}_e] \wedge \alpha \in \text{Zlv}_U(\xi_{\check{A}})$	
$r = stk$						

Table 5.3: SMT encoding of $\alpha \in \Sigma_P^r$ for DYNAMO's proof obligation O with outgoing assembly path $\xi_{\check{A}}$.

clearer presentation, e.g., instead of repeating $\text{sel}_1(\mathcal{L}_C, \alpha) = r$ for full-array encoding of $r \in \{hp, cl\} \cup G \cup \text{Zlv} \cup Y \cup \text{Z}_a \cup \text{Zls}$, we merge the cells for each of these entries into a single cell. As $\Sigma_C^r = \emptyset$ for $r \in F \cup \{cs, stk\}$, $\alpha \in \Sigma_C^r$ is encoded as **false**.

5.3.3 Soundness of Interval Encoding

Let the Hoare triple representation of a proof obligation O generated by DYNAMO be $\{pre\}(\xi_{\check{A}}; [\xi_C]_{\mathcal{D}_X}^{e_X})\{post\}$, where $\xi_{\check{A}} = (n_{\check{A}} \rightarrow n_{\check{A}}^t)$ and either $\xi_C = \epsilon$ or $\xi_C = (n_C \rightarrow n_C^t)$; both $\xi_{\check{A}}$ and ξ_C are I/O-free execution paths in \check{A} and C respectively; $n_X = (n_{\check{A}}, n_C) \in \mathcal{N}_X^{\text{DW}}$ is an error-free node; if $\xi_C = (n_C \rightarrow n_C^t)$, then $e_X = (n_X \xrightarrow{\xi_{\check{A}}; \xi_C} n_X^t) \in \mathcal{E}_X$ and $n_X^t = (n_{\check{A}}^t, n_C^t) \in \mathcal{N}_X$.

Let $HP(\xi_{\check{A}})$, $CL(\xi_{\check{A}})$, and $CS(\xi_{\check{A}})$ be defined as described in previous section.

Let $O' = \{pre\}(\xi_{\check{A}}; \xi_C)\{post\}$ be obtained by strengthening precondition pre to $pre' = pre \wedge (\Sigma_{\check{A}}^{hp} = HP(\xi_{\check{A}})) \wedge (\Sigma_{\check{A}}^{cl} = CL(\xi_{\check{A}})) \wedge (\Sigma_{\check{A}}^{cs} = CS(\xi_{\check{A}}))$ in O' . We need to show that $O \Leftrightarrow O'$ holds.

(\Rightarrow) Proving $O \Rightarrow O'$ is trivial, as O' requires a stronger precondition than O (with everything else identical).

(\Leftarrow) Assume that O' holds. We are interested in showing that O holds. Assume a machine state σ of product program X that satisfies the weaker precondition pre , and

executes to completion over $\xi_{\ddot{A}}$ and $\xi_{\ddot{C}}$. We are interested in showing that σ satisfies the postcondition *post* after completing the execution.

We define “error-free execution” to be the case where the execution on a state σ across $(\xi_{\ddot{A}}; \xi_{\ddot{C}})$ does not end at an error node in \mathbb{X} .

Lemma 5.3.2 ($HP(\xi_{\ddot{A}}), CL(\xi_{\ddot{A}})$ overapproximate hp, cl). $(\Sigma_{\ddot{A}}^{hp} \subseteq HP(\xi_{\ddot{A}})) \wedge (\Sigma_{\ddot{A}}^{cl} \subseteq CL(\xi_{\ddot{A}}))$ holds on σ for an error-free execution.

Proof. Recall that $pre \Rightarrow \phi_{n_x}$. If $\Sigma_{\ddot{A}}^{hp} \supset HP(\xi_{\ddot{A}})$ or $\Sigma_{\ddot{A}}^{cl} \supset CL(\xi_{\ddot{A}})$, then either at least one of `NoOverlapA` or `NoOverlapC` will evaluate to `false` in ϕ_{n_x} (and *pre*), or during the execution of path $\xi_{\ddot{A}}$; error \mathcal{W} will be triggered in \ddot{A} because either the allocation of stack space through stackpointer decrement will overstep $\Sigma_{\ddot{A}}^{\{hp, cl\}}$ (OP-ESP’), or the virtual allocation of a local variable will overstep $\Sigma_{\ddot{A}}^{\{hp, cl\}}$ (ALLOCV). However, by assumption, σ satisfies *pre* (and ϕ_{n_x}) and executes $\xi_{\ddot{A}}$ and $\xi_{\ddot{C}}$ to completion to an error-free node; thus proved by contradiction. \square

Lemma 5.3.3 ($CS(\xi_{\ddot{A}})$ underapproximates *cs*). $(\Sigma_{\ddot{A}}^{cs} \supseteq CS(\xi_{\ddot{A}}))$ holds on σ for an error-free execution.

Proof. Follows from lemma 5.3.2 and `csBd` ($\Sigma_{\ddot{A}}^{cs} = [\text{stk}_e + 1_{i_{32}}, \text{cs}_e] \setminus \Sigma_{\ddot{A}}^{cl}$). \square

Lemma 5.3.4 ($HP(\xi_{\ddot{A}})$ and $CL(\xi_{\ddot{A}})$ borrow from the `free` and `cs` regions). The following hold on σ for an error-free execution.

1. $(HP(\xi_{\ddot{A}}) \setminus \Sigma_{\ddot{A}}^{hp}) \subseteq \Sigma_{\ddot{A}}^{\text{free}} \subseteq \Sigma_{\ddot{C}}^{\text{free}}$
2. $(CL(\xi_{\ddot{A}}) \setminus \Sigma_{\ddot{A}}^{cl}) \subseteq \Sigma_{\ddot{A}}^{cs} \subseteq \Sigma_{\ddot{C}}^{\text{free}}$

Proof. The proof follows from the definition of $HP(\xi_{\ddot{A}})$ and $CL(\xi_{\ddot{A}})$, as these sets are not allowed to overlap with $\Sigma_{\ddot{A}}^{BUFUS}$ or $\Sigma_{\ddot{C}}^{BUFUS}$. \square

Construct a state σ' that is identical to σ with the following modifications made in sequence:

1. The region identified by addresses (that would belong to region `free` in \ddot{C} by lemma 5.3.4) $(HP(\xi_{\ddot{A}}) \cup CL(\xi_{\ddot{A}})) \setminus \Sigma_{\ddot{A}}^{\{hp, cl\}}$ in σ' 's $M_{\ddot{C}}$ is updated through $M_{\ddot{C}} := \text{upd}_{(HP(\xi_{\ddot{A}}) \cup CL(\xi_{\ddot{A}})) \setminus \Sigma_{\ddot{A}}^{\{hp, cl\}}}(M_{\ddot{C}}, M_{\ddot{A}})$.

2. The address sets $\Sigma_{\ddot{A}}^{hp}$, $\Sigma_{\ddot{A}}^{cl}$, $\Sigma_{\mathbb{C}}^{hp}$, and $\Sigma_{\mathbb{C}}^{cl}$ are expanded and the address set $\Sigma_{\ddot{A}}^{cs}$ is shrunk so that $\Sigma_{\ddot{A}}^{hp} = \Sigma_{\mathbb{C}}^{hp} = HP(\xi_{\ddot{A}})$, $\Sigma_{\ddot{A}}^{cl} = \Sigma_{\mathbb{C}}^{cl} = CL(\xi_{\ddot{A}})$, and $\Sigma_{\ddot{A}}^{cs} = CS(\xi_{\ddot{A}})$ (this involves the transfer of addresses from the **free** region to **hp** and **cl** regions in \mathbb{C} , and from the **free** and **cs** regions to **hp** and **cl** regions respectively in \ddot{A} (lemma 5.3.4)).

The constructed state σ' thus satisfies the stronger precondition pre' .

Let Σ_{σ}^{hp} ($\Sigma_{\sigma'}^{hp}$), Σ_{σ}^{cl} ($\Sigma_{\sigma'}^{cl}$), Σ_{σ}^{cs} ($\Sigma_{\sigma'}^{cs}$), and $\Sigma_{\sigma}^{\text{free}}$ ($\Sigma_{\sigma'}^{\text{free}}$) denote the values of $\Sigma_{\ddot{A}}^{hp}$, $\Sigma_{\ddot{A}}^{cl}$, $\Sigma_{\ddot{A}}^{cs}$, and $\Sigma_{\ddot{A}}^{\text{free}}$ in state σ (σ') respectively. Similarly, let $M_{\ddot{A}}^{\sigma}$ ($M_{\mathbb{C}}^{\sigma}$) and $M_{\ddot{A}}^{\sigma'}$ ($M_{\mathbb{C}}^{\sigma'}$) represent the state of procedure \ddot{A} 's (\mathbb{C} 's) memory $M_{\ddot{A}}$ ($M_{\mathbb{C}}$) in machine states σ and σ' respectively.

To relate σ and σ' , we define relation $sim(\sigma, \sigma')$ as the conjunction of the following conditions:

1. (**hp** subset in σ) $\Sigma_{\sigma}^{hp} \subseteq \Sigma_{\sigma'}^{hp}$.
2. (**cl** subset in σ) $\Sigma_{\sigma}^{cl} \subseteq \Sigma_{\sigma'}^{cl}$.
3. (**cs** superset in σ) $\Sigma_{\sigma}^{cs} \supseteq \Sigma_{\sigma'}^{cs}$.
4. (**free** superset in σ) $\Sigma_{\sigma}^{\text{free}} \supseteq \Sigma_{\sigma'}^{\text{free}}$.
5. (\ddot{A} 's memory states are equal) $M_{\ddot{A}}^{\sigma} = M_{\ddot{A}}^{\sigma'}$.
6. (\mathbb{C} 's memory states are equal except at the updated regions) $M_{\mathbb{C}}^{\sigma} =_{\text{comp}(\Sigma_{\sigma'}^{\{hp,cl\}} \setminus \Sigma_{\sigma}^{\{hp,cl\}})} M_{\mathbb{C}}^{\sigma'}$.
7. The remaining state elements have equal values in σ and σ' .

By construction, $sim(\sigma, \sigma')$ holds.

Lemma 5.3.5 ($sim(\sigma, \sigma')$ is preserved for error-free execution across all non-I/O edges in $\mathcal{E}_{\ddot{A}}$). *If a non-I/O edge $e_{\ddot{A}} \in \mathcal{E}_{\ddot{A}}$ is executed on both machine states σ and σ' , and if $sim(\sigma, \sigma')$ holds before the execution, and if the execution on σ completes without error, then there exists a sequence of non-deterministic choices during the execution on σ' such that the execution is error-free and $sim(\sigma, \sigma')$ holds at the end of both error-free executions.*

Proof. For each non-I/O \ddot{A} instruction that does not refer to the $\{hp, cl, cs, \text{free}\}$ regions ((OP-NESP), (ALLOCS), (DEALLOCS), (CALL $_{\ddot{A}}$), (RET $_{\ddot{A}}$), (DEALLOCV)), the execution will have identical behaviour on both σ and σ' , as identical values will be observed in σ and σ' . Thus, if an execution on σ' makes the same non-deterministic choice as the execution on σ , the execution on σ' will complete without error and

$sim(\sigma, \sigma')$ will hold at the end of both executions.

We consider each remaining non-I/O instruction in \ddot{A} below:

- (ENTRY \ddot{A}). Consider the overlap conditions $\Upsilon_1 = ov(\Sigma_{\ddot{A}}^{hp}, \Sigma_{\ddot{A}}^{cl}, \dots, i_{\ddot{A}}^g, \dots, \Sigma_{\ddot{A}}^f, \dots, i_{\ddot{A}}^y, \dots, \Sigma_{\ddot{A}}^{vrdc})$ (due to $\neg addrSetsAreWF$), $\Upsilon_2 = ov([\text{esp}, \text{esp} + 3_{i_{32}}], \Sigma_{\ddot{A}}^{BUF})$, $\Upsilon_3 = ov([\text{stk}_e + 1_{i_{32}}, \text{cs}_e], \Sigma_{\ddot{A}}^{\{hp\} \cup G \cup F})$, and $\Upsilon_4 = ov(\Sigma_{\ddot{A}}^{cl}, \text{comp}([\text{stk}_e + 1_{i_{32}}, \text{cs}_e]))$ (due to $stkIsWF$). During an execution on σ , all four conditions must evaluate to **false**, as we assume an error-free execution on σ . For the same non-deterministic choices made in both executions (over σ and σ'), by the definitions of $HP(\xi_{\ddot{A}})$ and $CL(\xi_{\ddot{A}})$, Υ_1 , Υ_2 , Υ_3 , and Υ_4 will also evaluate to **false** for an execution on σ' — recall that $HP(\xi_{\ddot{A}})$ cannot overlap with $[\text{esp}, \text{cs}_e]$ (which includes the arguments) and global variable regions (due to lemma 5.3.4); and $CL(\xi_{\ddot{A}})$ is a subset of $[\text{stk}_e + 1_{i_{32}}, \text{cs}_e]$ (by definition). Further, because all other state elements observed during the execution of the non-I/O edges in (ENTRY \ddot{A}) are identical in both σ and σ' , $sim(\sigma, \sigma')$ will hold at the end of error-free executions.
- (OP-ESP). The negated subset check $\Upsilon = \neg([t, \text{esp} - 1_{i_{32}}] \subseteq \Sigma_{\ddot{A}}^{\text{free}} \cup \Sigma_{\ddot{A}}^{Zl|v})$ (due to $\neg intrvlInSet(t, \text{esp} - 1_{i_{32}}, \Sigma_{\ddot{A}}^{\text{free}} \cup \Sigma_{\ddot{A}}^{Zl|v})$) depends (indirectly) on the addresses of the set $\Sigma_{\ddot{A}}^{\{hp, cl\}}$ (as **free** is defined as complement of the allocated region). The execution on σ must evaluate Υ to **false** as we assume an error-free execution. By the definitions of $HP(\xi_{\ddot{A}})$ and $CL(\xi_{\ddot{A}})$, for the same non-deterministic choices made in both executions (over σ and σ'), Υ will also evaluate to **false** for an execution on σ' — recall that $(HP(\xi_{\ddot{A}}) \cup CL(\xi_{\ddot{A}}))$ cannot overlap with $[SP_{min}(\xi_{\ddot{A}}), \text{stk}_e]$, and the latter includes $[t, \text{esp} - 1_{i_{32}}]$. All other state elements observed in the other instructions of (OP-ESP) are identical in both σ , σ' and $sim(\sigma, \sigma')$ will hold at the end of error-free executions.
- (ALLOCV). Consider the negated subset check $\Upsilon = \neg([v]_w \subseteq \Sigma_{\ddot{A}}^{\text{comp}(B)})$ (due to $\neg intrvlInSet_a(v, v + w - 1_{i_{32}}, \Sigma_{\ddot{A}}^{\text{comp}(B)})$). The execution on σ must evaluate Υ to **false** as we assume an error-free execution. By the definitions of $HP(\xi_{\ddot{A}})$ and $CL(\xi_{\ddot{A}})$, for the same non-deterministic choices made in both executions (over σ and σ'), Υ will also evaluate to **false** for an execution on σ' — recall that $(HP(\xi_{\ddot{A}}) \cup CL(\xi_{\ddot{A}}))$ cannot overlap with $Zlv_U(\xi_{\ddot{A}})$, and the latter includes the interval $[v]_w$. All other state elements observed in the other instructions of (ALLOCV) are identical in both σ , σ' and $sim(\sigma, \sigma')$ will hold at the end of error-free executions.
- (LOAD \ddot{A}) and (STORE \ddot{A}). The overlap checks, $ov([p]_w, (\Sigma_{\ddot{A}}^{Zl|v}) \setminus (\Sigma_{\ddot{A}}^F \cup [\text{esp}, \text{cs}_e]))$

for $(\text{LOAD}_{\bar{A}})$ and $\text{ov}([p]_w, (\Sigma_{\bar{A}}^{Zl}|^v) \setminus (\Sigma_{\bar{A}}^{Fw} \cup [\text{esp}, \text{cs}_e]))$ for $(\text{STORE}_{\bar{A}})$, in the modified semantics of $(\text{LOAD}_{\bar{A}})$ and $(\text{STORE}_{\bar{A}})$ will evaluate to **false** for σ due to the assumption of error-free execution. As these checks do not refer to the potentially modified regions $\{hp, cl, cs, \text{free}\}$, σ' must also evaluate the check to **false** (for the same sequence of non-deterministic choices). Notice that this reasoning relies on the safety-relaxed semantics, and would not hold on the original semantics. All other state elements observed in the other instructions of $(\text{LOAD}_{\bar{A}})$ and $(\text{STORE}_{\bar{A}})$ are identical in both σ, σ' and $\text{sim}(\sigma, \sigma')$ will hold at the end of error-free executions.

□

Recall that the DYNAMO algorithm populates the deterministic choice map \mathcal{D}_X such that the result of the *choose* instruction $(\theta(i_{32}))$ for α_b in an `alloc` instruction in ξ_C matches the address v in an `allocs,v` instruction in $\xi_{\bar{A}}$ and the result of the *choose* instruction for memory contents $(\theta(i_{32} \rightarrow i_8))$ of the freshly allocated interval $[\alpha_b, \alpha_e]$ matches the memory contents of the interval $[v]_w$ (in the `alloc` and `allocs,v` instructions respectively). We use this fact in the following theorem on the execution of $[\xi_C]_{\mathcal{D}_X}^{ex}$.

Lemma 5.3.6 (*sim*(σ, σ') is preserved for error-free execution across all non-I/O edges in \mathcal{E}_C). *If a non-I/O edge $e_C \in \mathcal{E}_C$ in the path $[\xi_C]_{\mathcal{D}_X}^{ex}$ is executed on both machine states σ and σ' , and if $\text{sim}(\sigma, \sigma')$ holds before the execution, and if the execution on σ , with non-deterministic choices determined by \mathcal{D}_X , completes without error, then, for the same sequence of non-deterministic choices, the execution on σ' completes without error and $\text{sim}(\sigma, \sigma')$ holds at the end of both error-free executions.*

Proof. For a non-I/O C instruction that does not refer to the $\{hp, cl, cs, \text{free}\}$ regions ((OP), (ASSIGNCONST), (DEALLOC), (VASTARTPTR), (CALLV), (CALL_C), (RETC), (RETV)), the execution will have identical behaviour on both σ and σ' as identical values will be observed in both σ and σ' . Thus, if an execution on σ' makes the same non-deterministic choice as the execution on σ , the execution on σ' will complete without error and $\text{sim}(\sigma, \sigma')$ will hold at the end of both executions.

