

Non-termination Proving at Scale

AZALEA RAAD, Imperial College London, UK and Bloomberg, USA
JULIEN VANEGUE, Bloomberg, USA and Imperial College London, UK
PETER O'HEARN, University College London, UK

Program termination is a classic non-safety property whose falsification cannot in general be witnessed by a finite trace. This makes testing for non-termination challenging, and also a natural target for symbolic proof. Several works in the literature apply non-termination proving to small, self-contained benchmarks, but it has not been developed for large, real-world projects; as such, despite its allure, non-termination proving has had limited practical impact. We develop a *compositional* theory for non-termination proving, paving the way for its *scalable* application to large codebases. Discovering non-termination is an under-approximate problem, and we present UNTER, a *sound and complete* under-approximate logic for proving non-termination. We then extend UNTER with separation logic and develop UNTER^{SL} for heap-manipulating programs, yielding a compositional proof method amenable to automation via under-approximation and bi-abduction. We extend the Pulse analyser from Meta and develop Pulse$^{\infty}$, an automated, compositional prover for non-termination based on UNTER^{SL}. We have run Pulse$^{\infty}$ on large codebases and libraries, each comprising hundreds of thousands of lines of code, including OpenSSL, libxml2, libxpm and CryptoPP; we discovered several previously-unknown non-termination bugs and have reported them to developers of these libraries.

CCS Concepts: • **Theory of computation** → **Programming logic**; **Separation logic**; **Program analysis**; *Program verification*; *Hoare logic*; • **Software and its engineering** → General programming languages.

Additional Key Words and Phrases: Divergence, non-termination, under-approximation, incorrectness logic

ACM Reference Format:
Azalea Raad, Julien Vanegue, and Peter O'Hearn. 2024. Non-termination Proving at Scale. *Proc. ACM Program. Lang.* 8, OOPSLA2, Article 280 (October 2024), 29 pages. https://doi.org/10.1145/3689720

1 Introduction

Why Prove Non-termination? Non-termination (divergence) is a fundamental problem in computer science, dating back to the halting problem. Assuming an unbounded memory or tape, neither it nor its complement is recursively enumerable, making it difficult to approach using testing. This makes non-termination an attractive target for symbolic proof techniques.

Apart from its fundamental nature, one can also ask: is non-termination a practical problem? To understand this better we manually evaluated the bugs in the Common Vulnerabilities and Exposures (CVE) database for security bugs that are due to non-termination, e.g. denial-of-service attacks. We found 916 such CVE's between 2000 and 2022 – see the extended version [Raad et al. 2024a, §A]. (For ongoing computations such as operating systems, potential non-termination is desirable and unavoidable. Here, we are concerned with buggy, unintended non-termination.)

Interestingly, we did not detect any reduction in non-termination CVE's during this period. For example, we found 4 such bugs from 2000 and 28 from 2022. We stress that our manual approach

Authors' Contact Information: Azalea Raad, azalea.raad@imperial.ac.uk, Imperial College London, UK and Bloomberg, USA; Julien Vanegue, jvanegue@bloomberg.net, Bloomberg, USA and Imperial College London, UK; Peter O'Hearn, p.ohearn@ucl.ac.uk, University College London, UK.



This work is licensed under a Creative Commons Attribution 4.0 International License.
© 2024 Copyright held by the owner/author(s).
ACM 2475-1421/2024/10-ART280
https://doi.org/10.1145/3689720

might have missed some non-termination CVE's, there is more code in 2022 than in 2000, and the classification of non-termination CVE's might be non-uniform. This data, however, motivated our work on the science and engineering of tools for detecting non-termination bugs.

Why Compositional? A compositional analysis is one where the analysis result of a composite program is computed from those of its constituent parts [Calcagno et al. 2011]. Compositionality enables program analysis to be deployed as part of a code review process, where code snippets in a pull request are analysed without the need to re-analyse the entire program (or even to have an entire program, which might not yet exist). A case study from Facebook [Distefano et al. 2019] describes how deploying a compositional static analysis tool on pull requests achieved a 70% fix rate, while the same analysis had a near 0% fix rate for a batch deployment (where a list of bugs is given outside of code review). This illustrates how a deployment of static analysis that meets programmers in their workflows can have considerable advantages over ones that ask them to leave their flow. (See the Facebook article [Distefano et al. 2019] and a related article from Google [Sadowski et al. 2018] for more information.)

It stands to reason that if an accurate non-termination prover is developed which is fast enough to be deployed at pull-request time, then it would have the potential to have more non-termination bugs fixed, early. We will not in this paper go so far as setting up an industrial deployment of non-termination proving in the CICD system of a company, but we take the Facebook/Google experience referenced above as motivation for our scientific goals: to establish a compositional proof method together with an algorithm which allow for automatic compositional program analysis, and initial experiments to probe its feasibility.

Our Approach. Proving non-termination is an *under-approximation* problem as the aim is to establish the *existence* of non-terminating executions. Therefore, for compositional reasoning it is natural to consider a formalism akin to incorrectness logic (IL) [O'Hearn 2019], which brings the compositional nature of Hoare logic to bug proving. It turns out the form of under-approximation we need is a reversed form of that in IL, based on what is called the 'backwards under-approximate triple' by Möller et al. [2021] and the 'total Hoare triple' by de Vries and Koutavas [2011].

The *backwards under-approximate* (BUA) triple $\vdash_{\mathsf{B}} \lfloor p \rfloor\ \mathsf{C}\ \lfloor ok\colon q \rfloor$ denotes that p is a *subset* of the states from which q can be reached executing C. That is, from any state in p it is possible to reach some state in q by executing C. This triple is forwards in terms of reachability, but backwards in terms of under-approximation (mirroring IL): p under-approximates the weakest *possible* precondition, wpp, of C on q: $p \subseteq \mathsf{wpp}(\mathsf{C}, q)$. Here, wpp is the inverse image of the C (relational) semantics, obtained by running Dijkstra's strongest post-condition on the reversal of C.

We next extend our BUA triples with under-approximate *divergence* triples. Specifically, we develop *under-approximate non-termination logic* (UNTER), where we write $\vdash \lfloor p \rfloor\ \mathsf{C}\ [\infty]$ to denote that every state in p leads to a divergent (infinite) execution via C. Note that this does not state that *every* execution diverges; rather, that each pre-state leads to *some* divergent execution.