We consider each remaining non-I/O instruction in C below:

- (ENTRY_C) Consider the overlap check $Y = \text{ov}(\Sigma_C^{hp}, \Sigma_C^{cl}, \dots, i_C^g, \dots, \Sigma_C^f, \dots, i_C^y, \dots, \Sigma_C^{\text{vrdc}})$ (due to `!addrSetsAreWF`). During an execution on σ , this condition must evaluate to **false**, as we assume an error-free execution on σ . For the same non-deterministic choices made in both executions (over σ and σ'), by the definitions

of $HP(\xi_{\bar{A}})$ and $CL(\xi_{\bar{A}})$, Υ will also evaluate to **false** for an execution on σ' — recall that $(HP(\xi_{\bar{A}}) \cup CL(\xi_{\bar{A}}))$ cannot overlap with other allocated regions (due to lemma 5.3.4). Further, because all other state elements observed during the execution of the non-I/O edges in $(ENTRY_C)$ are identical in both σ and σ' , $sim(\sigma, \sigma')$ will hold at the end of error-free executions.

- (ALLOC) Consider the negated subset check $\Upsilon = \neg([\alpha_b, \alpha_e] \subseteq \Sigma_C^{\text{free}})$ (due to $\neg \text{intrvlInSet}_a(\alpha_b, \alpha_e, \Sigma_C^{\text{free}})$). The execution on σ must evaluate Υ to **false** as we assume an error-free execution. By the definitions of $HP(\xi_{\bar{A}})$ and $CL(\xi_{\bar{A}})$, for the same non-deterministic choices made in both executions (over σ and σ'), Υ will also evaluate to **false** for an execution on σ' — recall that during execution on σ , the deterministic choice map \mathcal{D}_X will be used for the non-deterministic choices of address α_b and memory $\pi_{[\alpha_b, \alpha_e]}(M_C)$ such that the freshly allocated interval $[\alpha_b, \alpha_e]$ matches (in both address and data) the allocated interval $[v]_w$ in an $\text{alloc}_{s,v}$ instruction in $\xi_{\bar{A}}$; because the same \mathcal{D}_X is used in both σ and σ' executions, Υ will also evaluate to **false** in σ' . All other state elements observed in the other instructions of (ALLOC) are identical in both σ, σ' .
- (LOAD_C) and (STORE_C). An $\text{accessIsSafeC}_{\tau,a}()$ check must evaluate to **true** for σ due to the assumption of error-free execution. Because the allocated space Σ_C^B can only be bigger in σ' (by lemma 5.3.2), the accessIsSafeC check will also evaluate to **true** for σ' (for the same sequence of non-deterministic choices). Further, for an execution on σ , the contents of the memory region $\pi_{\Sigma_{\sigma'}^{\{hp,cl\}} \setminus \Sigma_{\sigma}^{\{hp,cl\}}}(M_C^{\sigma})$ cannot be observed on an error-free path; and because all other state elements observed in (LOAD_C) and (STORE_C) are identical in both σ and σ' , the contents of the memory region $\pi_{\Sigma_{\sigma'}^{\{hp,cl\}} \setminus \Sigma_{\sigma}^{\{hp,cl\}}}(M_C^{\sigma'})$ will also remain unobserved during an execution on σ' (that uses the same sequence of non-deterministic choices as an execution on σ). All other state elements observed in the other instructions of (LOAD_C) and (STORE_C) are identical in both σ, σ' .

□

Lemma 5.3.7 ($sim(\sigma, \sigma')$ is preserved for error-free execution across $\xi_{\bar{A}}; \xi_C$). *Recall that $\xi_{\bar{A}}$ contains only non-I/O instructions (by assumption). Thus, due to the (SingleIO) requirement, ξ_C also contains only non-I/O instructions.*

If $\xi_{\bar{A}}$ is executed on machine states σ and σ' , and if the execution of σ completes without error, then there exists a sequence of non-deterministic choices during the

execution of σ' such that the execution is error-free and $\text{sim}(\sigma, \sigma')$ holds at the end of both error-free executions.

Similarly, if ξ_C is next executed on machine states σ and σ' , and if the execution of σ completes without error, then there exists a sequence of non-deterministic choices during the execution of σ' such that the execution is error-free and $\text{sim}(\sigma, \sigma')$ holds at the end of both error-free executions.

Proof. To show this, we execute the sequence of paths $(\xi_{\bar{A}}; \xi_C)$ in lockstep on both σ and σ' , i.e., in a single step, one instruction is executed on both states modifying the states in place. The proof proceeds by induction on the number of steps. The base case holds by assumption. For the inductive step, we rely on lemmas 5.3.5 and 5.3.6. \square

Lemma 5.3.8 (σ and σ' execute the same path in \bar{A}). *If $\xi_{\bar{A}}$ executes to completion on state σ , it will also execute to completion on σ' .*

Proof. By case analysis on all edge conditions in figs. 2.6 to 2.8, 2.10 and 2.11. For $\xi_{\bar{A}} = n_{\bar{A}} \rightarrow \mathcal{U}_{\bar{A}}$ due to $(\text{LOAD}_{\bar{A}})$ and $(\text{STORE}_{\bar{A}})$, the proof relies on the safety-relaxed semantics, and would not hold on the original semantics. \square

Lemma 5.3.9 (σ and σ' execute the same non- \mathcal{U} path in C). *If ξ_C does not terminate in \mathcal{U}_C , and σ executes ξ_C to completion, then σ' will also execute ξ_C to completion.*

Proof. By case analysis on all edge conditions in figs. 2.4 and 2.5 with same arguments as used in lemma 5.3.6. \square

Lemma 5.3.10 ($\text{post}(\sigma') \wedge \text{sim}(\sigma, \sigma') \Rightarrow \text{post}(\sigma)$ holds for an error-free node $(n_{\bar{A}}^t, n_C^t)$). *For two states σ and σ' at node $(n_{\bar{A}}^t, n_C^t)$, where $n_{\bar{A}}^t$ and n_C^t are error-free nodes, $\text{post}(\sigma') \wedge \text{sim}(\sigma, \sigma') \Rightarrow \text{post}(\sigma)$ holds.*

Proof. The *post* condition that may appear in a Hoare triple proof obligation generated by DYNAMO can be one of the following:

- (Coverage C) where $\text{post} = \bigvee_{1 \leq j \leq m} \text{pathcond}([\xi_C^j]_{\mathcal{D}_X}^{e_X^j})$ for $e_X^j = (n_X \xrightarrow{\xi_{\bar{A}}; \xi_C^j} (n_{\bar{A}}^t, n_C^t)) \in \mathcal{E}_X$ ($1 \leq j \leq m$).
- (Inductive) where *post* is one of the predicate shapes listed in fig. 4.2. Note that the MemEq shape in fig. 4.2 represents the proof obligation for the (MemEq) requirement.

- (Equivalence) where $post$ is either $\Omega_{\ddot{A}} = \Omega_C$ or $T_{\ddot{A}} =_{st} T_C$. I/O free paths do not mutate world states so $\Omega_{\ddot{A}} = \Omega_C$ holds trivially for these cases. Further, the only I/O free paths that may modify trace must contain `halt` instruction, appearing as the last edge of the sequence. As the generated trace event for `halt` does not observe any procedure state variable, we ignore this case.
- (MAC) where $post$ checks the address of each memory access in \ddot{A} against the addresses of a set of memory accesses in C for equality. Also, (MAC) checks if a memory access overlaps with address regions $\Sigma_{\ddot{A}}^{G \cup F} \cup [esp, \boxed{stk_e}]$ or $\Sigma_{\ddot{A}}^{G_w \cup F_w} \cup [esp, \boxed{stk_e}]$.

Case: When $post$ is one of the predicate shapes in fig. 4.2 or is a (MAC) proof obligation.

- The predicate shapes `affine`, `ineqC`, `ineq`, `spOrd`, `zEmpty`, `spzBd`, `spzBd'`, and a (MAC) proof obligation do not involve operations over address sets $\{hp, cl, cs, free\}$ or memory operations in the updated region $\Sigma_{\sigma'}^{\{hp, cl\}} \setminus \Sigma_{\sigma}^{\{hp, cl\}}$. Thus, $post(\sigma') \wedge sim(\sigma, \sigma') \Rightarrow post(\sigma)$ holds in this case.
- Consider the case when $post$ is `AllocEq`. Due to (Equivalence), `AllocEq` is guaranteed to be in pre and therefore $\Sigma_{\ddot{A}}^{hp} = \Sigma_C^{hp}$ and $\Sigma_{\ddot{A}}^{cl} = \Sigma_C^{cl}$ hold for σ' . Due to $sim(\sigma, \sigma')$, σ and σ' agree on the remaining state elements, including the address sets for each region $z \in Z$. Thus, $post(\sigma') \wedge sim(\sigma, \sigma') \Rightarrow post(\sigma)$ holds in this case.
- Consider the case when $post$ is `MemEq`. $sim(\sigma, \sigma')$ ensures that the address sets of regions $\{hp, cl\}$ in σ are a subset of respective address sets in σ' . Further, due to $sim(\sigma, \sigma')$, the memory states of A in σ and σ' are identical, $M_{\ddot{A}}^\sigma = M_{\ddot{A}}^{\sigma'}$, and the memory states of C in σ and σ' disagree only over the updated (expanded) address sets, $M_C^\sigma =_{\text{comp}(\Sigma_{\sigma'}^{\{hp, cl\}} \setminus \Sigma_{\sigma}^{\{hp, cl\}})} M_C^{\sigma'}$. Because the allocated regions in σ do not belong to these (expanded) addresses, $post(\sigma)$ follows from $post(\sigma')$.

Case: When $post$ is a proof obligation for (CoverageC). In this case, $post$ must be of the form $\bigvee_{1 \leq j \leq m} pathcond([\xi_C^j]_{\mathcal{D}_X}^{e_X^j})$ for $e_X^j = ((n_{\ddot{A}}, n_C) \xrightarrow{\xi_{\ddot{A}}; \xi_C^j} (n_{\ddot{A}}^t, n_C^t)) \in \mathcal{E}_X$ ($1 \leq j \leq m$). The edge conditions in C are independent of the regions $\{hp, cl, cs, free\}$, except for (LOAD_C) and (STORE_C). If the edge condition is independent of these address regions, then $post(\sigma)$ follows trivially from $post(\sigma')$. Consider the other case now: for an error-free node, the \mathcal{U} -maximal set of paths $\{\xi_C^1, \dots, \xi_C^m\}$ includes both the paths that evaluate `accessIsSafeC\tau, a` to `true` and `false` respectively. Thus, even in this case,

$post(\sigma)$ holds if $post(\sigma')$ holds.

□

Lemma 5.3.11 ($post(\sigma') \Rightarrow post(\sigma)$ for $n_{\ddot{A}}^t = \mathcal{W}_{\ddot{A}}$). For two states σ and σ' at node $(\mathcal{W}_{\ddot{A}}, n_{\ddot{C}}^t)$, $post(\sigma') \Rightarrow post(\sigma)$ holds.

Proof. The $post$ condition of this type may appear in a Hoare triple proof obligation generated by DYNAMO for one of the following:

- (Coverage \mathbb{C}) where $post = \bigvee_{1 \leq j \leq m} pathcond([\xi_{\mathbb{C}}^j]_{\mathcal{D}_X}^{e_X^j})$ for $e_X^j = (n_X \xrightarrow{\xi_{\ddot{A}}; \xi_{\mathbb{C}}^j} (n_{\ddot{A}}^t, n_{\mathbb{C}}^t)) \in \mathcal{E}_X$ ($1 \leq j \leq m$).
- (MAC) where $post$ checks the address of each memory access in \ddot{A} against the addresses of a set of memory accesses in \mathbb{C} for equality. Also, (MAC) checks if a memory access overlaps with address regions $\Sigma_{\ddot{A}}^{G_{UF}} \cup [esp, \boxed{stk_e}]$ or $\Sigma_{\ddot{A}}^{G_w \cup F_w} \cup [esp, \boxed{stk_e}]$.

The proof arguments for both these cases are identical to the ones made in the proof for lemma 5.3.10. □

Lemma 5.3.12 ($post(\sigma') \Rightarrow post(\sigma)$ for $n_{\ddot{A}}^t = \mathcal{U}_{\ddot{A}}$). For two states σ and σ' at node $(\mathcal{U}_{\ddot{A}}, n_{\mathbb{C}}^t)$, $post(\sigma') \Rightarrow post(\sigma)$ holds.

Proof. The $post$ condition of this type may appear in only one type of proof obligation generated by DYNAMO:

- (Coverage \mathbb{C}) where $post = \bigvee_{1 \leq j \leq m} pathcond([\xi_{\mathbb{C}}^j]_{\mathcal{D}_X}^{e_X^j})$ for $e_X^j = (n_X \xrightarrow{\xi_{\ddot{A}}; \xi_{\mathbb{C}}^j} (\mathcal{U}_{\ddot{A}}, n_{\mathbb{C}}^t)) \in \mathcal{E}_X$ ($1 \leq j \leq m$).

Let the (Coverage \mathbb{C}) proof obligation be $\{\phi_{n_X}\}(\xi_{\ddot{A}}; \epsilon)\{\bigvee_{1 \leq j \leq m} pathcond([\xi_{\mathbb{C}}^j]_{\mathcal{D}_X}^{e_X^j})\}$. Due to (Safety), each path $\xi_{\mathbb{C}}^j$ must end at $\mathcal{U}_{\mathbb{C}}$.

From the semantics in figs. 2.4 and 2.5, if the path condition for $\xi_{\mathbb{C}}^j$ evaluates to **true** on σ' (for some j), it must also evaluate to **true** on σ — in other words, whenever σ' transitions to $\mathcal{U}_{\mathbb{C}}$, σ is guaranteed to transition to $\mathcal{U}_{\mathbb{C}}$. This is because the edge conditions in \mathbb{C} will evaluate either identically on σ and σ' (due to $\{\xi_{\mathbb{C}}^1, \dots, \xi_{\mathbb{C}}^m\}$ being a \mathcal{U} -maximal set), or in the case of $\neg accessIsSafeC_{\tau, a}()$, the edge condition will evaluate to **true** on σ if it evaluates to **true** on σ' .

Thus, if $post(\sigma')$ evaluates to **true**, $post(\sigma)$ also evaluates to **true**. □

Proof for (\Leftarrow). Follows from lemmas 5.3.7 to 5.3.12. □

Proof of theorem 5.3.1. Follows from (\Rightarrow) and (\Leftarrow). □

5.4 Semantics with Simpler SMT Encoding for stk Region of \ddot{A}

In our implementation, we use a slightly revised semantics for \ddot{A} for a more efficient SMT encoding. We define a new region stk^+ such that stk^+ is large enough to contain the stk region at every point during \ddot{A} 's execution. A key property that stk^+ satisfies is that $\alpha \in \Sigma_{\ddot{A}}^{stk}$ can be rewritten in terms of $\alpha \in \Sigma_{\ddot{A}}^{stk^+}$ while $\alpha \in \Sigma_{\ddot{A}}^{stk^+}$ itself having simpler SMT encoding than $\Sigma_{\ddot{A}}^{stk}$.

(ENTRY $_{\ddot{A}}$ '), (OP-ESP''), (ALLOCS''), and (DEALLOCS'') in fig. 5.1 present the new semantics for procedure-entry, a stackpointer updating instruction, and `alloc`, and `dealloc`, instructions respectively. As with previous presentations, we only show the changes (with respect to figs. 2.6 to 2.8 and 2.11) with appropriate context in fig. 5.1; additions are highlighted and deletions are canceled.

(ENTRY $_{\ddot{A}}$ ') initializes the address set $\Sigma_{\ddot{A}}^{stk^+}$ such that $\Sigma_{\ddot{A}}^{stk} \subseteq \Sigma_{\ddot{A}}^{stk^+}$ holds and $\Sigma_{\ddot{A}}^{stk^+}$ does not overlap with other allocated regions in \ddot{A} . The lower bound of $\Sigma_{\ddot{A}}^{stk^+}$ is defined by a ghost variable `stkb` which is required to be below `esp` at all times for an error-free execution. Intuitively, `stkb` corresponds to the lowest stackpointer value seen during the execution of \ddot{A} ⁵.

(OP-ESP'') shows the updated semantics for a stackpointer (`esp`) updating instruction. Our modification is limited to the address set argument (third argument) of the $\{\mathcal{W}, \mathcal{U}\}$ -guarding $\neg \text{intrvlInSet}(\dots)$ predicate in both stack allocation and deallocation cases of the instruction. $\text{intrvlInSet}(\alpha_b, \alpha_e, i)$ returns **false** if the interval $[\alpha_b, \alpha_e]$ does not completely lie within the address set i . In the stack allocation case (under `if (isPush(...))`), the address set argument $\Sigma_{\ddot{A}}^{\text{free}} \cup ((\Sigma_{\ddot{A}}^{cv} \cup (\Sigma_{\ddot{A}}^{Zl|v})) \setminus \Sigma_{\ddot{A}}^F)$ to $\text{intrvlInSet}(\dots)$ is replaced with $\Sigma_{\ddot{A}}^{stk^+}$. This ensures that the stk^+ region is always big enough to accommodate a stack push and, consequently, imply that $\Sigma_{\ddot{A}}^{stk} \subseteq \Sigma_{\ddot{A}}^{stk^+}$

⁵Note the similarity with $SP_{min}(\xi_{\ddot{A}})$ from section 5.3.2

$$\begin{array}{c}
\text{(ENTRY}_{\check{A}}\text{'}) \frac{p_{\check{A}}^j : \text{def } \check{A}(\vec{\tau})}{\dots} \\
\begin{array}{l}
\Sigma_{\check{A}}^{stk} := [\text{esp}, \text{stk}_e] \setminus \Sigma_{\check{A}}^Y; \\
\dots \\
\text{stk}_b := \theta(i_{32}); \\
\text{if } (\neg(\text{stk}_b \leq_u \text{esp}) \vee \text{ov}([\text{stk}_b, \text{stk}_e], \Sigma_{\check{A}}^{GUFU\{hp\}})); \\
\text{halt}(\mathcal{W}); \\
\Sigma_{\check{A}}^{stk^+} := [\text{stk}_b, \text{stk}_e] \setminus \Sigma_{\check{A}}^Y;
\end{array} \\
\text{(OP-ESP'')} \frac{p_{\check{A}}^j : \text{esp} := \text{op}(\vec{x})}{\dots} \\
\begin{array}{l}
\text{if } (\text{isPush}(p_{\check{A}}^j, \text{esp}, t)) \{ \\
\text{if } (\neg \text{intrvlInSet}(t, \text{esp} - 1_{i_{32}}, \Sigma_{\check{A}}^{stk^+} \setminus \Sigma_{\check{A}}^{\text{free}} \cup ((\Sigma_{\check{A}}^{\text{cv}} \cup (\Sigma_{\check{A}}^{\text{ZU}})) \setminus \Sigma_{\check{A}}^E))) \\
\text{halt}(\mathcal{W}); \\
\Sigma_{\check{A}}^{stk} := \Sigma_{\check{A}}^{stk} \cup [t, \text{esp} - 1_{i_{32}}]; \\
\} \text{elseif } (t \neq \text{esp}) \{ \\
\text{if } (\neg \text{intrvlInSet}(\text{esp}, t - 1_{i_{32}}, \Sigma_{\check{A}}^{stk^+} \setminus \Sigma_{\check{A}}^{stk})) \\
\text{halt}(\mathcal{U}); \\
\Sigma_{\check{A}}^{stk} := \Sigma_{\check{A}}^{stk} \setminus [\text{esp}, t - 1_{i_{32}}]; \\
\} \\
\dots
\end{array} \\
\text{(ALLOCS'')} \frac{p_{\check{A}}^j : \text{alloc}_s e_v, e_w, a, z}{\dots} \quad \text{(DEALLOCS'')} \frac{p_{\check{A}}^j : \text{dealloc}_s z}{\dots} \\
\begin{array}{l}
\Sigma_{\check{A}}^z |^s, \\
\Sigma_{\check{A}}^{stk}, \\
\Sigma_{\check{A}}^{stk^+} \\
\dots \\
\Sigma_{\check{A}}^z, \\
\Sigma_{\check{A}}^{stk}, \\
\Sigma_{\check{A}}^{stk^+} \\
\dots
\end{array}
:=
\begin{array}{l}
\Sigma_{\check{A}}^z |^s \cup [v]_w \\
\Sigma_{\check{A}}^{stk} \setminus [v]_w, \\
\Sigma_{\check{A}}^{stk^+} \setminus [v]_w; \\
\dots \\
\Sigma_{\check{A}}^z \cup [v]_w, \\
\Sigma_{\check{A}}^{stk} \setminus [v]_w, \\
\Sigma_{\check{A}}^{stk^+} \setminus [v]_w; \\
\dots
\end{array}
\quad
\begin{array}{l}
\Sigma_{\check{A}}^z |^s, \\
\Sigma_{\check{A}}^{stk}, \\
\Sigma_{\check{A}}^{stk^+} \\
\dots \\
\Sigma_{\check{A}}^z, \\
\Sigma_{\check{A}}^{stk}, \\
\Sigma_{\check{A}}^{stk^+} \\
\dots
\end{array}
:=
\begin{array}{l}
\emptyset, \\
\Sigma_{\check{A}}^{stk} \cup \Sigma_{\check{A}}^z |^s, \\
\Sigma_{\check{A}}^{stk^+} \cup \Sigma_{\check{A}}^z |^s; \\
\dots \\
\emptyset, \\
\Sigma_{\check{A}}^{stk} \cup \Sigma_{\check{A}}^z, \\
\Sigma_{\check{A}}^{stk^+} \cup \Sigma_{\check{A}}^z; \\
\dots
\end{array}
\end{array}$$

Figure 5.1: Revised translation rules for the new stk^+ -based semantics for \check{A} .

continues to hold after a stack allocation. Similarly, in the stack deallocation case (under elseif $(t \neq \text{esp})$), $\Sigma_{\check{A}}^{stk}$ is replaced with $\Sigma_{\check{A}}^{stk^+}$. Our modifications are carefully designed so that this modified check (for stack deallocation case) will evaluate identically to the original check in an execution of \check{A} .

(ALLOCS'') and (DEALLOCS'') shows the updated semantics for alloc_s and dealloc_s instructions. In both cases, $\Sigma_{\check{A}}^{stk^+}$ is updated identically to $\Sigma_{\check{A}}^{stk}$ for a local (de)allocation from(to) stack.

Theorem 5.4.1. *The following property holds at every error-free, non-entry node $n_{\ddot{A}} \in \mathcal{N}_{\ddot{A}}^{\text{SW}}$ in an execution of \ddot{A} with the semantics presented in fig. 5.1 (and semantics presented earlier in figs. 2.5, 2.6, 2.8, 2.10 and 2.11).*

$$\alpha \in \Sigma_{\ddot{A}}^{stk} \Leftrightarrow \alpha \in \Sigma_{\ddot{A}}^{stk^+} \wedge (\alpha \geq_u \text{esp})$$

Proof sketch: By induction on the number of transitions executed in \ddot{A} with the base case defined by the first transition out of (ENTRY $_{\ddot{A}}$) in fig. 5.1. \square

Using the above property, the SMT encoding of a verification condition over \ddot{A} is rewritten to replace references to $\Sigma_{\ddot{A}}^{stk}$ with references to $\Sigma_{\ddot{A}}^{stk^+}$. All assignments to $\Sigma_{\ddot{A}}^{stk}$ are made vacuous and only $\Sigma_{\ddot{A}}^{stk^+}$ is tracked in the allocation state of \ddot{A} . Because $\Sigma_{\ddot{A}}^{stk^+}$ is not updated due to a stackpointer updating instruction, the resulting SMT expressions are simpler.

Chapter 6

Evaluation

This chapter discusses the implementation specifics and evaluation of a prototype tool based on the DYNAMO algorithm described in the previous chapters.

6.1 Implementation of DYNAMO

6.1.1 System components

Figure 6.1 shows the high-level component blocks and the flowchart of our tool, also named DYNAMO. The two programs, C source code and its (potentially optimized) x86 assembly, are translated to their respective graph representations before they are passed as input to the tool. The C source is first translated to LLVM_d (which is based on unoptimized LLVM IR; see section 2.1.1) using a modified Clang/LLVM [48]. The explicit `dealloc` instructions, present only in LLVM_d, are also added during this translation using the `stacksave` and `stackrestore` intrinsics — these intrinsics are generated for the purpose of stack related code generation, we use them for deriving scope information of allocated locals. The generated LLVM_d program is then lowered to its graph representation **C** (described in section 2.3.1) and a may-point-to analysis is performed (described in section 4.1). The x86 assembly is similarly translated to its graph representation **A** (described in section 2.3.2); the translation uses the safety-relaxed semantics (described in section 3.5). Both **C** and **A** use the callers’ virtual smallest semantics (described in section 3.4).

The various components blocks of the tool are shown in fig. 6.1. The tool interfaces

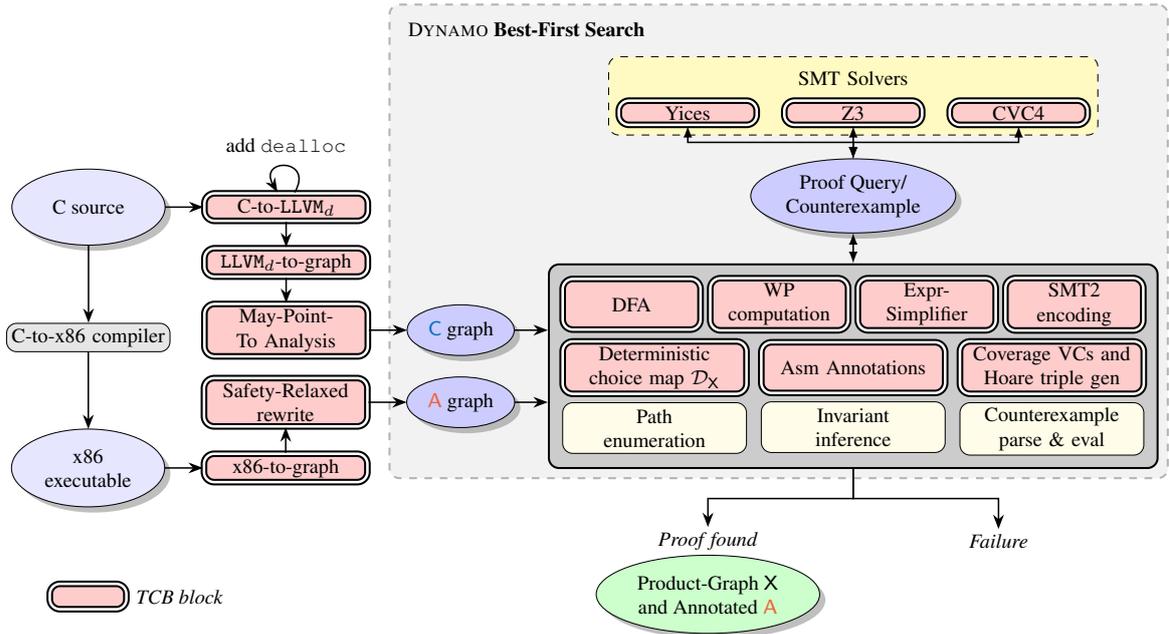


Figure 6.1: High-level components of DYNAMO implementation. Trusted Code Base (TCB) blocks have double border and a red background.

with third-party SMT solvers (“SMT Solvers” in the diagram) for discharging its proof obligations. The “SMT2 encoding” component encodes a Hoare triple into SMT-LIB2 format [7] that is used as a uniform input format for the SMT solvers. A model or counterexample returned by a solver is parsed, evaluated, and translated for guiding the correlation and invariant inference (“Counterexample parse & eval”). The other prominent components of the system shown in the diagram are: a common data-flow analysis (DFA) framework (“DFA”) a weakest-precondition computation module (“WP computation”), a syntactic expression simplifier (“Expr-Simplifier”), and a verification condition generator (“Coverage VCs and Hoare triple gen”). The modules for the deterministic choice map (described in section 3.2) and assembly annotations are shown separately as well (“Deterministic choice map \mathcal{D}_X ” and “Asm Annotations”).

Trusted Code Base

The soundness of a verification effort is critically dependent on the correctness of the *trusted computing base* (TCB) of the verifier — in fig. 6.1 the blocks belonging to the TCB are marked with double borders and have a red background. An error in the TCB may result in a soundness error.

DYNAMO is around 400K Source Lines of Code (SLOC) in C/C++ that has been in development for over a decade and used in prior work [4, 13, 42]. Out of the 400K, the TCB is around 70K SLOC. Within the TCB, around 20K SLOC is due to the expression handling and simplification logic, less than 10K for SMT encoding, less than 10K SLOC for the graph representation and the weakest-precondition logic and less than 1K SLOC for a common data-flow analysis framework. The x86-to-graph translation is around 18K SLOC of C code (for disassembly) and 5K SLOC of OCaml code (for logic encoding), and the IR-to-graph translation, including the addition of `dealloc` and the may-point-to analysis (section 2.3), is around 6K SLOC of C++ code. We rely on the Clang framework for the C-to-IR translation — one can imagine replacing this with a verified frontend, such as CompCert’s [30]. The modeling of the deterministic choice map, generation of Hoare triple and coverage verification conditions is relatively simple (less than 1K SLOC total).

6.1.2 Discharging Proof Obligations

Recall that we use quantified logic over bitvectors, arrays, and uninterpreted functions (AUFBV in SMT-LIB [7]) in our SMT encoding. For procedures without local (de)allocations and procedure calls, it is possible to use a quantifier-free encoding by modeling the stack as a separate memory array. Using a lightweight analysis on the **A** graph, each memory access may be labeled as either to stack region or *non-stack* (other than stack) regions. Utilizing the full-interval encoding (section 5.3), the various address set relations can then be rewritten in quantifier-free logic. Prior work [39, 42] has used similar approach for achieving quantifier-free encoding.

Use of multiple solvers: DYNAMO uses four SMT solvers running in parallel for discharging proof obligations: `Z3-4.8.7`, `Z3-4.8.14`, `Yices2-45e38fc`, and `CVC4-1.7`. Though `Yices2` does not support quantified logic (and cannot be used in the general setting), we use it fruitfully in the quantifier-free setting described in previous paragraph. For instance, in our experiments, we observed that it was consistently faster than the others. Because both our best-first search and invariant inference algorithm may potentially converge faster with a larger set of counterexamples, we invoke two versions of the `Z3` SMT solver: `Z3-4.8.7` and `Z3-4.8.14`. This decision was guided by an empirical observation where we noted that the two solvers have different performance characteristics and produce *sufficiently different* counterexamples that contribute to faster convergence of the invariant inference algorithm [42].

Handling difficult queries: As an optimization, our implementation makes use of the *query-decomposition* technique and simplification rewrites described in [18] for discharging queries difficult for the SMT solvers. We invoke query-decomposition only after the SMT solvers fail to discharge the proof query within a specified timeout period (we provide the timeout period used in each experiment in section 6.2).

6.1.3 Pseudo-register allocation in LLVM_d

Before checking refinement, if the address of a local variable l is never taken in \mathbb{C} , we transform \mathbb{C} to register-allocate l using LLVM’s `mem2reg` pass. This reduces the proof effort, at the cost of having to trust the pseudo-register allocation logic. However, `mem2reg` does not register-allocate local arrays and aggregates (`structs`) in LLVM_d, even though an optimizing compiler may register-allocate them in assembly — virtual allocations help validate such translations. Another case where virtual allocations are helpful is described in section 6.2.4 where a `va_list` variable (used in variadic procedures) is not (pseudo-)register allocated in \mathbb{C} by `mem2reg` but is register-allocated in optimized assembly. Thus, after `mem2reg`, virtual allocations are only required for validating register-allocation of arrays, aggregates, and `va_list` variables.

6.1.4 Instrumentation of Clang/LLVM for generating annotation hints

We instrumented the Clang/LLVM v12.0.0 compiler to generate annotation hints for the whitebox setting. The instrumentation, which required only a couple of lines of code, prints the constant offset (with respect to the stackpointer at entry of the procedure) of the allocated stack slots for *non-VLA* allocations — for VLA allocations we still rely on blackbox enumeration. These offsets do not map to a named local in \mathbb{C} and are only used as additional address options for an `allocs` annotation (section 4.1.3).

6.2 Experiments

We design our experiments to validate the capabilities and explore the limitations of the DYNAMO algorithm in our prototype tool. In particular, the experiments explore the following aspects of the tool (and the algorithm):

1. Handling of different constructs available for local allocation in the C language and

common extensions (section 6.2.1).

2. Performance of the full-interval and partial-interval SMT encoding in comparison to the naive full-array SMT encoding (section 6.2.1).
3. Overhead of identifying the required annotation for modeling local (de)allocations (section 6.2.2).
4. Performance of the tool on a real-world program (section 6.2.3).

6.2.1 Evaluating efficacy of DYNAMO

We first evaluate the efficacy of DYNAMO to handle the diverse programming patterns seen with local allocations in C. Table 6.1 lists the programming patterns we test in this experiment. Each programming pattern corresponds to one benchmark (i.e., a C procedure) except `vs1N` and `vilN` where we substitute N with 1, 2, 3 to obtain `vs11`, `vs12`, `vs13` and `vil1`, `vil2`, `vil3` respectively in each case. The programming patterns we include in this experiment have:

- use of address-taken local variables and parameters in benchmarks `ats`, `atss`, `atps`, `atpss`, and `atc`.
- use of constant- and variable-length arrays in benchmarks `ata`, `atail`, `fib`, `vin`, `vcu`, `vs11`, `vs12`, and `vs13`.
- variadic procedures in benchmarks `vw1`, `vw2`, `mp`, and `ms`.
- use of variable-length arrays (VLAs) allocated inside loops in benchmarks `vil1`, `vil2`, `vil3`, `vilcc`, and `vilce`.
- unconditional and conditional use of `alloca()` in `as`, `ac`, and `ams`.
- use of `alloca()` to create a linked-list in `all`.
- mixed use of constant-sized and variable-length array in `rod`.

There are total 27 benchmarks and we use three compilers, Clang/LLVM v12.0.0, GCC v8.4.0, and ICC v2021.8.0, to generate 32-bit x86 executables at `O3` optimization level to create 81 procedure-pairs. We manually disable interprocedural and unrolling/vectorization optimizations in each invocation using the compiler’s source-level `pragmas` and/or command-line flags¹. A refinement check is performed for each of the 81 procedure-pairs. We use an unroll-factor of one $\mu = 1$ for all benchmarks except `all` for GCC and ICC, and `fib` for Clang/LLVM where we use an unroll-factor of two $\mu = 2$. Each refinement

¹The full command-line used for generating the assembly in each case is provided in section B.1.