We can then state a proof rule for divergence as shown across. The idea behind this rule is simple. As $p \wedge B$ holds initially, after one loop iteration we can get to a state where $p \wedge B$ continues to hold because of the BUA triple in the premise. And in that case we can take one more step, *ad infinitum*.

$$\frac{\vdash_{\mathsf{B}} \lfloor p \wedge B \rfloor\ \mathsf{C}\ \lfloor ok\colon p \wedge B \rfloor}{\lfloor p \wedge B \rfloor\ \mathsf{while}\ (B)\ \mathsf{C}\ [\infty]}$$

This proof method is related to a method of non-termination testing whereby one looks for a concrete state to which a loop returns: this would witness divergence as one can get back to the same state again. As a testing method this approach is incomplete, in the presence of unbounded resources (e.g. a Turing machine tape) which gives rise to infinitely many states: then it is possible to diverge without returning to the same state twice. But the proof method uses a logical assertion

and not a concrete state, and is indeed complete for proving non-termination as we show later (take p to be the set of all states that lead to divergence). The proof method is also related to the idea of 'recurrence sets' by Gupta et al. [2008] (see §8 for the relation to their and other work).

Our aim is to *automate* divergence proof rules such as that above. There are several key observations in our approach. First, and remarkably, if we apply the strategy used commonly in abstract interpretation, namely iterating the abstract semantics of loops until we reach a fixpoint, then we will have proven non-termination of a loop when a fixpoint is reached. In abstract interpretation this would not imply divergence, but with our under-approximate UNTER logic it does. However, while we can employ the usual method of fixpoint iteration, since not all loops diverge, we additionally need a way to stop the analysis before a fixpoint is reached. It turns out that we can employ similar techniques to IL and bounded model checking, by simply stopping after some fixed number of iterations even when we do not have a fixpoint. This flexibility is not available in Hoare logic, or in over-approximate abstract interpretation, where stopping early is unsound.

Second, by detailing the relationship to the original IL we reveal additional possibilities for automation. Indeed, the BUA proof system is almost the same as that of IL, with the difference limited to the rule of consequence (see §2, §3). The use of the backwards predicate transformer wpp perhaps suggests to attempt a backwards program analysis, at least for a whole-program analysis: given a post, such an analysis would compute an under-approximation of backwards reachability at each program point; in a sense, the mirror image of Floyd's method of calculating over-approximations for forwards reachability. However, a forwards-running analysis is also possible, as long as we *abduce* preconditions as we go forwards: this semantics calculates a collection of triples at each program point, connecting procedure-entry to the program point. In addition to furnishing a compositional inter-procedural analysis, abduction is necessary here: there is no forwards predicate transformer semantics, evidenced by the fact that for some programs C and pre-conditions p there is no post-condition delivering a valid triple $\vdash_{\mathsf{B}} [p]\ \mathsf{C}\ [ok\colon ??]$.

The third key point for automation is that the close connection between the BUA and original IL proof theories suggests a method of automation that leverages *separation logic* [Ishtiaq and O'Hearn 2001], and which is obtained by small changes and a fundamental addition to the existing Pulse program analyser [Le et al. 2022] from Facebook. We observe that Pulse uses a restricted version of the rule of consequence, making it compatible both with BUA and IL triples. We thus develop UNTER$^{\mathrm{SL}}$ as an extension of UNTER (with divergent triples) with separation logic

Our Pulse$^{\infty}$ Prototype. To demonstrate the feasibility of UNTER$^{\mathrm{SL}}$, we have developed Pulse$^{\infty}$, a prototype compositional non-termination prover underpinned by UNTER$^{\mathrm{SL}}$, as an extension of the existing Pulse program analyser (which is underpinned by the ISL theory [Raad et al. 2020] and is compatible with BUA reasoning). To evaluate Pulse$^{\infty}$, we have compared its performance against cutting edge tools such as DynamiTe [Le et al. 2020] by running it on the state-of -the-art non-linear arithmetic extension of the SV-COMP benchmark. While Pulse$^{\infty}$ is not comparable to these tools in the divergence bugs it found, in that Pulse$^{\infty}$ successfully found divergence bugs that were undetected by these tools, while missed others found by these tools (see Table 2 on p. 22), it reported *zero false positives* thanks to its under-approximate nature, in contrast to DynamiTe which suffered several false positives.

More significantly, we have successfully run Pulse$^{\infty}$ on large codebases and libraries, each comprising hundreds of thousands of lines of code (LOC), including OpenSSL, libxml2, CryptoPP and libxpm. To our knowledge, Pulse$^{\infty}$ is *the first automated tool* for detecting divergence bugs in *large code bases and libraries*. As we discuss in the related work (§8), existing tools either focus on a small class of (integer) C programs *without functions calls* (thus excluding libraries), or stipulate

```
1  int http_server_get_asn1_req(const ASN1_ITEM *it, ASN1_VALUE **preq,
2                               char **ppath, BIO **pcbio, BIO *acbio, (...)) {
3    for (;;) {
4      char *key, *value;
5      len = BIO_gets(cbio, inbuf, sizeof(inbuf));
6      if (len <= 0) goto out;
7      key = inbuf;
8      value = strchr(key, ':');
9      if (value == NULL) goto out;
10     *(value++) = '\0'; }
11 }
```

Listing 1. A divergence bug found by Pulse$^\infty$ in OpenSSL

the presence of a `main` procedure (once again ruling out libraries), or require that the code under analysis be *confined to a single file* (thus precluding large code bases).

Using Pulse$^\infty$, we have automatically analysed each of these libraries *within minutes*. For instance, running Pulse$^\infty$ on OpenSSL (804 kLOC) completed *under two minutes* and found four hitherto-unknown divergence bugs. Indeed, we have found new divergence bugs in OpenSSL, libxml2, CryptoPP and libxpm. In the cases of CryptoPP and libxpm we have submitted pull requests with patches. In the case of libxml2, we shared our findings with the development team, who suggested the reported code branches may be unreachable, and thus should be removed entirely from the codebase. For OpenSSL, we present a divergence bug found by Pulse$^\infty$ in Listing 1, containing a potential infinite `for` loop on lines 3–10. Specifically, the function shown can keep reading more data with `BIO_gets` (line 5) and never break from the loop. This analysis requires inter-procedural, heap and arithmetic reasoning all at once, which is not uncommon in real-world code. We contacted a senior OpenSSL developer about this bug, who confirmed that the code should be made more restrictive and enforce an upper bound on the amount of data read at this location.

Contributions and Outline. In §2 we present an intuitive overview of BUA and IL reasoning, and describe how we extend them to reason about non-termination. In §3 we present UNTer as a BUA proof system and extend it to account for non-termination, yielding a compositional proof method. In §4 we present several examples of divergence and show how we can detect them using UNTer. In §5 we present the semantic model of UNTer and show that it is *sound* and *complete*. In §6 we develop UNTer$^{\text{SL}}$ by extending UNTer with separation logic for heap reasoning. In §7 we extend the under-approximate reasoning framework of Pulse to develop Pulse$^\infty$, an automated, compositional prover for non-termination; we evaluate Pulse$^\infty$ against other tools in the literature, and report our results of running Pulse$^\infty$ on large libraries. We discuss related work in detail in §8.