#	Name	Programming pattern	Code
1	ats	Address-taken local scalar	<code>int ats() { int ret; foo(&ret); return ret; }</code>
2	atss	Address-taken local struct	<code>int atss() { struct Point p1, p2; foo(&p1, &p2); return ...; }</code>
3	atps	Address-taken scalar parameter	<code>int atps(int a) { char x; scanf("%c %d", &x, &a); return ...; }</code>
4	atpss	Address-taken struct parameter	<code>int atpss(struct Point p1, ...) { ...; scanf(..., &p1.x); ... }</code>
5	atc	Address taken conditionally	<code>int atc(int* p) { int x; if (!p) p = &x; foo(p); return *p }</code>
6	ata	Local array	<code>int ata() { char ret[8]; foo(ret); return bar(ret, 0, 16); }</code>
7	atail	Local array alloc. in loop	<code>int atail(..){..for(..){ char a[4096]; f(a..); b(a..);...}..}</code>
8	fib	Program from fig. 2.1	
9	vin	VLA's in nested scopes	<code>void vin(int n, int m){int v1[n]; { int v2[m]; foo(v2); } bar(v1); }</code>
10	vcu	VLA conditional use	<code>int vcu(int n,int k){ int a[n]; if (...) { /*rd/wr to a*/}...}</code>
11	vs1N	N VLA(s)	<code>int vs1N(..){ .. int v1[n], ... vN[n]; for(i=0;i<n;++i) { v1[i]=..a[i]..; vN[i]=..a[i]..; } return fooN(...); }</code>
12	vilN	N VLA(s) in a loop	<code>int vilN(..){..for(i=1;i<n;++i) { int v1[4*i], ... vN[4*i]; fooN(..); ..}</code>
13	vilcc	VLA in loop with continue	<code>int vilcc(..){..while(i<n){ char v[i];...if(..) continue;...}..}</code>
14	vilce	VLA in loop with break	<code>int vilce(..){..while(i<n){ char v[i];...if(..) break;...}..}</code>
15	vwl	Variadic procedure	<code>int vwl(int n,...){ va_list a; va_start(a, n); for(..){/* va_arg(a,int) */}..}</code>
16	vwc	Variadic procedure using va_copy()	<code>vwc(int n,...) { va_list a, b; va_start(a,n); va_copy(b,a); for(..){ /* va_arg(b) */ va_copy(b,a); foo(b); ...}</code>
17	mp	minprintf procedure adapted from K&R [24]	
18	ms	minscanf procedure similar to minprintf	
19	as	alloca()	<code>int as(int n){...int* p=alloca(n*sizeof(n)); for(..){/*write to p*/}...}</code>
20	ac	alloca() conditional use	<code>int ac(char*a) {..if (!a) a=alloca(n); for(..){/*r/w to a*/}</code>
21	all	alloca() in a loop to form a linked list	<code>all() { ...hd=NULL; for(..){ ...n=alloca(..); n->nxt=hd; hd=n; } while(..){ /* traverse the list starting at hd */ }</code>
22	ams	alloca() & malloc()	<code>int ams(..){..if(..){p=alloca(..);}else{p=malloc(..);}/* r/w to p ..*/..}</code>
23	rod	A local char array initialized using a string and a VLA and a for loop	Available in ??.

Table 6.1: Benchmarks and their programming patterns. N in `vs1N` and `vilN` is substituted to obtain `vs11`, `vs12`, `vs13` and `vil1`, `vil2`, `vil3` respectively. Full program listings available in chapter C.

check is run with a refinement check timeout of two hours and a per SMT query timeout of 120 seconds, i.e., the refinement check is terminated automatically after two hours, resulting in a refinement failure, and, similarly, an SMT solver spawned for discharging an SMT query is automatically terminated after 120 seconds, yielding a “timeout” result which is interpreted as a failure in the discharge of the corresponding proof obligation — a failed proof obligation may eventually result in refinement failure.

Figure 6.2 shows graph of the refinement check run times for the 81 procedure-pairs. The X-axis lists the benchmarks and the Y-axis represents the total time taken in seconds (using a log scale) for a refinement check. To study the performance implications, we

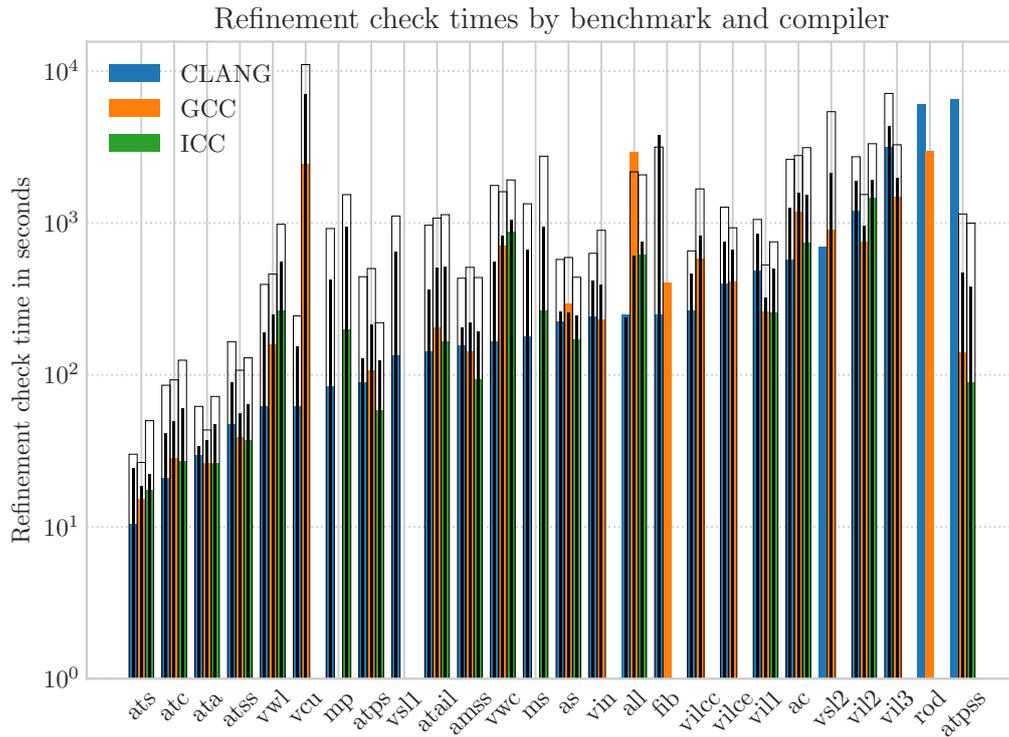


Figure 6.2: Comparison of running times of procedures in table 6.1 with full-interval (filled bars), partial-interval (thick black lines), and full-array (empty bars) encoding. Y-axis is logarithmically scaled.

run a check with all three encodings for these benchmarks:

- A filled bar (■) represents the time taken with, if it was possible to use, a full-interval encoding.
- A thick black line centered on a bar (▬) represents the time taken with a *forced* partial-interval encoding, i.e., partial-interval encoding is forced even if a full-interval encoding is possible.
- An empty bar (□) represents represents the time taken with a *forced* full-array encoding, i.e., a full-array encoding is forced even if a full-interval or a partial-interval encoding is possible.
- A missing bar or a missing thick line (for partial-interval encoding) represents a failure in proof search with respective encoding.

Of the 81 procedure pairs, our tool is able to find a refinement proof for 66, 61, and

60 procedure-pairs while using full-interval, partial-interval, and full-array encodings respectively. For benchmarks where a refinement check succeeds for all encodings, the full-interval encoding performs 2.07 and 3.88 times faster on average (for each compiler) than the partial-interval and full-array encodings respectively.

Encoding	# of successful refinement checks			
	Clang/LLVM	GCC	ICC	Total
Full-interval encoding	26	23	17	66
Partial-interval encoding	23	21	17	61
Full-array encoding	22	21	17	60

Figure 6.3: Summary of refinement check results for the programs in table 6.1.

Figure 6.3 summarizes the results of this experiment, showing the number of successful refinement checks for each compiler and encoding pair. DYNAMO is unable to find refinement a proof in 15 cases for the full-interval encoding. We analyze these failures in detail in section 6.2.4. A failure particular to partial-interval or full-array encoding (i.e., the refinement check succeeded with a full-interval encoding) is either due to SMT solver timeout causing failure of a crucial proof obligation or a refinement check timeout in that invocation.

Two benchmarks, `vilcc` and `vilce`, require multiple `deallocs` instructions to be added to **A** for a single `dealloc` in **C**. An `allocv` annotation is required for the ‘`va_list a`’ variable in the GCC and ICC compilations of `vw1` and `vw2` (see table 6.1) — while GCC and ICC register-allocate `a`, it is allocated in memory using `alloc` in LLVM_d even after `mem2reg` (section 6.1.3). The time spent in constructing the correct product graph forms around 70-80% of the total search time.

6.2.2 Evaluating modeling cost of local allocations

Our next experiment evaluates the cost of modeling local allocations. Towards this, we run DYNAMO on the TSVC suite of vectorization benchmarks with arrays and loops [33]. The benchmarks in this suite are used to evaluate the vectorization capabilities of optimizing compilers and typically have C functions with (nested) loops and array accesses. A variation of this suite, where the floating-point types are replaced with integer types, was used in prior work on translation validation [8, 17] and we adopt the same modified version for our use — we select 25 procedures (as used in [17]) for our experiments.

```
int a[4000]; // global array -- 'a' is locally allocated in 'locals' version
int b[4000];
int s122(int n1, int n3) {
    int j, k;
    j = 1;
    k = 0;
    // int a[4000]; // this is uncommented in 'locals' version
    init_local1(a); // this call initializes 'a' in 'locals' version
    for (int i = n1-1; i < 4000; i += n3) {
        k += j;
        a[i] += b[4000 - k];
    }
    print_local1(a); // this call observes (outputs) 'a' in 'locals' version
    return 0;
}
```

Figure 6.4: Procedure `s122` from ‘globals’ version of (modified) TSVC suite.

We create two versions of this suite for our experiments:

1. ‘globals’ where global variables are used for storing the output array values.
2. ‘locals’ where local array variables are used for storing the output array values.

In both versions, a procedure call is added at the end of the procedure body to print the contents of the modified output array variables. Further, if the output array variable is being read before being assigned, we add a procedure call before the read to initialize the variable. Figure 6.4 shows an example.

We use Clang/LLVM v12.0.0 at optimization level 03 with vectorization enabled using `-msse 4.2` flag for compiling the 25 C procedures in each version. The compiler performs the same vectorizing transformations on both versions. Unlike `globals`, the benchmarks in `locals` version additionally require the automatic identification of required annotation. We use an unroll-factor of 64, a global timeout of two hours, and a per SMT query timeout of one minute for each run of DYNAMO on the 50 procedure-pairs.

Figure 6.5 shows the execution times of DYNAMO for these two versions of the TSVC benchmark. DYNAMO is able to successfully validate these compilations. Similar to fig. 6.2, we show execution times for both full-interval and *forced* partial-interval encodings for the ‘locals’ benchmarks using filled bars and empty bars respectively (we omit the *forced* full-array encoding numbers in this case). Compared to `globals`,

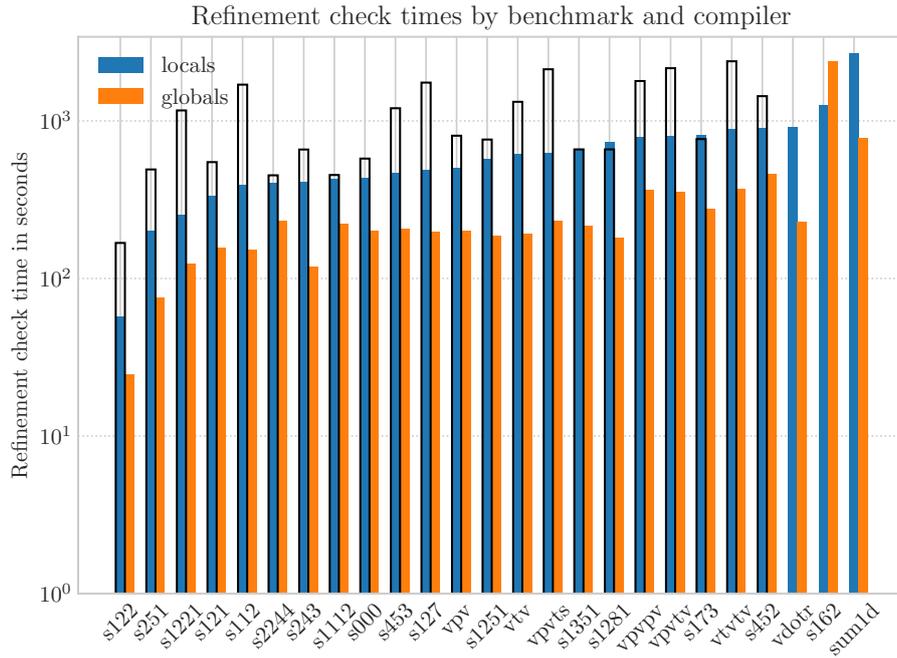


Figure 6.5: Comparison of running times of TSVC benchmarks with exactly same code modulo allocation strategy. Y-axis is logarithmically scaled.

refinement checks are 2.5x slower for `locals` (on average) due to the extra overhead of identifying the required annotations. The full-interval encoding is 1.9 times faster (on average) than partial-interval encoding. This relative speed-up of full-interval encoding over partial-interval encoding is similar to the number obtained in previous experiment (section 6.2.1). With a forced partial-interval encoding, DYNAMO fails to validate 3 (out of 25) benchmarks within the assigned time budget.

6.2.3 Evaluating DYNAMO on a real-world program

Our last experiment is on SPEC CPU2000’s `bzip2`[20] program compiled using Clang/LLVM v12.0.0 at three optimization levels: `O1`, `O2`, and `O1-`. `O1-` is a custom optimization level configured by us that enables all optimizations at `O1` *except* the following.

- (a) merging of multiple procedure calls on different paths into a single call,
- (b) early-CSE (common subexpression elimination),
- (c) loop-invariant code motion at both LLVM IR and Machine IR,
- (d) dead-argument elimination,

- (e) interprocedural sparse conditional constant propagation,
- (f) dead-code elimination of procedure calls.

These passes were chosen based on our knowledge of the limitations of our invariant inference algorithm. `bzip2` runs 2% slower with `01-` than with `01`; this is still 5% faster than the executable produced by CompCert, for example. Of all 72 procedures in `bzip2`, DYNAMO successfully validates the translations for 64, 63, and 57 procedures at `01-`, `01`, and `02` respectively at unroll-factor $\mu = 2$ and a per SMT query timeout of three minutes. At `01-`, DYNAMO takes around 8.7 CPU hours to compute refinement proofs for the 64 procedures. DYNAMO times out for the remaining 8 procedures, all of which are bigger than 142 ALOC.

Name	SLOC	ALOC	$\#_{al}$	$\#_{call}$	T	$ \mathcal{N}_X $	$ \mathcal{E}_X $	EXP	BT	$\#_q$	Avg. qT
<code>generateMTFValues</code>	76	144	1	1	4k	14	30	60	16	3860	0.56
<code>recvDecodingTables</code>	70	199	2	10	3k	38	66	102	15	5611	0.21
<code>undoReversible-Transformation.fast</code>	116	221	1	6	2k	21	34	43	6	2998	0.23

Table 6.2: Statistics obtained by running DYNAMO on procedures in the `bzip2` program.

Three of `bzip2`'s procedures for which refinement proofs are successfully computed at both `01-` and `01` optimization levels contain at least one local array, and table 6.2 presents statistics for the `01-` validation experiments for these procedures. For each procedure, we show the number of source lines of code in C (SLOC), the number of assembly instructions in A (ALOC), the number of local variables ($\#_{al}$), and the number of procedure calls ($\#_{call}$). The T column shows the validation times (in seconds). The $|\mathcal{N}_X|$ and $|\mathcal{E}_X|$ columns show the number of nodes and edges in the final product graph, and BT and EXP is the number of best-first search backtrackings and the number of (partial) candidate product graphs explored by best-first search in DYNAMO respectively. $\#_q$ is the total number of SMT queries discharged, and Avg. qT is the average time taken by an SMT query in seconds for the refinement check.

In the experiment with `02` optimization level some of the loops in these procedures are vectorized. With an unroll-factor $\mu = 8$, DYNAMO is able to validate `undoReversible-Transformation.fast`'s assembly implementation within ≈ 38 minutes. However, the BFS search diverges away from the correct solution for the other `bzip2` functions, eventually resulting in a timeout. These experiments indicate the need for more robust search algorithms and/or faster logical models/solvers in the presence of aggressive

optimizations on large programs.

In a separate experiment, we manually split the large procedures in `bzip2` into smaller procedures, so that DYNAMO successfully validates the O1- compilation of the full modified `bzip2` program in ≈ 16 hours: the splitting disables some compiler transformations and also reduces the correlation search space.

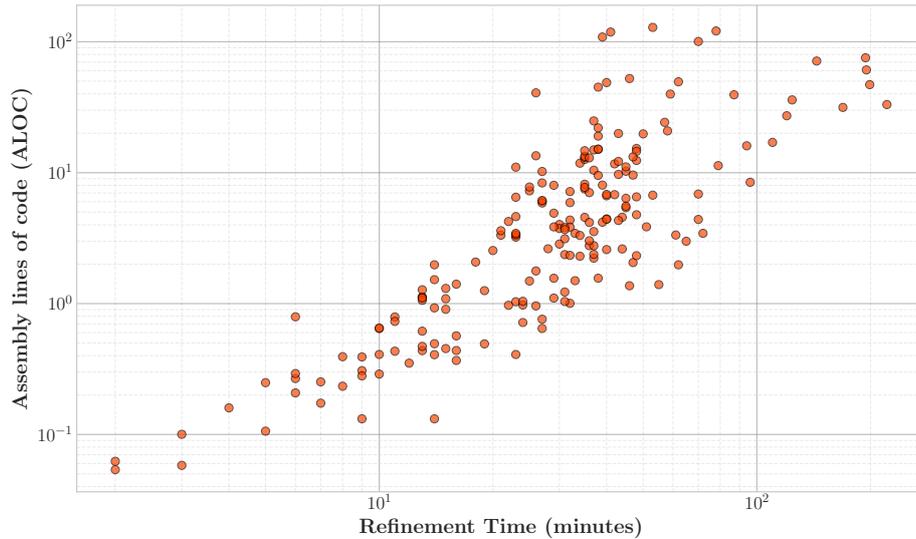


Figure 6.6: Scatter plot of refinement time (in minutes) vs assembly lines of code (ALOC). Both axes are logarithmically scaled.

Figure 6.6 shows the scatter plot of refinement time (in minutes) versus the assembly lines of code (ALOC) of all procedure-pairs considered so far. While the relationship is almost linear, the range of values is quite wide, especially at higher ALOC.

6.2.4 Analysis of Failures

Table 6.3 lists the 15 refinement check failures with full-interval encoding for the refinement checks in fig. 6.2. We have grouped the reasons for these failures into three categories: (1) limitation of blackbox annotation; (2) incompleteness due to affine invariant inference grammar; and (3) incompleteness in affine invariant inference due to chosen set of procedure variables. We discuss each reason in detail in following subsections.

#	Benchmark	Compiler	Failure reason
1	vs11	GCC	Limitation of <code>dealloc_s</code> annotation
2	vs11		
3	vs12		
4	vil3		
5	vilcc		
6	vilce	ICC	Non-affine invariant required
7	fib		
8	rod		
9	vin		
10	vcu		
11	vs13		
12	vs13	Clang	
13	vs13	GCC	
14	mp	GCC	Incompleteness in affine invariant inference due to the chosen set of procedure variables
15	ms		

Table 6.3: Failure reasons for the refinement checks shown in fig. 6.2.

Limitation of the `allocs/deallocs` annotation algorithm in the blackbox setting

Recall that in the blackbox setting of DYNAMO, when hints from the compiler are not available, the annotation algorithm (`asmAnnotOpts()`) limits the insertion of a `deallocs` instruction to only those PCs that occur just before an instruction that updates the stackpointer register `esp`. This limitation may cause a refinement check to fail in some (not all) of the situations where a compiler implements merging of multiple allocations (resp. deallocations) into a single stackpointer decrement (resp. increment) instruction. This is the reason for the failure to validate GCC’s compilation of `vs11`.

Figure 6.7 shows the `vs11` procedure (fig. 6.7a) and a sketch of the control-flow graph (CFG) of the assembly procedure generated by GCC at `03` optimization level (fig. 6.7b). The assembly path $S \rightarrow B \rightarrow J \rightarrow E$ represents the case when $(n \leq 0)$ and the procedure exits early (without allocating any local variable). PC with label L represents the loop head in the CFG, and the allocation and deallocation of the VLA v is supposed to happen before entering the loop and after exiting the loop respectively.

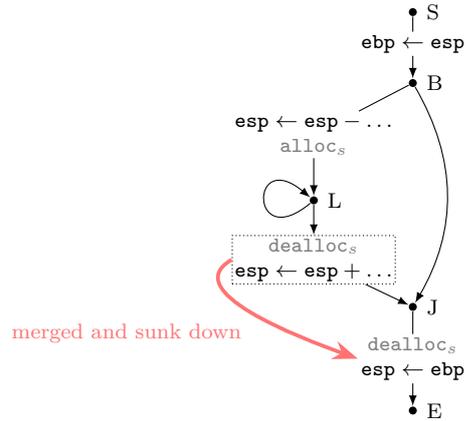
On the assembly CFG path $L \rightarrow J$, the assembly instruction that reclaims stack space (by incrementing the stackpointer) for deallocating v has been merged by the compiler with an instruction that restores the stackpointer to its value at the beginning of

```

int vs11(int *a, int n)
{
  if (n <= 0)
    return 0;
  int v1[n];
  for (int i = 0; i < n; ++i)
  {
    v1[i] = a[i]+a[i];
  }
  return foo1(v1);
}

```

(a) C source code



(b) Abbreviated CFG of GCC compiled assembly

Figure 6.7: `vs11` procedure from table 6.1 (appears as `vs1N`) and the control-flow graph (CFG) of its GCC compiled assembly.

the function (stored in register `ebp`). Thus, while the original stackpointer increment instruction would have been at the end of the $L \rightarrow J$ path, the merged instruction is sunk by the compiler to lie within the path $J \rightarrow E$. As can be seen, this transformation saves an extra instruction to update the stack pointer on the path $L \rightarrow J \rightarrow E$.

In the absence of compiler hints (blackbox setting), our tool considers the annotation of a `deallocs` instruction in assembly only at a PC that immediately precedes an instruction that updates the stackpointer. In this example, the only candidate PC for annotating `deallocs` (considered by our blackbox algorithm) is on the path $J \rightarrow E$. However, the required position of the `deallocs` instruction was at the end of the path $L \rightarrow J$ (which is not considered because there is no instruction that updates the stackpointer at the end of the path $L \rightarrow J$). Thus, our blackbox algorithm cannot find a refinement proof. On the other hand, providing a manual hint to the tool that it should consider annotating a `deallocs` instruction at the end of the $L \rightarrow J$ path causes the algorithm to successfully return a refinement proof for GCC's compilation of `vs11`.

For completeness, let us consider what happens when the tool annotates a `deallocs` instruction just before the instruction that updates the stackpointer to `ebp` on the $J \rightarrow E$ path. Such an annotation violates the (stuttering) trace equivalence condition on the procedure path $S \rightarrow B \rightarrow J \rightarrow E$: in the C procedure, there is no deallocation (or allocation) on the early exit path (when $n \leq 0$), but this annotation will cause a `deallocs` instruction to be executed on the correlated path ($S \rightarrow B \rightarrow J \rightarrow E$) in the

assembly procedure. Because a `deallocs` instruction generates a (non-silent) trace event through the `wr` instruction, this candidate annotation therefore fails to show the equivalence of traces on at least one pair of correlated paths. Thus, this candidate annotation is discarded by our algorithm.

Requirement of a non-affine invariant in some ICC compiled benchmarks

The ICC compiler (and in two cases GCC and Clang as well) generates a certain pattern of assembly code in some (not all) cases of VLA containing C source code that necessitate need of a non-affine invariant shape for completing the refinement proof.

In some cases, ICC uses the following pattern of instruction sequence for the allocation of a VLA of size `vlaSz`:

$$\begin{aligned} reg_1 &\leftarrow vlaSz \\ reg_2 &\leftarrow (reg_1 + C) \& \sim C \\ esp &\leftarrow esp - reg_2 \end{aligned}$$

Here, `reg1` and `reg2` are assembly registers other than `esp`, `esp` is the stackpointer register, `C` is a bitvector constant, and `~` denotes the bitwise complement operator. After execution of the instruction sequence, the value in `reg1` matches the allocation size of the VLA in the C procedure. For example, for a VLA declaration `int v[n]`, `reg1` would have value `n*4`². The value in `reg2` is the allocation amount after adjusting for alignment requirements, e.g., `v` (declared as `int v[n]`) would have an alignment of at least 4 in 32-bit x86. The last instruction shifts `esp reg2` bytes below its original value, thus *allocating* the VLA³.

At time of deallocation, the stackpointer register is simply incremented by the same value as it was decremented by during the allocation:

$$\begin{aligned} reg_1 &\leftarrow vlaSz \\ reg_2 &\leftarrow (reg_1 + C) \& \sim C \\ esp &\leftarrow esp + reg_2 \end{aligned}$$

²Recall that 4 is the size of an `int` in our 32-bit configuration

³`esp` must have been originally aligned by the required alignment for the resulting value to be correctly aligned.

Notice that the two assembly sequences are identical except for the last instruction, where in the deallocation sequence $+$ is used in place of $-$.

Recall that the execution of a stackpointer increment instruction may trigger \mathcal{U} if the deallocated interval does not belong to the stack region ((OP-ESP) in fig. 2.6). Thus, in order to prove that the assembly procedure does not trigger \mathcal{U} in the above deallocation sequence, we must have an invariant stating that the original stackpointer value (just before executing the above deallocation sequence), was at least reg_2 bytes (where reg_2 is derived from $vlaSz$) below some address in the stack region *and* the interval spanned by the two stackpointer values belongs entirely to the stack region. The global invariants $\boxed{\text{StkBd}}$, $\boxed{\text{zlIntvl}}$, and $\boxed{\text{zaBd}}$ (fig. 4.3) and, invariants $\boxed{\text{affine}}$ and $\boxed{\text{ineq}}$ (fig. 4.2) over the ghost variables for stackpointer ($\boxed{\text{sp.p}_{\ddot{A}}^j}$) and local regions ($\boxed{\text{lb.z}}$, $\boxed{\text{ub.z}}$, $\boxed{\text{lstSz.z}}$) help in proving separation of intervals defined by stackpointer ghost variables and stack-allocated local regions, and are usually sufficient to discharge the latter check. Proving the former condition, that the gap between two stackpointer values is larger than a value derived from $vlaSz$, however, requires an invariant that cannot be inferred by our predicate grammar.

Let us look at the allocation and deallocation instruction sequences again. At the end of the allocation sequence, the new stackpointer value is related to the old stackpointer value by ($\text{esp}_a = \text{esp}_b - reg_2$) where esp_b and esp_a are the stackpointer values just before and after executing the stackpointer decrement instruction respectively. Due to translation rule (OP-ESP) in fig. 2.6, both esp_b and esp_a will have ghost variables, say, $\boxed{\text{sp.p}_{\ddot{A}}^{j_1}}$ and $\boxed{\text{sp.p}_{\ddot{A}}^{j_2}}$, holding value equal to them. Thus, the relation between the two stackpointer values can be expressed as $\boxed{\text{sp.p}_{\ddot{A}}^{j_2}} = \boxed{\text{sp.p}_{\ddot{A}}^{j_1}} - reg_2$, an affine shape that is a member of $\boxed{\text{affine}}$. On the other hand, the relation $reg_2 = (vlaSz + C) \ \& \ \sim C$ is not an affine shape and cannot be captured by any other shape in our predicate grammar. Further, the alignment adjustment computation is not performed in C so there is no C variable that can be related to reg_2 through an affine relation. Consequently, instead of a precise affine invariant relating reg_2 and $vlaSz$, an imprecise inequality invariant of the form $\boxed{\text{sp.p}_{\ddot{A}}^{j_2}} \leq_u \boxed{\text{sp.p}_{\ddot{A}}^{j_1}} - vlaSz$, is inferred ($\boxed{\text{spOrd}}$ in fig. 4.2). This inferred invariant is sufficient for proving the well-formedness of the `alloc` instruction (so that C does not go to \mathcal{W}_C) and proving that \ddot{A} does not go to $\mathcal{U}_{\ddot{A}}$ in (ALLOCS).

In the deallocation sequence, instead of reusing reg_2 or resetting the stackpointer to older value (as happens in some compilations of GCC), the required value for stackpointer increment is recomputed again using $vlaSz$. At this point, the product graph invariants

do not imply a precise relationship between $vlaSz$ and the reg_2 . Thus, the stackpointer is incremented by a value which, according to the inferred invariants, is completely unrelated to the value by which stackpointer was decremented. Consequently, the attempt to prove that the deallocated interval belongs to the stack region fails and we are unable to prove that \ddot{A} does not go to $\mathcal{U}_{\ddot{A}}$ in the stackpointer increment instruction.

It can be observed that the necessary invariant required for the proof to go through is $(reg_2 = (vlaSz + C) \ \& \ \sim C)$, where $vlaSz$ can be substituted with the ghost variable for the size of the VLA ($\boxed{\text{1stSz.z}}$ for allocation site z of the VLA). This invariant shape, though required in this case, is rather specific and may not be useful otherwise. On the other hand, choosing to include it in the predicate grammar is highly likely to cause an increase in the runtime of the tool. We make the choice of omitting it, choosing to sacrifice completeness in favor of (relatively better) runtimes — our choice is informed by casual observations on scalability bottlenecks due to inference of non-affine invariant shapes.

Failure due to choice of program variables for invariant inference

In the $\boxed{\text{affine}} \sum_i c_i v_i = c$ invariant shape of the predicate grammar (fig. 4.2), the program variables v_i are drawn from a set V that includes the pseudo-registers in C and registers and stack slots in A (section 4.1.5). The candidate variables for correlation in V do not include “memory slots” of shape $\text{se1}_{sz}(M_C, \alpha)$ (little-endian concatenation of sz bytes starting at α in the array M_C) to avoid an explosion in the number of candidate invariants and, consequently, the running time of the algorithm.

This causes a failure while validating the GCC compilations (at 03) of the variadic `mp` and `ms` benchmarks in table 6.1. GCC register-allocates the `va_list` variable (that maintains the current position in the variadic argument). On the other hand, the LLVM_d IR maintains this pointer value in a local variable (allocated using an `alloc` instruction) — the loads and stores to this local variable $\langle ap \rangle$ can be seen in fig. 2.3. Thus, for a refinement proof to succeed, a validator must relate the assembly register’s value with the value stored inside the local variable’s memory region ($\text{se1}_4(M_C, \langle ap \rangle)$). Because our invariant inference algorithm does not consider memory slots in C , this required relation is not identified, resulting in a proof failure.

It may be worth asking the question: why does our choice of program variables work for the other benchmarks? Due to the `mem2reg` pass used in C before computing equivalence,

the only memory slots that remain in procedure `C` pertain to potentially address-taken variables. Our requirements on the product graph X ensure that the memory regions corresponding to address-taken local variables (and global variables) of `C` and `A` are equated in X . Thus, relating the *addresses* of potential memory accesses in `C` and `A` using affine invariants and considering only the memory slots from `A` largely suffices for invariant inference to validate most compilations (but not for GCC’s compilation of `mp` and `ms`).

6.3 Other Applications

A translation validator has applications beyond compiler validation. This section explores some applications of our tool to domains other than translation validation.

Applicability in a Superoptimizer

An end-to-end translation validation is a crucial ingredient in a synthesis-driven superoptimizer [5, 9, 38] which attempts to generate “optimal” code for a particular program (or specification). In a superoptimizer, a synthesis tool proposes candidate solutions and a translation validation tool validates them against a high-level input specification. The efficacy of superoptimization depends both on the synthesis tool, through its ability to generate effective solutions, and the translation validation tool, through its ability to validate the proposed solutions. With advances in synthesis, e.g. through Large Language Models (LLMs), the burden of effective superoptimization will inevitably shift towards translation validation where DYNAMO-like tools will find application.

Checking enforcement of calling conventions

Consider the case of alignment used for certain types. Compilers often use higher alignment factors than those necessitated by the C standard, e.g., the “`long long`” type is often aligned at eight-byte boundaries to reduce cache misses. Our tool can easily check use of such higher than required alignments by changing the well-formedness condition for alignment (section 2.1.3) to reflect the higher alignment value. By using `long long` type in our first set of benchmarks (containing different programming patterns) in section 6.2.1, we validated that all the three production compilers (Clang/LLVM, GCC, ICC) ensure that `long long` variables are eight-byte aligned for these benchmarks. In contrast, we found that the ACK compiler [47] only ensures the mandated four-byte

alignment.

Use in fuzzing

A translation validator can also be used as a part of compiler-fuzzing tools such as CSmith [52] and EMI [27]. Validation usually takes much longer than testing, yet provides greater (maximum) coverage. It remains to be seen if validation can be used in conjunction with fuzzing to uncover more compiler bugs. We confirmed that our validator is able to detect the previous bugs (involving local memory variables) reported by these compiler fuzzing tools. These bugs were hitherto not possible to uncover through validation due to the lack of support for local variable modeling in prior work.

Detecting performance anomaly

The validator also helped us identify a small performance anomaly in the code generated by a recent version of GCC. It turns out that in the presence of VLAs, GCC emits assembly code to redundantly align an already aligned pointer. This is easily checked by a lightweight static checker on the final product graph that checks if an alignment operation (characterized by the `and` opcode) occurs at an assembly PC where the product graph invariants already ensure that alignment.

Bug in SMT solver

Through our experiments, we uncovered and reported a bug in recent versions of `z3`, including `z3-4.8.14` and `z3-4.12.5`, where for an input satisfiability query Q , the SMT solver returns an unsound model (counterexample) that evaluates Q to false [53]. When a modern SMT solver is used to validate compilations produced by a mature compiler, a bug may be found on either side.

Chapter 7

Conclusion

The overarching goal of this work is to investigate the applicability of translation validation as an alternative to verified compilation. Toward this, we identify the modeling of dynamic allocation and deallocation of local memory as an important unaddressed sub-problem and make multiple contributions within this space.

7.1 Summary

We formalize the refinement from an unoptimized representation (IR) of a C procedure to a 32-bit x86 assembly procedure as a relation over the observable events (termination, procedure-call invocation, (de)allocation, etc.) produced by the execution of the respective procedures (section 2.4). A key aspect of this formalization is the observation of a (de)allocation in the IR procedure as a distinct (observable) event. This enables the identification of a similar (de)allocation event in the assembly procedure through a verifiable *annotation* instruction (sections 2.5.1 and 2.6.1). The identification of correlated events provides a basis for a *lockstep correlation*-based proof method, such as product program construction, where the execution of both procedures matches step-by-step, ensuring event correspondence. We demonstrate that, unlike prior work, this annotation-based modeling is not tied to a particular local (de)allocation strategy and is applicable in different settings, including validation of programs with dynamic variable-sized or constant-sized allocations, e.g., through C99 VLAs or the `alloca()` operator (sections 2.5.4 and 2.6.3).

We generalize the product program-based proof method, used in prior work for estab-

lishing equivalence, to a determinized product program that is applicable for correlating two programs in the presence of non-determinism and, thus, can cater to a generalized refinement setting (section 3.2). We identify the key requirements for witnessing refinement through a determinized product program and present an automatic algorithm for constructing the same (section 3.3).

Our automatic algorithm is based on prior heuristic-guided best-first search that constructs the required product program incrementally (section 4.1). The automatic algorithm supports external untrusted hints for guiding the construction and thus can make use of information from sources such as compiler instrumentation (sections 4.1.3 and 6.1.4). We provide an efficient SMT encoding that is up to 4 times faster than a naive encoding for discharging the generated proof obligations (section 5.3). We show that our prototype implementation is able to validate compilations by production optimizing compilers, Clang/LLVM, GCC, and ICC (section 6.2.1). We also evaluate our prototype on SPEC2000's `bzip2` program, where we are able to successfully validate an optimized compilation of procedures with up to 142 assembly lines of code, failing to validate just 8 procedures out of 72 (section 6.2.3).