2 Overview

Incorrectness Logic and Under-Approximate Reasoning. As Godefroid [2005] argues, the main value of analysis tools lies in the discovery of bugs, not in the proof of program correctness. A bug presented to a developer is often a more convincing utility of a tool than a correctness proof, which is often carried out under certain assumptions that may not hold. This is evidenced by the recent trend in under-approximate reasoning techniques [O'Hearn 2019; Raad et al. 2020, 2022] and their significant success at finding bugs on an *industrial scale* [Blackshear et al. 2018; Le et al. 2022]. Specifically, Incorrectness logic (IL) [O'Hearn 2019] presented an under-approximate formal foundation for bug detection. It was later extended to enable compositional bug detection in heap-manipulating programs [Raad et al. 2020], and to support concurrency [Raad et al. 2022,

2023]. IL and its later extensions are instances of under-approximate reasoning and are associated with *no-false-positives theorems*, ensuring that all bugs identified by them are true positives.

Intuitively, the under-approximate nature of IL stems from considering a *subset* of program behaviours. More concretely, given a program C whose behaviours (traces) is given by the set S, IL reasoning considers a subset (under-approximated) $S_u \subseteq S$ of the C behaviours. This makes IL ideally suited for bug-detection as it guarantees no-false-positives: if one detects a bug in the smaller set S_u, then the bug is also guaranteed to be in S and thus exhibited by C. This is in contrast to over-approximate reasoning techniques such as Hoare logic, where one considers a superset (over-approximated) set $S_o \supseteq S$ of C behaviours, making them ideal for verification (as they guarantee no false negatives): if one can show that the larger set S_o contains only correct behaviours, then the smaller set S also contains correct behaviours only.

An IL triple, also referred to as a *forward, under-approximate* (FUA) triple, is of the form $\vdash_{\mathsf{F}} [p]\ \mathsf{C}\ [\epsilon : q]$, where F hints at its *forward* direction, denoting that q is a subset of program behaviours when C is run (forward) from the states in p. In other words, an FUA triple describes *backward reachability*: *every* post-state in q is *reachable* by running C forward on *some* pre-state in p. The ϵ is an *exit condition* and may be either *ok*, to denote a normal execution or *er* to denote an erroneous execution. For instance, executing an explicit error (e.g. assert(false)) terminates erroneously and the underlying states are unchanged: $\vdash_{\mathsf{F}} [p]$ error $[er : p]$. The under-approximate nature of FUA triples is best illustrated by their rules for reasoning about branches and loops. To show that a behaviour is possible when executing $\mathsf{C}_1 + \mathsf{C}_2$ (where + denotes non-deterministic choice), it is sufficient to show the behaviour is possible when executing one of the branches, i.e. executing C_i for *some* (rather than all) $i \in \{1, 2\}$, as shown in ChoiceF below (left). Similarly, to show a behaviour is possible when executing $\mathsf{C}^\star$ (where $\mathsf{C}^\star$ denotes a non-deterministic loop, executing C for zero or more iterations), it suffices to show it is possible when executing C for a particular number $n \in \mathbb{N}$ of iterations, as shown in LoopF below (right), where C^n denotes executing C for n times.

$$\textsc{ChoiceF}\quad \frac{\vdash_{\mathsf{F}} [p]\ \mathsf{C}_i\ [\epsilon : q] \quad \text{for some } i \in \{1, 2\}}{\vdash_{\mathsf{F}} [p]\ \mathsf{C}_1 + \mathsf{C}_2\ [\epsilon : q]} \qquad \textsc{LoopF}\quad \frac{\vdash_{\mathsf{F}} [p]\ \mathsf{C}^n\ [\epsilon : q] \quad \text{for some } n \in \mathbb{N}}{\vdash_{\mathsf{F}} [p]\ \mathsf{C}^\star\ [\epsilon : q]}$$

Non-termination and Under-Approximate Reasoning. Existing literature includes a large body of work [Berdine et al. 2007, 2006; Chawdhary et al. 2008; Cook et al. 2006a,b; da Rocha Pinto et al. 2016; D'Osualdo et al. 2021; Liang and Feng 2016] on *termination* analysis, proving that a program C always terminates by showing that *all* traces of C terminate for *all* given inputs. Showing that a program C terminates is compatible with *over-approximate* reasoning frameworks. Specifically, when the traces of C are given by the set S, showing that all traces in a larger set $S_o \supseteq S$ terminate is sufficient for showing that all traces in S terminate. Showing termination is difficult in the presence of loops: to show that a loop L terminates typically involves the challenging task of establishing a *loop invariant* as well as a *well-founded measure* (a.k.a. a ranking function) that is decreased after each iteration. Establishing such invariants and measures is far from straightforward and typically involves reasoning about *ordinal* (rather than natural) numbers.

Showing that a program C does not terminate is compatible with *under-approximate* reasoning: when the traces of C are given by the set S, showing that the traces in a smaller (under-approximate), possibly singleton, set $S_u \subseteq S$ do not terminate is sufficient for showing that C does not terminate.

Inspired by the success of under-approximate analysis techniques and their industrial application of detecting bugs at scale, we develop ***u**nder-approximate, **n**on-**ter**mination logic* (UNTer) as the first *formal, under-approximate foundation* for detecting non-termination bugs. As with existing under-approximate techniques, UNTer is associated with a no-false-positives theorem, ensuring that all non-termination bugs identified are true positives. More concretely, UNTer enables deriving

under-approximate, *divergent* triples of the form $[p]\ \mathsf{C}\ [\infty]$, stating that starting from the states in p program C has divergent (non-terminating) traces. Note that $[p]\ \mathsf{C}\ [\infty]$ does not state that C never terminates (i.e. that *all* traces of C are divergent), but rather that it is possible for C not to terminate (i.e. *some* traces of C are divergent). For instance, given the program $\mathsf{C} \triangleq \mathsf{skip} + (\mathsf{while}\ (\mathsf{true})\ \mathsf{skip})$, the triple $[\mathsf{true}]\ \mathsf{C}\ [\infty]$ is valid, since starting from any state (in true) C can always diverge by taking the right branch, even though taking the left branch would immediately lead to termination.

Divergent Triples and FUA Triples. As in the existing formal systems for reasoning about programs (be they over- or under-approximate), we should ideally reason about non-termination in a *compositional* fashion. For instance, given $\mathsf{C}_L \triangleq x := 1; \mathsf{while}\ (x > 0)\ x{+}{+}$ and an arbitrary initial value v, to show that the triple $[x = v]\ \mathsf{C}\ [\infty]$ holds (i.e. C_L does not terminate starting from states satisfying $x = v$), we should ideally show that 1) running $x := 1$ on states in which $x = v$ terminates and modifies the states to those where $x = 1$; and 2) running $\mathsf{while}\ (x > 0)\ x{+}{+}$ on states where $x = 1$ diverges, i.e. $[x = 1]\ \mathsf{while}\ (x > 0)\ x{+}{+}\ [\infty]$. To do (1), we need to reason about *non-divergent* (terminating) program executions in an *under-approximate* fashion. At first glance, this seems an ideal job for FUA triples as they under-approximate reachable program behaviours upon termination; as such, to establish (1), we could simply show $\vdash_{\mathsf{F}} [x = v]\ x := 1\ [ok\colon x = 1]$.