7.2 Limitations and Directions for Future Work

While our work represents a significant step toward realizing a practical translation validator, it has several limitations. Although our execution model and refinement definition admit a broad range of transformations and allocation strategies, our (determinized) product program construction and SMT encoding operate under specific constraints. In particular, we assume:

- In allocations and procedure calls that reuse stack space, their relative order is preserved in both the original and compiled programs (section 2.5.4).
- An `alloca()` is always stack-allocated (section 2.6.1).
- The compiler does not specialize allocation-containing paths such that one specialization uses stack allocation and another uses register allocation (section 2.6.1).

While the generalized execution model and refinement definition presented in section 2.7 lift these constraints, they leave the challenge of efficient product program construction and SMT encoding to future work.

Additionally, our over-approximate modeling of procedure calls, which treats them as arbitrary mutations of callee-observable state, lacks the precision necessary for validating interprocedural transformations (section 2.5.4).

Our experiments on the `bzip2` benchmark highlight the need for a more robust and scalable search algorithm and/or SMT encoding. Addressing these scalability challenges presents a promising direction for future work toward a fully realized translation validator.

Appendices

Appendix A

Soundness and Completeness Implications of `isPush()` Choice

An update to the stackpointer `esp` in the assembly procedure `A` can be through any arbitrary instruction, such as `esp := Y`. If the previous `esp` value, just before this instruction was executed, was `X`, then the stackpointer update distance is $\mathbb{D} = \mathbb{X} - \mathbb{Y}$. In general, it is impossible to tell whether this instruction intends a stack growth by \mathbb{D} bytes (push) or a shrink by $(2^{32} - \mathbb{D})$ bytes (pop). The modeling for the two cases is different: for stack push, an overlap of the interval representing the push with non-stack region causes a \mathcal{W} error, while for stack pop, the stackpointer going outside stack region causes \mathcal{U} error. Refinement is trivially proven if `A` terminates with \mathcal{W} error. Unfortunately, this seems impossible to disambiguate just by looking at the assembly code – to tackle this dilemma, we assume an oracle function, `isPush($p_A^j, \mathbb{X}, \mathbb{Y}$)`, that returns `true` iff the assembly instruction at PC p_A^j represents a stack push.

In section 2.3.2, we define an `isPush($p_A^j, \mathbb{X}, \mathbb{Y}$)` operator for an assembly instruction at p_A^j based on thresholding of the update distance $\mathbb{D} = \mathbb{X} - \mathbb{Y}$ by a threshold value $\mathbb{K} = 2^{31} - 1$:

$$\text{isPush}(p_A^j, \mathbb{X}, \mathbb{Y}) \Leftrightarrow \mathbb{X} - \mathbb{Y} \leq_u \mathbb{K}$$

Here, \mathbb{K} represents the threshold value for the stack update distance $\mathbb{X} - \mathbb{Y}$, below which we consider the update to be a push.

If \mathbb{K} is smaller than required, then we risk misclassifying stack pushes (stack growth) as stack pops (stack shrink). On the other hand, if \mathbb{K} is bigger than required, then we

risk misclassifying stack pops (stack shrink) as stack pushes (stack growth). In the latter case (when \mathbb{K} is bigger than required), we would incorrectly trigger \mathcal{W} , instead of \mathcal{U} , and that would cause the refinement proof to complete incorrectly (soundness problem). In the extreme case, if $\mathbb{K} = 2^d - 1$ (where the address space has size 2^d), then even 4-byte stack pops (e.g., through the x86 `pop` instruction) would be considered as stack pushes (growth), and we would incorrectly trigger in every situation where \mathcal{U} was expected, and the refinement proof would complete trivially (and unsoundly).

On the other hand, if \mathbb{K} is smaller than required, we may incorrectly count some stack growth operations as stack pops. In these cases, we will have show to absence of \mathcal{U} (as part of (Safety)) for a stack pop for which a stack push never happened. This would result in an refinement failure (completeness problem).

A.1 \mathbb{K} needs to be at least 2^{d-1} in the presence of VLAs

Consider a VLA declaration, “`char v[n]`” in `C`. In this case, `n` could be any positive integer $\leq_u \text{INT_MAX}$; this upper bound of `INT_MAX` comes from the variable size limits imposed by the C language. The corresponding allocation statement in assembly code would be something like “ $p_A^j: \text{esp} := \text{esp} - n$ ”. The resulting condition for not triggering \mathcal{U} is (from (OP-ESP) of fig. 2.6):

$$\begin{aligned} & \neg(\neg \text{isPush}(p_A^j, \text{esp}, \text{esp} - n) \\ & \quad \wedge \text{esp} \neq \text{esp} - n \\ & \quad \wedge \neg \text{intrvlInSet}(\text{esp}, \text{esp} - n, \Sigma_A^{stk})) \end{aligned}$$

or equivalently,

$$\begin{aligned} (n >_u \mathbb{K}) \Rightarrow (& n = 0_{i_{32}} \\ & \vee (\text{esp} \neq 0_{i_{32}} \\ & \quad \wedge (\text{esp} \leq_u \text{esp} - n) \\ & \quad \wedge [\text{esp}, \text{esp} - n] \subseteq \Sigma_A^{stk})) \end{aligned} \tag{A.1}$$

Now, if \mathbb{K} is smaller than the biggest possible value of `n`, then there exist values of `n` where the left clause (left of \Rightarrow) of eq. (A.1) would evaluate to `true`. Consequently, there exist values of `n` for which the right clause has to be proven `true`, i.e., prove that

the stack region is at least $2^d - n$ bytes large. It may not be possible to prove such strong conditions in all cases and thus we get false refinement check failures. Because the C language constrains n to be $\leq_u \text{INT_MAX}(= 2^{d-1} - 1)$, $\mathbb{K} \geq_u 2^{d-1} - 1$ seems sufficient to be able to validate such translations.

However, $\mathbb{K} = 2^{d-1} - 1$ is also insufficient, because typically the code generated by a compiler for “`char v[n]`” also aligns n using a rounding factor $r = 2^i$: “`esp := esp - (\lceil \frac{n}{r} \rceil \cdot r)`”. In this scenario, even though $n \leq_u (2^{d-1} - 1)$, it is possible for $\mathbb{D} = \lceil \frac{n}{r} \rceil \cdot r$ to be greater than $(2^{d-1} - 1)$. Thus, if $\mathbb{K} = 2^{d-1} - 1$, there exist legal values of n for which stack region is at least $2^d - n$ bytes large has to be proven to demonstrate absence of \mathcal{U} . The choice $\mathbb{K} = 2^{d-1}$ allows for such alignment padding, and thus allows the refinement proof to be completed in these situations.

A.2 $\mathbb{K} = 2^{d-1}$ can still lead to completeness problems

If a single stack update allocates two VLAs at once, we can incorrectly classify a stack growth as a stack shrink.

Consider two C statements in sequence, “`char v1[m]; char v2[n];`”. In this case both m and n can individually be as large as $2^{d-1} - 1$. If the compiler decides to use a single assembly instruction to allocate both these variables, then it is possible for a single stack update distance \mathbb{D} to be greater than $\mathbb{K} = 2^{d-1}$. Thus, in these cases, the refinement proof may fail if we are not able to prove that stack is large enough to contain $2^d - \mathbb{D}$ bytes (for the classified stack pop). This is a completeness problem.

A.3 $\mathbb{K} = 2^{d-1}$ can also lead to soundness problems

If a single stack update deallocates two VLAs at once, we can incorrectly classify a stack shrink as a stack growth.

Consider two C statements in sequence, “`char v1[2d-1 - 1]; char v2[2];`”. If during deallocation, the compiler decides to use a single instruction to deallocate both the arrays, e.g., “`esp := esp + (2d-1 - 1) + 2`” for a total update distance of:

$$\mathbb{D} = -((2^{d-1} - 1) + 2) = 2^{d-1} - 1 \quad (\text{mod } 2^d)$$

Here, because $2^{d-1} - 1 \leq_u \mathbb{K}$ we will classify this “deallocation” as a stack push

(allocation) of $(2^{d-1} - 1)$ bytes and trigger \mathcal{W} if allocation of $(2^{d-1} - 1)$ bytes is not possible. This is a soundness problem because triggering \mathcal{W} under such a weaker condition may lead the refinement proof to succeed incorrectly.

A.4 Solution

Thus, it seems impossible in general to be able to distinguish a push from a pop in a sound manner. This problem is unavoidable in the presence of VLAs. CompCert side-stepped this problem by disabling VLA support and thus being able to statically bound the overall stack size. For a bounded stack, it becomes possible to distinguish pushes from pops. But it is not possible to bound the stack in the presence of a VLA.

Thus we propose that the compiler must explicitly emit trustworthy information that distinguishes a push from a pop. Hence, `isPush()` can simply leverage this information emitted by the compiler.

As explained in section 2.3.2, in our work, we use a threshold of $2^{31} - 1$ on the update distance to disambiguate stack pushes from pops. We rely on manual verification for soundness.

Appendix B

More details of the experiments

B.1 Command-line used for compiling benchmarks in experiments

1. Programs in table 6.1

- Clang/LLVM v12.0.0

```
clang -m32 -S -no-integrated-as -g -Wl,--emit-relocs -fdata-sections -g
  -fno-builtin -fno-strict-aliasing -fno-optimize-sibling-calls -
  fwrapv -fno-strict-overflow -ffreestanding -fno-jump-tables -fcf-
  protection=none -fno-stack-protector -fno-inline -fno-inline-
  functions -D_FORTIFY_SOURCE=0 -D__noreturn__=__no_reorder__ -I/usr/
  include/x86_64-linux-gnu/c++/9/32 -I/usr/include/x86_64-linux-gnu/c
  ++/9 -mllvm -enable-tail-merge=false -mllvm -nomerge-calls -std=c11
  -O3 <file.c> -o <file.s>
```

- GCC v8.4.0

```
gcc-8 -m32 -S -g -Wl,--emit-relocs -fdata-sections -g -no-pie -fno-pie -
  fno-strict-overflow -fno-unit-at-a-time -fno-strict-aliasing -fno-
  optimize-sibling-calls -fkeep-inline-functions -fwrapv -fno-reorder-
  blocks -fno-jump-tables -fno-caller-saves -fno-inline -fno-inline-
  functions -fno-inline-small-functions -fno-indirect-inlining -fno-
  partial-inlining -fno-inline-functions-called-once -fno-early-
  inlining -fno-whole-program -fno-ipa-sra -fno-ipa-cp -fcf-protection
```

```
=none -fno-stack-protector -fno-stack-clash-protection -
D_FORTIFY_SOURCE=0 -D__noreturn__=__no_reorder__ -fno-builtin-printf
-fno-builtin-malloc -fno-builtin-abort -fno-builtin-exit -fno-
builtin-fscanf -fno-builtin-abs -fno-builtin-acos -fno-builtin-asin
-fno-builtin-atan2 -fno-builtin-atan -fno-builtin-calloc -fno-
builtin-ceil -fno-builtin-cosh -fno-builtin-cos -fno-builtin-exp -
fno-builtin-fabs -fno-builtin-floor -fno-builtin-fmod -fno-builtin-
fprintf -fno-builtin-fputs -fno-builtin-frexp -fno-builtin-isalnum -
fno-builtin-isalpha -fno-builtin-iscntrl -fno-builtin-isdigit -fno-
builtin-isgraph -fno-builtin-islower -fno-builtin-isprint -fno-
builtin-ispunct -fno-builtin-isspace -fno-builtin-isupper -fno-
builtin-isxdigit -fno-builtin-tolower -fno-builtin-toupper -fno-
builtin-labs -fno-builtin-ldexp -fno-builtin-log10 -fno-builtin-log
-fno-builtin-memchr -fno-builtin-memcmp -fno-builtin-memcpy -fno-
builtin-memset -fno-builtin-modf -fno-builtin-pow -fno-builtin-
putchar -fno-builtin-puts -fno-builtin-scanf -fno-builtin-sinh -fno-
builtin-sin -fno-builtin-sprintf -fno-builtin-sprintf -fno-builtin-
sqrt -fno-builtin-sscanf -fno-builtin-strcat -fno-builtin-strchr -
fno-builtin-strcmp -fno-builtin-strcpy -fno-builtin-strcspn -fno-
builtin-strlen -fno-builtin-strncat -fno-builtin-strncmp -fno-
builtin-strncpy -fno-builtin-strpbrk -fno-builtin-strrchr -fno-
builtin-strspn -fno-builtin-strstr -fno-builtin-tanh -fno-builtin-
tan -fno-builtin-vfprintf -fno-builtin-vsprintf -fno-builtin -I/usr/
include/x86_64-linux-gnu/c++/9/32 -I/usr/include/x86_64-linux-gnu/c
++/9 -fno-tree-tail-merge --param max -tail-merge-comparisons=0 --
param max-tail-merge-iterations=0 -std=c11 -O3 <file.c> -o <file.s>
```

- ICC v2021.8.0

```
icc -m32 -D_Float32=__Float32 -D_Float64=__Float64 -D_Float32x=
__Float32x -D_Float64x=__Float64x -S -g -Wl,--emit-relocs -fdata-
sections -g -no-ip -fno-optimize-sibling-calls -fargument-alias -no-
ansi-alias -falias -fno-jump-tables -fno-omit-frame-pointer -fno-
strict-aliasing -fno-strict-overflow -fwrapv -fabi-version=1 -nolib-
inline -inline-level=0 -fno-inline-functions -finline-limit=0 -no-
inline-calloc -no-inline-factor=0 -fno-builtin-printf -fno-builtin-
malloc -fno-builtin-abort -fno-builtin-exit -fno-builtin-fscanf -fno-
builtin-abs -fno-builtin-acos -fno-builtin-asin -fno-builtin-atan2
```

```

-fno-builtin-atan -fno-builtin-calloc -fno-builtin-ceil -fno-builtin
-cosh -fno-builtin-cos -fno-builtin-exp -fno-builtin-fabs -fno-
builtin-floor -fno-builtin-fmod -fno-builtin-fprintf -fno-builtin-
fputs -fno-builtin-frexp -fno-builtin-isalnum -fno-builtin-isalpha -
fno-builtin-iscntrl -fno-builtin-isdigit -fno-builtin-isgraph -fno-
builtin-islower -fno-builtin-isprint -fno-builtin-ispunct -fno-
builtin-isspace -fno-builtin-isupper -fno-builtin-isxdigit -fno-
builtin-tolower -fno-builtin-toupper -fno-builtin-labs -fno-builtin-
ldexp -fno-builtin-log10 -fno-builtin-log -fno-builtin-memchr -fno-
builtin-memcmp -fno-builtin-memcpy -fno-builtin-memset -fno-builtin-
modf -fno-builtin-pow -fno-builtin-putchar -fno-builtin-puts -fno-
builtin-scanf -fno-builtin-sinh -fno-builtin-sin -fno-builtin-
snprintf -fno-builtin-sprintf -fno-builtin-sqrt -fno-builtin-sscanf
-fno-builtin-strcat -fno-builtin-strchr -fno-builtin-strcmp -fno-
builtin-strcpy -fno-builtin-strcspn -fno-builtin-strlen -fno-builtin
-strncat -fno-builtin-strncmp -fno-builtin-strncpy -fno-builtin-
strpbrk -fno-builtin-strrchr -fno-builtin-strspn -fno-builtin-strstr
-fno-builtin-tanh -fno-builtin-tan -fno-builtin-vfprintf -fno-
builtin-vsprintf -fno-builtin -D_FORTIFY_SOURCE=0 -D__noreturn__=
__no_reorder__ -qno-opt-multi-version-aggressive -ffreestanding -
unroll0 -no-vec -I/usr/include/x86_64-linux-gnu/c++/9/32 -I/usr/
include/x86_64-linux-gnu/c++/9 -std=c11 -O3 <file.c> -o <file.s>

```

2. TSVC

```

clang -m32 -S -no-integrated-as -g -Wl,--emit-relocs -fdata-sections -g -
fno-builtin -fno-strict-aliasing -fno-optimize-sibling-calls -fwrapv -
fno-strict-overflow -ffreestanding -fno-jump-tables -fcf-protection=
none -fno-stack-protector -fno-inline -fno-inline-functions -
D_FORTIFY_SOURCE=0 -D__noreturn__=__no_reorder__ -I/usr/include/x86_64-
linux-gnu/c++/9/32 -I/usr/include/x86_64-linux-gnu/c++/9 -msse4.2 -
mllvm -enable-tail-merge=false -mllvm -nomerge-calls -std=c11 -O3
<file.c> -o <file.s>

```

3. bzip2 01-

```

clang -m32 -S -g -Wl,--emit-relocs -fno-unroll-loops -fdata-sections -fno-
inline -fno-inline-functions -fcf-protection=none -fno-stack-protector

```

```

-mllvm -enable-tail-merge=false -O1 -mllvm -nomerge-calls -mllvm -no-
early-cse -mllvm -no-licm -mllvm -no-machine-licm -mllvm -no-dead-arg-
elim -mllvm -no-ip-sparse-conditional-constant-prop -mllvm -no-dce-
fcalls -mllvm -replexitval=never -std=c11 -fno-builtin -fno-strict-
aliasing -fno-optimize-sibling-calls -fwrapv -fno-strict-overflow -
ffreestanding -fno-jump-tables -D_FORTIFY_SOURCE=0 -D__noreturn__=
__no_reorder__ -fno-builtin-printf -fno-builtin-malloc -fno-builtin-
abort -fno-builtin-exit -fno-builtin-fscanf -fno-builtin-abs -fno-
builtin-acos -fno-builtin-asin -fno-builtin-atan2 -fno-builtin-atan -
fno-builtin-calloc -fno-builtin-ceil -fno-builtin-cosh -fno-builtin-
cos -fno-builtin-exp -fno-builtin-fabs -fno-builtin-floor -fno-builtin-
fmod -fno-builtin-fprintf -fno-builtin-fputs -fno-builtin-frexp -fno-
builtin-isalnum -fno-builtin-isalpha -fno-builtin-iscntrl -fno-builtin-
isdigit -fno-builtin-isgraph -fno-builtin-islower -fno-builtin-isprint
-fno-builtin-ispunct -fno-builtin-isspace -fno-builtin-isupper -fno-
builtin-isxdigit -fno-builtin-tolower -fno-builtin-toupper -fno-
builtin-labs -fno-builtin-ldexp -fno-builtin-log10 -fno-builtin-log -
fno-builtin-memchr -fno-builtin-memcmp -fno-builtin-memcpy -fno-
builtin-memset -fno-builtin-modf -fno-builtin-pow -fno-builtin-putchar
-fno-builtin-puts -fno-builtin-scanf -fno-builtin-sinh -fno-builtin-
sin -fno-builtin-snprintf -fno-builtin-sprintf -fno-builtin-sqrt -fno-
builtin-sscanf -fno-builtin-strcat -fno-builtin-strchr -fno-builtin-
strcmp -fno-builtin-strcpy -fno-builtin-strcspn -fno-builtin-strlen -
fno-builtin-strncat -fno-builtin-strncmp -fno-builtin-strncpy -fno-
builtin-strpbrk -fno-builtin-strrchr -fno-builtin-strspn -fno-builtin-
strstr -fno-builtin-tanh -fno-builtin-tan -fno-builtin-vfprintf -fno-
builtin-vsprintf -fno-builtin -I/usr/include/x86_64-linux-gnu/c++/9/32
-I/usr/include/x86_64-linux-gnu/c++/9 bzip2.c -o bzip2.s

```

4. bzip2 01

```

clang -m32 -S -no-integrated-as -g -Wl,--emit-relocs -fdata-sections -g -
fno-builtin -fno-strict-aliasing -fno-optimize-sibling-calls -fwrapv -
fno-strict-overflow -ffreestanding -fno-jump-tables -fcf-protection=
none -fno-stack-protector -fno-inline -fno-inline-functions -
D_FORTIFY_SOURCE=0 -D__noreturn__=__no_reorder__ -I/usr/include/x86_64-
linux-gnu/c++/9/32 -I/usr/include/x86_64-linux-gnu/c++/9 -fno-unroll-
loops -mllvm -enable-tail-merge=false -mllvm -nomerge-calls -std=c11 -

```

```
01 bzip2.c -o bzip2.s
```

5. bzip2 02

```
clang -m32 -S -no-integrated-as -g -Wl,--emit-relocs -fdata-sections -g -
  fno-builtin -fno-strict-aliasing -fno-optimize-sibling-calls -fwrapv -
  fno-strict-overflow -ffreestanding -fno-jump-tables -fcf-protection=
  none -fno-stack-protector -fno-inline -fno-inline-functions -
  D_FORTIFY_SOURCE=0 -D__noreturn__=__no_reorder__ -I/usr/include/x86_64-
  linux-gnu/c++/9/32 -I/usr/include/x86_64-linux-gnu/c++/9 -fno-unroll-
  loops -mllvm -enable-tail-merge=false -mllvm -nomerge-calls -std=c11 -
02 bzip2.c -o bzip2.s
```

B.2 Full results for the bzip2 experiment

Table B.1 shows the full list of `bzip2` procedures with their assembly lines of code (ALOC) and validation times (in seconds) for the three Clang/LLVM compiler configurations (01-, 01, 02).

Name	ALOC			Validation time (seconds)		
	01-	01	02	01-	01	02
<code>allocateCompressStructures</code>	47	47	51	43.2	47.2	50.6
<code>badBGLengths</code>	13	13	13	23.2	25.3	28.5
<code>badBlockHeader</code>	13	13	13	21.9	23.0	27.3
<code>bitStreamEOF</code>	13	13	13	22.7	21.2	26.1
<code>blockOverrun</code>	13	13	13	23.5	25.3	27.2
<code>bsFinishedWithStream</code>	22	22	25	24.0	23.2	26.0
<code>bsGetInt32</code>	4	4	4	6.0	6.0	7.2
<code>bsGetIntVS</code>	6	6	6	7.3	8.1	9.8
<code>bsGetUChar</code>	5	5	5	6.0	7.4	7.1
<code>bsGetUInt32</code>	24	24	24	13.3	16.9	20.6
<code>bsPutInt32</code>	6	6	6	6.6	7.6	9.0
<code>bsPutIntVS</code>	6	6	6	9.6	9.8	11.1
<code>bsPutUChar</code>	8	8	8	7.8	8.4	10.9
<code>bsPutUInt32</code>	32	32	32	30.8	30.7	36.4
<code>bsR</code>	46	46	46	42.8	42.8	51.0
<code>bsSetStream</code>	9	9	9	3.7	3.7	4.7

bsW	31	32	36	34.4	33.8	40.0
cadvise	6	6	6	20.3	20.2	26.9
cleanUpAndFail	48	46	46	187.9	179.0	227.8
compressOutOfMemory	14	14	14	33.6	33.0	42.6
compressStream	124	124	124	342.0	369.2	402.6
compressedStreamEOF	16	16	16	21.5	24.0	27.6
crcError	15	15	15	31.4	30.9	36.6
debug_time	2	2	2	1.6	1.8	2.0
doReversibleTransformation	48	49	47	93.3	102.9	129.2
fullGtU	120	113	113	363.0	375.4	404.4
generateMTFValues	144	144	166	1909.3	10441.4	✗
getAndMoveToFrontDecode	299	296	305	✗	✗	✗
getFinalCRC	3	3	3	1.9	2.2	2.3
getGlobalCRC	2	2	2	2.1	1.8	2.2
getRLEpair	72	73	73	144.6	✗	✗
hbAssignCodes	37	37	37	296.4	325.4	330.7
hbCreateDecodeTables	94	94	107	1610.3	1622.3	✗
hbMakeCodeLengths	261	249	292	✗	✗	✗
indexIntoF	23	23	23	30.6	32.0	41.2
initialiseCRC	2	2	2	2.3	1.9	2.3
ioError	15	15	15	17.9	18.3	23.1
loadAndRLEsource	96	96	96	336.2	366.7	✗
main	190	132	183	✗	✗	✗
makeMaps	16	16	16	14.5	15.8	17.9
med3	14	14	14	3.8	4.1	4.2
moveToFrontCodeAndSend	9	9	9	15.1	16.0	15.5
mySIGSEGVorSIGBUScatcher	35	23	23	178.3	✗	✗
mySignalCatcher	10	10	10	16.9	18.9	25.2
panic	13	13	13	32.3	36.1	30.3
qSort3	297	297	363	✗	✗	✗
randomiseBlock	35	37	38	155.1	177.9	✗
recvDecodingTables	199	193	295	2539.8	2690.8	✗
sendMTFValues	691	692	832	✗	✗	✗
setDecompressStructureSizes	79	79	81	426.1	351.8	345.0
setGlobalCRC	3	3	3	2.8	3.0	3.1
showFileNames	8	8	8	15.6	18.4	17.0
simpleSort	194	185	215	✗	✗	✗
sortIt	409	406	421	✗	✗	✗
spec_compress	11	11	11	16.0	16.0	16.0
spec_getc	29	29	29	40.6	43.7	46.6
spec_init	48	49	49	120.4	134.1	123.7
spec_initbufs	9	9	9	11.0	9.6	11.0
spec_load	110	105	105	512.4	499.8	524.1

spec_putc	29	29	29	52.5	51.2	57.3
spec_read	44	46	46	133.5	131.2	166.7
spec_reset	16	16	16	21.8	20.1	23.4
spec_rewind	5	5	5	3.4	3.3	3.5
spec_uncompress	10	10	10	15.0	16.4	14.3
spec_ungetc	45	48	48	176.1	188.2	183.7
spec_write	34	34	34	73.3	77.0	73.8
testStream	195	194	196	1619.5	✗	✗
uncompressOutOfMemory	14	14	14	48.6	50.5	45.8
uncompressStream	169	174	176	1010.5	✗	✗
undoReversibleTransformation_fast	221	223	248	1794.0	1836.8	✗
undoReversibleTransformation_small	273	271	281	✗	✗	✗
vswap	27	27	27	63.3	61.1	54.0

Table B.1: List of `bzip2` procedures with their assembly lines of code (ALOC) and validation times (in seconds) for the three Clang/LLVM compiler configurations (O1-, O1, O2). ✗ denotes validation failure for that procedure-compiler pair.

Appendix C

Full source code of the benchmarks

We provide the full source code of the benchmarks from table 6.1 in figs. C.1 to C.4 below (the source code for `fib` is already listed in fig. 2.1a).

The loops of validated `bzip2` benchmarks are shown in fig. C.5.

```

// substitute  $\mathcal{N}$  with 1, 2, 3
// to obtain vsl1, vsl2, vsl3
int vsl $\mathcal{N}$ (int n)
{
    if (n <= 0)
        return 0;
    int v1[n], ..., v $\mathcal{N}$ [n];
    for (int i = 0; i < n; ++i) {
        v1[i] = F1(a[i]);
        ...
        v $\mathcal{N}$ [i] = F $\mathcal{N}$ (a[i]);
    }
    return foo $\mathcal{N}$ (v1, ...,v $\mathcal{N}$ );
}

int vcu(int n, int k)
{
    int a[n];
    if (k > 0 && k <= n) {
        a[0] = 0;
        a[k-1] = 10;
        return a[0];
    }
    return 0;
}

int vilcc(int n)
{
    int ret = 0;
    int i = 1;
    while (i < n) {
        char t[i];
        if (init(t, i) < 0)
            continue;
        ret += t[i-1];
        ++i;
    }
    return ret;
}

// substitute  $\mathcal{N}$  with 1, 2, 3
// to obtain vil1, vil2, vil3
int vil $\mathcal{N}$ (unsigned n)
{
    int r = 0;
    for (unsigned i = 1; i < n; ++i) {
        int v1[4*i], ... v $\mathcal{N}$ [4*i];
        r += foo $\mathcal{N}$ (v1, ...,v $\mathcal{N}$ , i);
    }
    return r;
}

int vilce(int n)
{
    int ret = 0;
    int i = 1;
    while (i < n) {
        char t[i];
        if (init(t, i) < 0)
            break;
        ret += t[i-1];
        ++i;
    }
    return ret;
}

```

Figure C.1: Benchmarks with VLAs.

```

#include <alloca.h>
int as(int n)
{
    if (n < 1) {
        return 0;
    }
    int* p = alloca(n*sizeof(int));
    for (int i = 0; i<n; ++i) {
        p[i] = i*i;
    }
    return p[0]*p[n-1];
}

int ams(int n)
{
    if (n <= 0)
        return -1;
    int* p;
    if (n < 4096) {
        p = alloca(sizeof(int)*n);
    } else {
        p = malloc(sizeof(int)*n);
        if (!p) return -1;
    }
    foo(p);
    int ret = p[0]+p[n/2]+p[n-1];
    if (!(n < 4096))
        free(p);
    return ret;
}

int ac(char* s, int fd, int* a)
{
    int n;
    if (!s || (n = strlen(s)) <= 0)
        return 0;
    if (!a) {
        a = alloca(sizeof(int)*n);
    }
    for (int i = 0; i < n; ++i) {
        a[i] = s[i] + 32;
    }
    return write(fd, a, n);
}

#include <alloca.h>
int n;
int all()
{
    typedef struct lln {
        int data;
        struct lln* next;
    } Node;
    if (n > 4096)
        return 0;
    Node* hd = 0;
    for (int i = 0; i < n; ++i) {
        Node* t = alloca(sizeof(Node));
        t->data = next_data();
        t->next = hd; hd = t;
    }
    Node* t = hd;
    int ret = 0;
    while (t != 0) {
        ret += t->data;
        t = t->next;
    }
    return ret;
}

```

Figure C.2: Benchmarks with use of alloca

```
const int cts[] = { 0x66, 0x65, 0x67, 0x60 };
int rod(int n)
{
    char zz[] = "0123456789";
    printf("Scanning_%d_chars", n);
    char t[n];
    scanf("%s",t);
    int ret = 0;
    for (int i = 0, j = 0; i < n; ++i) {
        printf("Round#...\n", i);
        zz[j] ^= t[i];
        if (++j >= sizeof zz) j = 0;
    }
    ret += zz[0] + cts[n%((sizeof cts)/sizeof(cts[0]))];
    printf("Returning_%d", ret);
    return ret;
}
```

Figure C.3: Benchmark rod with mixed use of VLA and address-taken variable.

```

#include <stdarg.h>
void mp(char *fmt, ...)
{
    va_list ap;
    char *p, *sval;
    int ival;

    va_start(ap, fmt);
    for (p = fmt; *p; p++) {
        check(p);
        if (*p != '%') {
            putchar(*p);
            continue;
        }
        switch (*++p) {
            case 'd':
                ival = va_arg(ap, int);
                print_int(ival);
                break;
            case 's':
                for (sval = va_arg(ap, char*);
                    *sval; sval++)
                    putchar(*sval);
                break;
            default:
                break;
        }
    }
    va_end(ap);
}

#include <stdarg.h>
int ms(char* fmt, ...)
{
    va_list ap;
    char *p;
    int ret = 0;

    va_start(ap, fmt);
    for (p = fmt; *p; ++p) {
        DBG();
        if (*p != '%') {
            if (!is_blank(*p)) {
                if (!match_char(*p))
                    break;
            }
            continue;
        }
        switch (*++p) {
            case 'd' :
                *va_arg(ap, int*) = read_int();
                ++ret;
                break;
            case 's' :
                read_string(va_arg(ap, char*));
                ++ret;
                break;
            case '%' :
                if (!match_char('%'))
                    goto end;
                break;
            default:
                goto end;
        }
    }
    end:
    va_end(ap);
    return ret;
}

```

Figure C.4: Benchmarks mp and ms with variable argument list. mp is adapted from minprintf of K&R[24]

```

void recvDecodingTables() {
    unsigned char inUse16[16];
    for (...) { /* write:inUse16 ... */ }
    for (...) { /* ... */ }
    for (...) { /* read:inUse16 ...*/
        for (...) { /* ... */ }
    }
    for (...) { while (...) { /* ... */ } }
    { unsigned char pos[6];
      for (...) { /* write:pos ... */ }
      for (...) { /* read,write:pos ... */
        while (...) { /* ... */ }
      }
    }
    for (...) {
      for (...) {
        while (...) { /* ... */ }
      }
    }
    for (...) { for (...) { /* ... */ } }
}

void generateMTFValues() {
    unsigned char yy[256];
    for (...) { /* ... */ }
    for (...) { /* write:yy ... */ }
    for (...) { /* read,write:yy ... */
        while (...) { /* ... */ }
        while (...) { /* ... */ }
    }
    while (...) { /* ... */ }
}

void undoReversibleTransformation_fast() {
    int cftab[257];
    for (...) { /* write:cftab ... */ }
    for (...) { /* read,write:cftab ... */ }
    for (...) { /* read,write:cftab ... */ }
    if (...) { while (...) for (...) { /* ... */ } }
    else { while (...) for (...) { /* ... */ } }
}

```

(a) Loops in `recvDecodingTables()` (b) Loops in `generateMTFValues()`

(c) Loops in `undoReversibleTransformation_fast()`

Figure C.5: Structure of bzip2's functions

List of Publications

- [1] Shubhani Gupta, Abhishek Rose, and Sorav Bansal. “Counterexample-Guided Correlation Algorithm for Translation Validation”. In: *Proc. ACM Program. Lang.* 4.OOPSLA (Nov. 2020). DOI: [10.1145/3428289](https://doi.org/10.1145/3428289). URL: <https://doi.org/10.1145/3428289>.
- [2] Vaibhav Kiran Kurhe, Pratik Karia, Shubhani Gupta, Abhishek Rose, and Sorav Bansal. “Automatic Generation of Debug Headers through BlackBox Equivalence Checking”. In: *2022 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*. 2022, pp. 144–154. DOI: [10.1109/CGO53902.2022.9741273](https://doi.org/10.1109/CGO53902.2022.9741273).
- [3] Abhishek Rose and Sorav Bansal. “Modeling Dynamic (De)Allocations of Local Memory for Translation Validation”. In: *Proc. ACM Program. Lang.* 8.OOPSLA1 (Apr. 2024). DOI: [10.1145/3649863](https://doi.org/10.1145/3649863). URL: <https://doi.org/10.1145/3649863>.

Biography

Abhishek Rose is a PhD student in the Department of Computer Science and Engineering at IIT Delhi. He obtained his B. Tech. from MSIT, GGSIPU and his M. Tech. from IIT Kanpur.

Bibliography

- [1] ‘*alloca*’ Instruction. *LLVM Language Reference Manual*. URL: <https://llvm.org/docs/LangRef.html#alloca-instruction>.
- [2] *alloca(3) Linux Programmer’s Manual*. URL: <https://man7.org/linux/man-pages/man3/alloca.3.html>.
- [3] Lars Ole Andersen. *Program Analysis and Specialization for the C Programming Language*. Tech. rep. 1994.
- [4] Sorav Bansal. “Peephole Superoptimization”. PhD thesis. Stanford University, 2008. URL: <https://sorav.compiler.ai/pubs/thesis.pdf>.
- [5] Sorav Bansal and Alex Aiken. “Automatic Generation of Peephole Superoptimizers”. In: *Proceedings of the 12th International Conference on Architectural Support for Programming Languages and Operating Systems*. ASPLOS XII. San Jose, California, USA: ACM, 2006, pp. 394–403. ISBN: 1-59593-451-0. DOI: [10.1145/1168857.1168906](https://doi.org/10.1145/1168857.1168906). URL: <http://doi.acm.org/10.1145/1168857.1168906>.
- [6] Clark Barrett, Yi Fang, Benjamin Goldberg, Ying Hu, Amir Pnueli, and Lenore Zuck. “TVOC: A Translation Validator for Optimizing Compilers”. In: *Computer Aided Verification*. Ed. by Kousha Etessami and Sriram K. Rajamani. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 291–295. ISBN: 978-3-540-31686-2.