A key feature of our UNTer framework is proof rules for establishing when a loop does not terminate. As a first naive attempt, we can propose the LoopBad rule below (left), stating that if initially the while condition B holds, and executing one iteration of the loop body C starting from p leaves the states (p) and the loop condition (B) unchanged, then while (B) C diverges.

$$\textsc{LoopBad}\ \frac{\vdash_{\mathsf{F}} [p \wedge B]\ \mathsf{C}\ [ok\colon p \wedge B]}{[p \wedge B]\ \mathsf{while}\ (B)\ \mathsf{C}\ [\infty]} \qquad \textsc{LoopFix}\ \frac{\vdash_{\mathsf{B}} [p \wedge B]\ \mathsf{C}\ [ok\colon p \wedge B]}{[p \wedge B]\ \mathsf{while}\ (B)\ \mathsf{C}\ [\infty]}$$

On closer inspection, however, this rule is unsound. Consider the program $\mathsf{while}\ (x > 0)\ x{-}{-}$; this program always terminates regardless of the value of x (for non-positive values the loop is never entered; positive values are eventually decremented to zero). As such, the triple $[x > 0]\ \mathsf{while}\ (x > 0)\ x{-}{-}\ [\infty]$ is invalid. Nevertheless, we can derive it using LoopBad by showing $\vdash_{\mathsf{F}} [x > 0]\ x{-}{-}\ [ok\colon x > 0]$. Specifically, the $\vdash_{\mathsf{F}} [x > 0]\ x{-}{-}\ [ok\colon x > 0]$ triple stipulates that every post-state in $x > 0$ be reachable from some pre-state in $x > 0$, which is indeed the case. More concretely, consider an arbitrary post-state $s_q \in x > 0$ and let $s_q(x) = v$ (i.e. x holds value v in s_q) for some $v > 0$. State s_q is then reachable by running $x{-}{-}$ on a state $s_p = s_q[x \mapsto v{+}1]$ and $s_p \in x > 0$ (as $v > 0$).

Backward Under-Approximate Triples. Intuitively, the problem lies in the backward reachability of FUA triples: it stipulates that each post-state be reachable from some pre-state, which does not necessarily lead to divergence. In other words, having a backward chain of C executions from $p \wedge B$ to $p \wedge B$ does not yield an infinite execution. Instead, we need a forward chain of C executions from $p \wedge B$ to $p \wedge B$, as we can then repeat this execution forward *ad infinitum*. This is captured in the LoopFix rule above (right), where a *backward, under-approximate* (BUA) triple $\vdash_{\mathsf{B}} [p]\ \mathsf{C}\ [\epsilon : q]$ states that every pre-state in p reaches some post-state in q by executing C. Therefore, if we show that each iteration of the loop body transitions each pre-state in $p \wedge B$ to some post-state also in $p \wedge B$, then we can repeat this transition infinitely, leading to divergence. Note that in the example above, we cannot show $\vdash_{\mathsf{B}} [x > 0]\ x{-}{-}\ [ok\colon x > 0]$ (unlike the $\vdash_{\mathsf{F}}$ variant): given state $s_p \in x > 0$ with $s_p(x) = 1$, running $x{-}{-}$ on s_p yields a state $s_q = s_p[x \mapsto 0]$, which is *not* in $x > 0$. As such, using LoopFix, we cannot derive the invalid triple $[x > 0]\ \mathsf{while}\ (x > 0)\ x{-}{-}\ [\infty]$. Note that while BUA triples describe *forward reachability*, they denote *backward under-approximation*: $p \subseteq \mathsf{wpp}(\mathsf{C}, q)$, where $\mathsf{wpp}(\mathsf{C}, q)$ denotes running C *backwards* from q. That is, BUA triples mirror FUA ones (which describe *backward reachability* but *forward under-approximation*).

In order to present our divergence proof rules in a compositional fashion, we thus use BUA triples to describe normal, terminating executions. For instance, in order to show that $C_1; C_2$ does not terminate starting from p, we can show either C_1 does not terminate starting from p (i.e. $[p]\ C_1\ [\infty]$), or C_1 terminates normally transforming the states to q, and C_2 does not terminate starting from q (i.e. $\vdash_B [p]\ C_1\ [ok: q]$ and $[q]\ C_2\ [\infty]$). This is captured by the Div-Seq1 and Div-Seq2 rules in Fig. 2 (§3), where we present our full set of proof rules for detecting divergence.

Forward versus Backward Under-Approximate Triples. As with FUA triples, BUA triples are also inherently under-approximate. Most notably, as we show in §3, the BUA rules for reasoning about branches and loops are identical to their FUA counterparts; i.e. the $\vdash_F$ in ChoiceF and LoopF above can simply be replaced with $\vdash_B$ (see Fig. 1). Indeed, almost all FUA and BUA proof rules coincide, and the only difference between FUA and BUA rules lie in their associated rules of consequence, namely the ConsF (for FUA) and ConsB (for BUA) rules in Fig. 1 (p. 10). However, as we describe shortly, in the practical context of industrially-deployed (under-approximate) bug detection tools such as Pulse [Le et al. 2022], it is straightforward to reconcile this difference between FUA and BUA and to develop a unified, under-approximate reasoning framework.

The main application of the FUA rule of consequence, ConsF, is in conjunction with the rule of disjunction, Disj in Fig. 1 (p. 10). More concretely, when a given program contains multiple branches, thanks to the ChoiceF rule, we can analyse each branch (and not necessarily all branches) in isolation and generate a separate triple. Subsequently, we can merge them into a single triple using Disj. However, when there are many branches (and subsequently many disjuncts in the pre- and post-states), we can simply use ConsF to drop some of the disjuncts in the *post-states*. (Note that using ConsB analogously allows us to drop some of the disjuncts in the *pre-states*.)

However, as our conversations with the lead engineer behind Pulse have revealed, in the practical setting of such tools this scenario rarely arises, and it is handled differently when it does. Specifically, different triples of a program are not merged very often, as it is simpler and more efficient to keep them separate. Second, when triples *are* merged, they are done so in a fashion that additionally *tracks* the correspondence between the disjuncts in the pre- and post-states. Specifically, note that the Disj rule is *lossy*: while in its premise we know that the post-states in q_1 (resp. q_2) are reached from the pre-states in p_1 (resp. p_2), we lose this correspondence in the conclusion and only know that the post-states in $q_1 \vee q_2$ are reached from the pre-states in $p_1 \vee p_2$. As such, when merging the triples $\vdash_F [p_1]\ C\ [\epsilon : q_1]$ and $\vdash_F [p_2]\ C\ [\epsilon : q_2]$ into $\vdash_F [p_1 \vee p_2]\ C\ [\epsilon : q_1 \vee q_2]$, Pulse additionally tracks the correspondence between p_1 and q_1 (resp. p_2 and q_2). This is beneficial when later dropping branches: when dropping the disjuncts in the post-states (e.g. q_2), we can also drop their associated pre-states (p_2). This allows us to avoid accumulating 'clutter' in the pre-states and is tantamount to dropping a full triple rather than its post-states only.