- [7] Clark Barrett, Pascal Fontaine, and Cesare Tinelli. *The Satisfiability Modulo Theories Library (SMT-LIB)*. <https://www.smt-lib.org>. 2016.
- [8] Berkeley Churchill, Oded Padon, Rahul Sharma, and Alex Aiken. “Semantic Program Alignment for Equivalence Checking”. In: *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*. PLDI 2019. Phoenix, AZ, USA: ACM, 2019, pp. 1027–1040. ISBN: 978-1-4503-6712-7. DOI: [10.1145/3314221.3314596](https://doi.org/10.1145/3314221.3314596). URL: <http://doi.acm.org/10.1145/3314221.3314596>.
- [9] Berkeley Churchill, Rahul Sharma, JF Bastien, and Alex Aiken. “Sound Loop Superoptimization for Google Native Client”. In: *Proceedings of the Twenty-Second International Conference on Architectural Support for Programming Languages and Operating Systems*. ASPLOS ’17. Xi’an, China: ACM, 2017, pp. 313–326. ISBN: 978-1-4503-4465-4. DOI: [10.1145/3037697.3037754](https://doi.org/10.1145/3037697.3037754).
- [10] *Clang C Language Family Frontend for LLVM*. URL: <https://clang.llvm.org/>.
- [11] Keith Cooper. “Live Range Splitting in a Graph Coloring Register Allocator”. In: (Jan. 1998). DOI: [10.1007/BFb0026430](https://doi.org/10.1007/BFb0026430).
- [12] Ron Cytron, Jeanne Ferrante, Barry K. Rosen, Mark N. Wegman, and F. Kenneth Zadeck. “Efficiently computing static single assignment form and the control dependence graph”. In: *ACM Trans. Program. Lang. Syst.* 13.4 (Oct. 1991), pp. 451–490. ISSN: 0164-0925. DOI: [10.1145/115372.115320](https://doi.org/10.1145/115372.115320). URL: <https://doi.org/10.1145/115372.115320>.
- [13] Manjeet Dahiya. “Black-box Equivalence Checking across Compiler Transformations”. PhD thesis. Indian Institute of Technology Delhi, 2018. URL: https://sorav.compiler.ai/pubs/manjeet_thesis.pdf.
- [14] Edsger W. Dijkstra. *A Discipline of Programming*. Prentice-Hall, 1976. ISBN: 013215871X.

- [15] *GCC, the GNU Compiler Collection*. URL: [%5Curl%7Bhttps://gcc.gnu.org/%7D](https://gcc.gnu.org/).
- [16] Dick Grune, Kees van Reeuwijk, Henri E. Bal, Cerie J.H. Jacobs, and Koen Langendoen. *Modern Compiler Design*. 2nd. Springer Publishing Company, Incorporated, 2012. ISBN: 1461446988.
- [17] Shubhani Gupta, Abhishek Rose, and Sorav Bansal. “Counterexample-Guided Correlation Algorithm for Translation Validation”. In: *Proc. ACM Program. Lang.* 4.OOPSLA (Nov. 2020). DOI: [10.1145/3428289](https://doi.org/10.1145/3428289). URL: <https://doi.org/10.1145/3428289>.
- [18] Shubhani Gupta, Aseem Saxena, Anmol Mahajan, and Sorav Bansal. “Effective Use of SMT Solvers for Program Equivalence Checking Through Invariant-Sketching and Query-Decomposition”. In: *Theory and Applications of Satisfiability Testing – SAT 2018*. Ed. by Olaf Beyersdorff and Christoph M. Wintersteiger. Cham: Springer International Publishing, 2018, pp. 365–382. ISBN: 978-3-319-94144-8.
- [19] Chris Hathhorn, Chucky Ellison, and Grigore Roşu. “Defining the undefinedness of C”. In: *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation*. PLDI ’15. Portland, OR, USA: Association for Computing Machinery, 2015, pp. 336–345. ISBN: 9781450334686. DOI: [10.1145/2737924.2737979](https://doi.org/10.1145/2737924.2737979). URL: <https://doi.org/10.1145/2737924.2737979>.
- [20] John L. Henning. “SPEC CPU2000: Measuring CPU performance in the new millenium”. In: *IEEE Computer* 33.7 (July 2000), pp. 28–35.
- [21] ISO. *ISO/IEC 9899:2011 Information technology — Programming languages — C*. Geneva, Switzerland: International Organization for Standardization, Dec. 2011, 683 (est.) URL: http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=57853.

- [22] Jeehoon Kang, Yoonseung Kim, Youngju Song, Juneyoung Lee, Sanghoon Park, Mark Dongyeon Shin, Yonghyun Kim, Sungkeun Cho, Joonwon Choi, Chung-Kil Hur, and Kwangkeun Yi. “Crellvm: Verified Credible Compilation for LLVM”. In: *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation*. PLDI 2018. Philadelphia, PA, USA: ACM, 2018, pp. 631–645. ISBN: 978-1-4503-5698-5. DOI: [10.1145/3192366.3192377](https://doi.org/10.1145/3192366.3192377). URL: <http://doi.acm.org/10.1145/3192366.3192377>.
- [23] Theodoros Kasampalis, Daejun Park, Zhengyao Lin, Vikram S. Adve, and Grigore Roşu. “Language-Parametric Compiler Validation with Application to LLVM”. In: *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*. ASPLOS 2021. Virtual, USA: Association for Computing Machinery, 2021, pp. 1004–1019. ISBN: 9781450383172.
- [24] Brian W. Kernighan and Dennis M. Ritchie. *The C Programming Language*. 2nd. Prentice Hall Professional Technical Reference, 1988. ISBN: 0131103709.
- [25] Ramana Kumar, Magnus O. Myreen, Michael Norrish, and Scott Owens. “CakeML: A Verified Implementation of ML”. In: *Principles of Programming Languages (POPL)*. ACM Press, Jan. 2014, pp. 179–191. DOI: [10.1145/2535838.2535841](https://doi.org/10.1145/2535838.2535841). URL: <https://cakeml.org/pop14.pdf>.
- [26] Vaibhav Kiran Kurhe, Pratik Karia, Shubhani Gupta, Abhishek Rose, and Sorav Bansal. “Automatic Generation of Debug Headers through BlackBox Equivalence Checking”. In: *2022 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*. 2022, pp. 144–154. DOI: [10.1109/CGO53902.2022.9741273](https://doi.org/10.1109/CGO53902.2022.9741273).
- [27] Vu Le, Mehrdad Afshari, and Zhendong Su. “Compiler Validation via Equivalence Modulo Inputs”. In: *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation*. PLDI ’14. Edinburgh, United

- Kingdom: ACM, 2014, pp. 216–226. ISBN: 978-1-4503-2784-8. DOI: [10.1145/2594291.2594334](https://doi.org/10.1145/2594291.2594334). URL: <http://doi.acm.org/10.1145/2594291.2594334>.
- [28] Juneyoung Lee, Dongjoo Kim, Chung-Kil Hur, and Nuno P. Lopes. “An SMT Encoding of LLVM’s Memory Model for Bounded Translation Validation”. In: *Computer Aided Verification*. Ed. by Alexandra Silva and K. Rustan M. Leino. Cham: Springer International Publishing, 2021, pp. 752–776. ISBN: 978-3-030-81688-9.
- [29] Xavier Leroy. “Formal certification of a compiler back-end, or: programming a compiler with a proof assistant”. In: *33rd ACM symposium on Principles of Programming Languages*. ACM Press, 2006, pp. 42–54. URL: <http://gallium.inria.fr/~xleroy/publi/compiler-certif.pdf>.
- [30] Xavier Leroy, Sandrine Blazy, Daniel Kästner, Bernhard Schommer, Markus Pister, and Christian Ferdinand. “CompCert – A Formally Verified Optimizing Compiler”. In: *ERTS 2016: Embedded Real Time Software and Systems*. SEE, 2016. URL: http://xavierleroy.org/publi/erts2016_compcert.pdf.
- [31] Nuno P. Lopes, Juneyoung Lee, Chung-Kil Hur, Zhengyang Liu, and John Regehr. “Alive2: Bounded Translation Validation for LLVM”. In: *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. PLDI 2021. Virtual, Canada: Association for Computing Machinery, 2021, pp. 65–79. ISBN: 9781450383912. DOI: [10.1145/3453483.3454030](https://doi.org/10.1145/3453483.3454030). URL: <https://doi.org/10.1145/3453483.3454030>.
- [32] H.J. Lu, David L Krietzer, Milind Girkar, and Zia Ansari. *System V Application Binary Interface. Intel386 Architecture Processor Supplement. Version 1.0*. <https://uclibc.org/docs/psABI-i386.pdf>. Feb. 2015.
- [33] Saeed Maleki, Yaoqing Gao, Maria J. Garzarán, Tommy Wong, and David A. Padua. “An Evaluation of Vectorizing Compilers”. In: *Proceedings of the 2011*

- International Conference on Parallel Architectures and Compilation Techniques*. PACT '11. Washington, DC, USA: IEEE Computer Society, 2011, pp. 372–382. ISBN: 978-0-7695-4566-0. DOI: [10.1109/PACT.2011.68](https://doi.org/10.1109/PACT.2011.68). URL: <https://doi.org/10.1109/PACT.2011.68>.
- [34] David Menendez, Santosh Nagarakatte, and Aarti Gupta. “Alive-FP: Automated Verification of Floating Point Based Peephole Optimizations in LLVM”. In: Sept. 2016, pp. 317–337. ISBN: 978-3-662-53412-0. DOI: [10.1007/978-3-662-53413-7_16](https://doi.org/10.1007/978-3-662-53413-7_16).
- [35] KedarS. Namjoshi and LenoreD. Zuck. “Witnessing Program Transformations”. English. In: *Static Analysis*. Ed. by Francesco Logozzo and Manuel Fähndrich. Vol. 7935. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2013, pp. 304–323. ISBN: 978-3-642-38855-2. DOI: [10.1007/978-3-642-38856-9_17](https://doi.org/10.1007/978-3-642-38856-9_17). URL: http://dx.doi.org/10.1007/978-3-642-38856-9_17.
- [36] George C. Necula. “Translation Validation for an Optimizing Compiler”. In: *Proceedings of the ACM SIGPLAN 2000 Conference on Programming Language Design and Implementation*. PLDI '00. Vancouver, British Columbia, Canada: ACM, 2000, pp. 83–94. ISBN: 1-58113-199-2. DOI: [10.1145/349299.349314](https://doi.org/10.1145/349299.349314). URL: <http://doi.acm.org/10.1145/349299.349314>.
- [37] Amir Pnueli, Michael Siegel, and Eli Singerman. “Translation Validation”. In: *Proceedings of the 4th International Conference on Tools and Algorithms for Construction and Analysis of Systems*. TACAS '98. London, UK, UK: Springer-Verlag, 1998, pp. 151–166. ISBN: 3-540-64356-7. URL: <http://dl.acm.org/citation.cfm?id=646482.691453>.
- [38] Raimondas Sasnauskas, Yang Chen, Peter Collingbourne, Jeroen Ketema, Jubi Taneja, and John Regehr. “Souper: A Synthesizing Superoptimizer”. In: *CoRR*

- abs/1711.04422 (2017). arXiv: [1711.04422](https://arxiv.org/abs/1711.04422). URL: <http://arxiv.org/abs/1711.04422>.
- [39] Thomas Sewell. “Translation Validation for Verified, Efficient and Timely Operating Systems”. PhD thesis. Sydney, Australia: UNSW, July 2017. URL: <https://trustworthy.systems/publications/papers/Sewell%3Aphd.pdf>.
- [40] Thomas Arthur Leck Sewell, Magnus O. Myreen, and Gerwin Klein. “Translation Validation for a Verified OS Kernel”. In: *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation*. PLDI ’13. Seattle, Washington, USA: Association for Computing Machinery, 2013, pp. 471–482. ISBN: 9781450320146. DOI: [10.1145/2491956.2462183](https://doi.org/10.1145/2491956.2462183). URL: <https://doi.org/10.1145/2491956.2462183>.
- [41] Rahul Sharma, Eric Schkufza, Berkeley Churchill, and Alex Aiken. “Data-driven Equivalence Checking”. In: *Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications*. OOPSLA ’13. Indianapolis, Indiana, USA: ACM, 2013, pp. 391–406. ISBN: 978-1-4503-2374-1. DOI: [10.1145/2509136.2509509](https://doi.acm.org/10.1145/2509136.2509509). URL: <http://doi.acm.org/10.1145/2509136.2509509>.
- [42] Shubhani. “Counterexample-Guided Equivalence Checking”. PhD thesis. Indian Institute of Technology Delhi, 2023. URL: https://sorav.compiler.ai/pubs/shubhani_thesis.pdf.
- [43] Bjarne Steensgaard. “Points-to analysis in almost linear time”. In: *Proceedings of the 23rd ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. 1996, pp. 32–41.
- [44] Michael Stepp, Ross Tate, and Sorin Lerner. “Equality-based Translation Validator for LLVM”. In: *Proceedings of the 23rd International Conference on Computer Aided Verification*. CAV’11. Snowbird, UT: Springer-Verlag, 2011, pp. 737–742.

ISBN: 978-3-642-22109-5. URL: <http://dl.acm.org/citation.cfm?id=2032305.2032364>.

- [45] Chengnian Sun, Vu Le, and Zhendong Su. “Finding Compiler Bugs via Live Code Mutation”. In: *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications*. OOPSLA 2016. Amsterdam, Netherlands: Association for Computing Machinery, 2016, pp. 849–863. ISBN: 9781450344449. DOI: [10.1145/2983990.2984038](https://doi.org/10.1145/2983990.2984038). URL: <https://doi.org/10.1145/2983990.2984038>.
- [46] Chengnian Sun, Vu Le, Qirun Zhang, and Zhendong Su. “Toward understanding compiler bugs in GCC and LLVM”. In: *Proceedings of the 25th International Symposium on Software Testing and Analysis*. ISSTA 2016. Saarbrücken, Germany: Association for Computing Machinery, 2016, pp. 294–305. ISBN: 9781450343909. DOI: [10.1145/2931037.2931074](https://doi.org/10.1145/2931037.2931074). URL: <https://doi.org/10.1145/2931037.2931074>.
- [47] Andrew S. Tanenbaum, Hans van Staveren, E. G. Keizer, and Johan W. Stevenson. “A Practical Tool Kit for Making Portable Compilers”. In: *Commun. ACM* 26.9 (Sept. 1983), pp. 654–660. ISSN: 0001-0782. DOI: [10.1145/358172.358182](https://doi.org/10.1145/358172.358182). URL: <https://doi.org/10.1145/358172.358182>.
- [48] The LLVM developers. *The LLVM Compiler Infrastructure*. URL: <https://llvm.org/>.
- [49] Jean-Baptiste Tristan, Paul Govereau, and Greg Morrisett. “Evaluating Value-graph Translation Validation for LLVM”. In: *Proceedings of the 32Nd ACM SIGPLAN Conference on Programming Language Design and Implementation*. PLDI '11. San Jose, California, USA: ACM, 2011, pp. 295–305. ISBN: 978-1-4503-0663-8. DOI: [10.1145/1993498.1993533](http://doi.acm.org/10.1145/1993498.1993533). URL: <http://doi.acm.org/10.1145/1993498.1993533>.

- [50] Jean-Baptiste Tristan and Xavier Leroy. “Formal verification of translation validators: A case study on instruction scheduling optimizations”. In: *Proceedings of the 35th ACM Symposium on Principles of Programming Languages (POPL’08)*. ACM Press, Jan. 2008, pp. 17–27. URL: <http://xavierleroy.org/publi/validation-scheduling.pdf>.
- [51] Xi Wang, Haogang Chen, Alvin Cheung, Zhihao Jia, Nickolai Zeldovich, and M. Frans Kaashoek. “Undefined behavior: what happened to my code?” In: *Proceedings of the Asia-Pacific Workshop on Systems. APSYS ’12*. Seoul, Republic of Korea: Association for Computing Machinery, 2012. ISBN: 9781450316699. DOI: [10.1145/2349896.2349905](https://doi.org/10.1145/2349896.2349905). URL: <https://doi.org/10.1145/2349896.2349905>.
- [52] Xuejun Yang, Yang Chen, Eric Eide, and John Regehr. “Finding and Understanding Bugs in C Compilers”. In: *Proceedings of the 32Nd ACM SIGPLAN Conference on Programming Language Design and Implementation. PLDI ’11*. San Jose, California, USA: ACM, 2011, pp. 283–294. ISBN: 978-1-4503-0663-8. DOI: [10.1145/1993498.1993532](http://doi.acm.org/10.1145/1993498.1993532). URL: <http://doi.acm.org/10.1145/1993498.1993532>.
- [53] *Z3 bug report for an unsound model*. <https://github.com/Z3Prover/z3/issues/7132>. 2024.
- [54] Anna Zaks and Amir Pnueli. “CoVaC: Compiler Validation by Program Analysis of the Cross-Product”. In: *Proceedings of the 15th International Symposium on Formal Methods. FM ’08*. Turku, Finland: Springer-Verlag, 2008, pp. 35–51. ISBN: 978-3-540-68235-6. DOI: [10.1007/978-3-540-68237-0_5](http://dx.doi.org/10.1007/978-3-540-68237-0_5). URL: http://dx.doi.org/10.1007/978-3-540-68237-0_5.
- [55] Jianzhou Zhao, Santosh Nagarakatte, Milo M.K. Martin, and Steve Zdancewic. “Formal Verification of SSA-based Optimizations for LLVM”. In: *Proceedings*

of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation. PLDI '13. Seattle, Washington, USA: ACM, 2013, pp. 175–186. ISBN: 978-1-4503-2014-6. DOI: [10.1145/2491956.2462164](https://doi.org/10.1145/2491956.2462164). URL: <http://doi.acm.org/10.1145/2491956.2462164>.

- [56] Jianzhou Zhao, Santosh Nagarakatte, Milo M.K. Martin, and Steve Zdancewic. “Formalizing the LLVM Intermediate Representation for Verified Program Transformations”. In: *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL '12. Philadelphia, PA, USA: ACM, 2012, pp. 427–440. ISBN: 978-1-4503-1083-3. DOI: [10.1145/2103656.2103709](https://doi.org/10.1145/2103656.2103709). URL: <http://doi.acm.org/10.1145/2103656.2103709>.
- [57] Lenore Zuck, Amir Pnueli, Yi Fang, and Benjamin Goldberg. “VOC: A Methodology for the Translation Validation of Optimizing Compilers”. In: 9.3 (Mar. 28, 2003), pp. 223–247.