We thus follow a similar approach here which allows us to unify FUA and BUA reasoning. More concretely, we introduce the notion of *indexed disjunctions*, $P, Q \in \mathbb{N} \xrightarrow{\text{fin}} \mathcal{P}(\text{State})$. Intuitively, an indexed disjunction P can be flattened into a standard disjunction as $\bigvee_{i \in dom(P)} P(i)$. We write $[P]\ C\ [\epsilon : Q]$ as a shorthand for $dom(P) = dom(Q) \wedge \forall i \in dom(P).\ [P(i)]\ C\ [\epsilon : Q(i)]$, denoting a merged set of triples. Note that a triple $[p]\ C\ [\epsilon : q]$ can be simply lifted to $[P]\ C\ [\epsilon : Q]$, where $dom(P) = dom(Q) = \{0\}$ with $P(0) = p$ and $Q(0) = q$. We can then use the DisjTrack rule (Fig. 1 on p. 10) to merge indexed disjuncts – note that the $dom(P_1) \cap dom(P_2) = \emptyset$ premise can be simply satisfied by renaming the domain of P_2. Observe that unlike the Disj rule, DisjTrack is not lossy and preserves the pre-post correspondence. Finally, the unified rule of consequence, Cons (Fig. 1), allows us to drop matching disjuncts from both the pre- and post-states, where $P \downarrow I$ denotes restricting the domain of P to I. The unified Cons rule can be used for both FUA and BUA reasoning.

Unified Triples and Bug Catching Tools. Note that the rules in Fig. 1, excluding ConsB, ConsF and Disj (and instead including Cons and DisjTrack) correspond to the reasoning principles used in the industrially deployed Pulse tool. That is, although Pulse is formally underpinned by IL (with FUA triples), it does not use ConsF and Disj, and instead uses Cons and DisjTrack, meaning that using our unified rules (suitable for both FUA and BUA reasoning) has no practical ramifications, and we can use Pulse as it is! This is indeed great news: in order to reason about divergence, we can extend Pulse without changing its underlying principles, and simply add our divergence rules.

Theoretical Connection between BUA and FUA Triples. Note that while it is useful to have both FUA and BUA triples, *theoretically* speaking, only the BUA triples are needed for proving non-termination. As such, the BUA proof system constitutes one of our main contributions (while the FUA proof system was previously developed by de Vries and Koutavas [2011]; O'Hearn [2019]. Moreover, as mentioned above, with the exception of their associated rules of consequence (ConsF and ConsB in Fig. 1) all other FUA and BUA reasoning principles and proof rules coincide. In §5 we bolster this intuition (Theorem 10) by showing that given any under-approximate triple $[p]\ \mathsf{C}\ [\epsilon : q]$, if $[p]\ \mathsf{C}\ [\epsilon : q]$ is a valid *FUA* triple *and* its pre-states (p) are *FUA-minimal*, then $[p]\ \mathsf{C}\ [\epsilon : q]$ is also a valid *BUA* triple. The pre-states p are FUA-minimal if for all smaller pre-states $p' \subset p$, the triple $[p']\ \mathsf{C}\ [\epsilon : q]$ is not a valid FUA triple. Intuitively, this ensures that pre-states p have not been arbitrarily weakened (grown) using ConsF.

Conversely, we show that given an under-approximate triple $[p]\ \mathsf{C}\ [\epsilon : q]$, if $[p]\ \mathsf{C}\ [\epsilon : q]$ is a valid *BUA* triple *and* its post-states (q) are *BUA-minimal*, then $[p]\ \mathsf{C}\ [\epsilon : q]$ is also a valid *FUA* triple. Analogously, q is BUA-minimal if for all smaller $q' \subset q$, the triple $[p]\ \mathsf{C}\ [\epsilon : q']$ is not a valid BUA triple. This ensures that the post-states q have not been arbitrarily weakened using ConsB.

Formal Interpretation of Divergent Triples. As discussed above, we write a divergent triple of the form $\lfloor p \rfloor\ \mathsf{C}\ [\infty]$ to denote that C has *some* divergent trace(s) (i.e. in an under-approximate fashion) starting from p. The next question to answer when interpreting such triples is whether there is some divergent trace starting from *every* state in p or *some* state in p. Observe that both interpretations are under-approximate as they pertain to *some* rather than *all* traces of C. Although the latter interpretation is a weaker statement, it is nevertheless sufficient for an under-approximate divergence detection framework: to establish divergence it suffices to show *some* divergent trace is possible from *some* initial state in p. However, under this weaker interpretation, inspecting a divergent triple $\lfloor p \rfloor\ \mathsf{C}\ [\infty]$ yields little information on how the divergence arises (which may be needed for debugging and fixing the cause of divergence): as p may contain many states, it is unclear which state(s) in p lead(s) to divergence (unless p describes a single state). On the other hand, the former, stronger interpretation provides more information for debugging and fixing the cause of divergence as it states that starting from any state in p the program has a divergent trace.

Although more useful, at first glance this stronger interpretation may seem too strong and antithetic to the spirit of under-approximation in UNTer. However, this additional strength is not accompanied by a theoretical or practical cost. In theoretical terms, rather than considering an arbitrarily large set of pre-states that contain some states that may lead to divergence, one can always shrink the pre-states to contain exactly those states that lead to divergence. More concretely, when starting from a state s executing C may diverge, one can establish $\lfloor p \rfloor\ \mathsf{C}\ [\infty]$ by defining p as the singleton set $\{s\}$, rather than an arbitrarily large set that contains s. In practical terms, this stronger interpretation incurs no additional cost when extending an existing under-approximate tool such as Pulse with divergence proof rules. In particular, the divergence rules in Fig. 2 (p. 12) fall into one of two categories: 1) base rules, where the premises contain BUA triples only (e.g. LoopFix above or Div-Loop in Fig. 2); or 2) inductive cases, where the premises contain other divergent triples (e.g. Div-Seq1 in Fig. 2) or a combination of divergent and BUA triples (e.g. Div-Seq2 in Fig. 2).

For the base cases such as LoopFix, thanks to the forward reachability of BUA triples, we already establish the desired result for *every* pre-state. Moreover, as discussed above, the BUA and FUA reasoning principles are almost identical and can be easily unified for practical purposes. As such, extending exiting under-approximate tools with a base case under a strong interpretation incurs no additional cost. Similarly, establishing an inductive case requires establishing its premises, and since neither their BUA premises (as argued above) nor their divergent premises (by inductive hypothesis) incur an additional cost, establishing an inductive case under a strong interpretation incurs no additional cost. We therefore opt for the stronger under-approximate interpretation of divergent triples: $\lfloor p \rfloor\ \mathsf{C}\ [\infty]$ denotes that *every* state in p leads to *some* divergent trace.

3 The UNTer Framework

We present the UNTer framework for detecting non-termination bugs. To present the key ideas underpinning UNTer more clearly, here we develop it as an analogue of Hoare logic/incorrectness logic (IL), in that UNTer enables *global* and not *local* (compositional) reasoning as in separation logic (SL) [Ishtiaq and O'Hearn 2001] and incorrectness separation logic (ISL) [Raad et al. 2020]. Later in §6 we develop an extension of UNTer that marries the compositionality of SL/ISL with the divergence reasoning of UNTer.

Programming Language. To keep our presentation concise, we employ a simple imperative programming language given by the C grammar below. Our language comprises the standard constructs of skip, assignment ($x := e$), assume statements ($\mathsf{assume}(B)$), scoped variable declaration ($\mathsf{local}\ x\ \mathsf{in}\ \mathsf{C}$), sequential composition ($\mathsf{C}_1; \mathsf{C}_2$), non-deterministic choice ($\mathsf{C}_1 + \mathsf{C}_2$) and loops ($\mathsf{C}^\star$), as well as explicit error statements (error, which can be thought of e.g. as $\mathsf{assert}(\mathsf{false})$).

$$\mathsf{C} ::= \mathsf{skip} \mid x := e \mid \mathsf{assume}(B) \mid \mathsf{local}\ x\ \mathsf{in}\ \mathsf{C} \mid \mathsf{error} \mid \mathsf{C}_1 + \mathsf{C}_2 \mid \mathsf{C}_1; \mathsf{C}_2 \mid \mathsf{C}^\star$$

As is standard, deterministic choice and loops can be encoded using their non-deterministic counterparts and assume statements. Specifically, if $(B)\ \mathsf{then}\ \mathsf{C}_1\ \mathsf{else}\ \mathsf{C}_2$ can be encoded as $(\mathsf{assume}(B); \mathsf{C}_1) + (\mathsf{assume}(\neg B); \mathsf{C}_2)$, and $\mathsf{while}\ (B)\ \mathsf{C}$ can be encoded as $(\mathsf{assume}(B); \mathsf{C})^\star; \mathsf{assume}(\neg B)$.

Assertions (Sets of States). The UNTer assertion language is given by the simple grammar below, comprising classical (first-order logic) and Boolean assertions, where $\oplus \in \{=, \neq, <, \leq, \cdots\}$. Other classical connectives can be encoded using existing ones (e.g. $\neg p \triangleq p \Rightarrow \mathsf{false}$). We use p, q, r and their variants (e.g. p') as metavariables for assertions. An assertion describes a set of states, where each state is a (variable) store in $\textsc{Store} \triangleq \textsc{Var} \to \textsc{Val}$, mapping program variables to values.

$$\textsc{Ast} \ni p, q, r ::= \mathsf{false} \mid p \Rightarrow q \mid \exists x.\ p \mid e \oplus e'$$

An expression e is interpreted under a variable store, written as $s(e)$; this interpretation is standard and elided here. We interpret assertions as sets of states, and thus write false for $\emptyset$, $p \vee q$ for $p \cup q$, $p \Rightarrow q$ for state set inclusion ($p \subseteq q$), and so forth. Similarly, $e \oplus e'$ denotes sets of states (stores) in which $s(e) \oplus s(e')$ holds. As discussed in §2, we introduce the notion of *indexed disjunctions*, $P, Q \in \mathbb{N} \xrightarrow{\text{fin}} \mathcal{P}(\textsc{State})$, as a map from numbers to assertions (disjuncts); i.e. $P \equiv \bigvee_{i \in dom(P)} P(i)$.

UNTer Under-Approximate Proof Rules for Termination. Recall from §2 that to reason about divergence in a piecemeal fashion, we reason about terminating sub-programs via (under-approximate) BUA triples. We present the UNTer under-approximate proof rules for terminating programs in Fig. 1. The rules denoted by $\vdash_\dagger$ are *FUA and BUA* rules in that they are valid when interpreted in either the forward ($\vdash_\mathsf{F}$) or backward ($\vdash_\mathsf{B}$) direction. Note that as discussed in §2, with the exception of ConsF and ConsB rules, all rules in Fig. 1 are valid FUA *and* BUA triples.

$$\textsc{Skip}\ \vdash_\dagger [p]\ \mathsf{skip}\ [ok: p] \qquad \textsc{Assign}\ \frac{y \notin \mathsf{fv}(p)}{\vdash_\dagger [p]\ x := e\ [ok: \exists y.\, p[y/x] \wedge x = e[y/x]]} \qquad \textsc{Assume}\ \vdash_\dagger [p \wedge B]\ \mathsf{assume}(B)\ [ok: p \wedge B]$$

$$\textsc{Error}\ \vdash_\dagger [p]\ \mathsf{error}\ [er: p] \qquad \textsc{Seq}\ \frac{\vdash_\dagger [p]\ \mathsf{C}_1\ [ok: r] \qquad \vdash_\dagger [r]\ \mathsf{C}_2\ [\epsilon : q]}{\vdash_\dagger [p]\ \mathsf{C}_1; \mathsf{C}_2\ [\epsilon : q]} \qquad \textsc{SeqEr}\ \frac{\vdash_\dagger [p]\ \mathsf{C}_1\ [er: q]}{\vdash_\dagger [p]\ \mathsf{C}_1; \mathsf{C}_2\ [er: q]}$$

$$\textsc{Choice}\ \frac{\vdash_\dagger [p]\ \mathsf{C}_i\ [\epsilon : q] \qquad \text{for some } i \in \{1,2\}}{\vdash_\dagger [p]\ \mathsf{C}_1 + \mathsf{C}_2\ [\epsilon : q]} \qquad \textsc{Loop0}\ \vdash_\dagger [p]\ \mathsf{C}^\star\ [ok: p] \qquad \textsc{Loop}\ \frac{\vdash_\dagger [p]\ \mathsf{C}^\star; \mathsf{C}\ [\epsilon : q]}{\vdash_\dagger [p]\ \mathsf{C}^\star\ [\epsilon : q]}$$

$$\textsc{Loop-Subvar}\ \frac{\forall n < k.\ \vdash_\dagger [p(n)]\ \mathsf{C}\ [ok: p(n{+}1)]}{\vdash_\dagger [p(0)]\ \mathsf{C}^\star\ [ok: p(k)]} \qquad \textsc{Local}\ \frac{\vdash_\dagger [p]\ \mathsf{C}\ [\epsilon : q]}{\vdash_\dagger [\exists x.\, p]\ \mathsf{local}\ x\ \mathsf{in}\ \mathsf{C}\ [\epsilon : \exists x.\, q]} \qquad \textsc{Subst}\ \frac{\vdash_\dagger [p]\ \mathsf{C}\ [\epsilon : q] \quad x \notin \mathsf{fv}(p, \mathsf{C}, q)}{(\vdash_\dagger [p]\ \mathsf{C}\ [\epsilon : q])[y/x]}$$

$$\textsc{Disj}\ \frac{\vdash_\dagger [p_i]\ \mathsf{C}\ [\epsilon : q_i] \quad \text{for all } i \in I}{\vdash_\dagger \left[\bigvee_{i \in I} p_i\right]\ \mathsf{C}\ \left[\epsilon : \bigvee_{i \in I} q_i\right]} \qquad \textsc{Constancy}\ \frac{\vdash_\dagger [p]\ \mathsf{C}\ [\epsilon : q] \quad \mathsf{fv}(r) \cap \mathsf{mod}(\mathsf{C}) = \emptyset}{\vdash_\dagger [p \wedge r]\ \mathsf{C}\ [\epsilon : q \wedge r]}$$

$$\textsc{ConsF}\ \frac{p' \subseteq p \qquad \vdash_\mathsf{F} [p']\ \mathsf{C}\ [\epsilon : q'] \qquad q \subseteq q'}{\vdash_\mathsf{F} [p]\ \mathsf{C}\ [\epsilon : q]} \qquad \textsc{ConsB}\ \frac{p \subseteq p' \qquad \vdash_\mathsf{B} [p']\ \mathsf{C}\ [\epsilon : q'] \qquad q' \subseteq q}{\vdash_\mathsf{B} [p]\ \mathsf{C}\ [\epsilon : q]}$$

$$\textsc{DisjTrack}\ \frac{\vdash_\dagger [P_1]\ \mathsf{C}\ [\epsilon : Q_1] \qquad \vdash_\dagger [P_2]\ \mathsf{C}\ [\epsilon : Q_2]}{\vdash_\dagger [P_1 \uplus P_2]\ \mathsf{C}\ [\epsilon : Q_1 \uplus Q_2]} \qquad \textsc{Cons}\ \frac{\vdash_\dagger [P]\ \mathsf{C}\ [\epsilon : Q] \qquad I \subseteq dom(P)}{\vdash_\dagger [P \downarrow I]\ \mathsf{C}\ [\epsilon : Q \downarrow I]}$$

$$\textsc{IfTrue}\ \frac{\vdash_\dagger [p \wedge B]\ \mathsf{C}_1\ [\epsilon : q]}{\vdash_\dagger [p \wedge B]\ \mathsf{if}\ (B)\ \mathsf{then}\ \mathsf{C}_1\ \mathsf{else}\ \mathsf{C}_2\ [\epsilon : q]} \qquad \textsc{IfFalse}\ \frac{\vdash_\dagger [p \wedge \neg B]\ \mathsf{C}_2\ [\epsilon : q]}{\vdash_\dagger [p \wedge \neg B]\ \mathsf{if}\ (B)\ \mathsf{then}\ \mathsf{C}_1\ \mathsf{else}\ \mathsf{C}_2\ [\epsilon : q]}$$

$$\textsc{ConsEq}\ \frac{p \Leftrightarrow p' \qquad \vdash_\dagger [p']\ \mathsf{C}\ [\epsilon : q'] \qquad q' \Leftrightarrow q}{\vdash_\dagger [p]\ \mathsf{C}\ [\epsilon : q]} \qquad \textsc{WhileFalse}\ \vdash_\dagger [p \wedge \neg B]\ \mathsf{while}\ (B)\ \mathsf{C}\ [ok: p \wedge \neg B]$$

$$\textsc{WhileSubvar}\ \frac{\forall n < k.\ \vdash_\dagger [p(n) \wedge B]\ \mathsf{C}\ [ok: p(n{+}1) \wedge B] \qquad \vdash_\dagger [p(k) \wedge B]\ \mathsf{C}\ [\epsilon : q \wedge \neg B]}{\vdash_\dagger [p(0) \wedge B]\ \mathsf{while}\ (B)\ \mathsf{C}\ [\epsilon : q \wedge \neg B]}$$

Fig. 1. Under-approximate proof rules where † in each rule can be instantiated as F or B; the highlighted rules can be derived from other rules (see the extended version [Raad et al. 2024a, §B]).

The Skip, Error, Seq, SeqEr, Choice, Loop0, Loop and Disj rules are identical to those of existing FUA logics [O'Hearn 2019; Raad et al. 2020, 2022]. Specifically, executing skip and error leave the state unchanged (Skip and Error), where the former terminates normally while the latter terminates erroneously; Disj allows us to merge multiple triples into one in a lossy fashion (as discussed in §2); the behaviour of a branching program can be under-approximated as the behaviour of *some* of its branches (Choice); and the behaviour of a loop can be under-approximated through bounded

unrolling as zero (Loop0) or more (Loop) iterations. Note that while in correctness frameworks we can over-approximate a loop behaviour via an *invariant*, i.e. an assertion that holds after *any* number of iterations (including zero), in FUA/BUA frameworks we can under-approximate a loop behaviour via a *subvariant* as an indexed assertion p, where $p(n)$ describes the state after n iterations. This is captured by Loop-Subvar: for an arbitrary k, if executing C terminates normally and transforms $p(n)$ to $p(n{+}1)$ for all $n < k$, then $p(k)$ can be reached by executing $\mathsf{C}^\star$ (i.e. executing C for k iterations) from the initial states $p(0)$. The SeqEr captures the short-circuiting behaviour of erroneous executions: if C_1 terminates erroneously, then $\mathsf{C}_1;\mathsf{C}_2$ also terminates erroneously. By contrast, Seq captures the case where executing C_1 does not encounter an error: if executing C_1 terminates normally transforming the states in p to those in r, and executing C_2 terminates as ϵ (either *ok* or *er*) and transforms r to q, then executing $\mathsf{C}_1;\mathsf{C}_2$ terminates as ϵ, transforming p to q.

The Assign rule is identical to the standard Floyd assignment rule and holds for both FUA and BUA. Observe that as noted by O'Hearn [2019], the Hoare assignment rule is not sound for FUA. That is, $\vdash_{\mathsf{F}} \left[p[e/x]\right]\, x := e\, \left[ok\colon p\right]$ is not sound (e.g. let $e = 42$ and p be $x = y$, then the state $s \in p$ such that $s(x) = s(y) = 17$ cannot be reached by executing $x := 42$ on any state in $p[42/x]$. By contrast, the Hoare assignment rule *is* sound for BUA, i.e. $\vdash_{\mathsf{B}} \left[p[e/x]\right]\, x := e\, \left[ok\colon p\right]$ is a sound BUA triple. However, this difference between BUA and FUA does not have a practical ramification as the Floyds assignment rule (in Assign) is sufficient to enable automated reasoning in Pulse.

The Assume, Local and Constancy rules are analogous to the FUA rules of [O'Hearn 2019]. Concretely, executing assume(B) terminates normally and leaves the state unchanged, provided that B holds beforehand. When executing the scoped variable declaration local x in C, the information about x is erased by existentially quantifying it in the pre- and post-states. The Constancy rule is used to adapt triples in different contexts and states: if an assertion r holds before executing C, it also holds afterwards provided that it does not refer to free variables that may have been modified by C. This is captured by the $\mathsf{fv}(r) \cap \mathsf{mod}(\mathsf{C}) = \emptyset$, where $\mathsf{fv}(r)$ denotes the free variables of r and $\mathsf{mod}(\mathsf{C})$ denotes the variables modified by C (i.e. those on the left-hand side of assignments).

As discussed in §2, ConsF and ConsB are the FUA and BUA rules of consequence, respectively. We reconcile the two in the unified rule of consequence, Cons, by using indexed disjunctions, where $dom(P \downarrow I) = I$ and $\forall i \in I.\ (P \downarrow I)(i) = P(i)$. Finally, using indexed disjunctions in DisjTrack we can merge triples in a non-lossy fashion, preserving the pre-post correspondence.

The remaining highlighted rules can be derived from existing rules (see the extended version [Raad et al. 2024a, §B]). The IfTrue (resp. IfFalse) is analogous to its non-deterministic counterpart (Choice) and requires that condition B hold (resp. not hold) at the beginning. The ConsEq simply replaces implication (subset inclusion) in the premises of ConsF and ConsB with equivalence. The WhileFalse states that the pre-states are unchanged by the loop if the condition B does not hold to begin with (i.e. the loop is never entered). The WhileSubvar is analogous to Loop-Subvar and states that if for all $n < k$ an execution of C transforms $p(n) \wedge B$ to $p(n{+}1) \wedge B$, i.e. loop condition B remains true in the first $k{-}1$ iterations, and the k^{th} iteration results in the states in $q \wedge \neg B$ (i.e. it invalidates the loop condition), then while (B) C terminates, transforming the initial states in $p(0) \wedge B$ to those in $q \wedge \neg B$.

UNTer Divergent Proof Rules for Non-Termination. We present the (syntactic) proof rules for divergence in Fig. 2. Recall from §2 that $\left[p\right]$ C $[\infty]$ states that every state in p leads to *some* divergent trace. We provide the formal semantic interpretation of divergent triples later in §5.

Note that skip, assignment, error and assume statements never diverge. In order to show that $\mathsf{C}_1;\mathsf{C}_2$ has a divergent trace starting from p, we can show either C_1 has a divergent trace starting from p (Div-Seq1), or C_1 terminates normally transforming the states to q and C_2 does not terminate starting from q (Div-Seq2). To show that the branching program $\mathsf{C}_1 + \mathsf{C}_2$ has a divergent trace

$$\textsc{Div-Seq1}\ \frac{\vdash [p]\ \mathsf{C}_1\ [\infty]}{\vdash [p]\ \mathsf{C}_1;\mathsf{C}_2\ [\infty]} \qquad \textsc{Div-Seq2}\ \frac{\vdash_{\mathsf{B}} [p]\ \mathsf{C}_1\ [ok: q] \quad \vdash [q]\ \mathsf{C}_2\ [\infty]}{\vdash [p]\ \mathsf{C}_1;\mathsf{C}_2\ [\infty]} \qquad \textsc{Div-Choice}\ \frac{\vdash [p]\ \mathsf{C}_i\ [\infty] \quad \text{for some } i \in \{1,2\}}{\vdash [p]\ \mathsf{C}_1 + \mathsf{C}_2\ [\infty]}$$

$$\textsc{Div-LoopUnfold}\ \frac{\vdash [p]\ \mathsf{C};\mathsf{C}^\star\ [\infty]}{\vdash [p]\ \mathsf{C}^\star\ [\infty]} \qquad \textsc{Div-Loop}\ \frac{\vdash_{\mathsf{B}} [p]\ \mathsf{C}\ [ok: q] \quad q \subseteq p}{\vdash [p]\ \mathsf{C}^\star\ [\infty]} \qquad \textsc{Div-Subvar}\ \frac{\forall n \in \mathbb{N}.\ \vdash_{\mathsf{B}} [p(n)]\ \mathsf{C}\ [ok: p(n{+}1)]}{\vdash [p(0)]\ \mathsf{C}^\star\ [\infty]}$$

$$\textsc{Div-Local}\ \frac{\vdash [p]\ \mathsf{C}\ [\infty]}{\vdash [\exists x.p]\ \text{local } x \text{ in } \mathsf{C}\ [\infty]} \qquad \textsc{Div-Subst}\ \frac{\vdash [p]\ \mathsf{C}\ [\infty] \quad y \notin \mathsf{fv}(p,\mathsf{C})}{\vdash ([p]\ \mathsf{C}\ [\infty])[y/x]} \qquad \textsc{Div-Cons}\ \frac{\vdash [p']\ \mathsf{C}\ [\infty] \quad p \subseteq p'}{\vdash [p]\ \mathsf{C}\ [\infty]}$$

$$\textsc{Div-Disj}\ \frac{\vdash [p_i]\,\mathsf{C}\,[\infty] \quad \text{for all } i \in I}{\vdash \left[\bigvee_{i\in I} p_i\right]\ \mathsf{C}\ [\infty]} \qquad \textsc{Div-LoopNest}\ \frac{\vdash [p]\ \mathsf{C}\ [\infty]}{\vdash [p]\,\mathsf{C}^\star\,[\infty]} \qquad \textsc{Div-While}\ \frac{\vdash_{\mathsf{B}} [p \wedge B]\,\mathsf{C}\,[ok: q \wedge B] \quad q \subseteq p}{\vdash [p \wedge B]\ \text{while } (B)\ \mathsf{C}\ [\infty]}$$

$$\textsc{Div-WhileNest}\ \frac{\vdash [p \wedge B]\ \mathsf{C}\ [\infty]}{\vdash [p \wedge B]\ \text{while } (B)\ \mathsf{C}\ [\infty]} \qquad \textsc{Div-WhileSubvar}\ \frac{\forall n \in \mathbb{N}.\ \vdash_{\mathsf{B}} [p(n) \wedge B]\,\mathsf{C}\,[ok: p(n{+}1) \wedge B]}{\vdash [p(0) \wedge B]\ \text{while } (B)\ \mathsf{C}\ [\infty]}$$

Fig. 2. The UNTer divergence rules, where the highlighted rules can be derived from other rules

starting from p, it suffices to show that *some* branch C_i has a divergent trace from p, i.e. in an under-approximate fashion. The Div-Cons denotes the rule of consequence for divergence: if C has some divergent trace starting from any state in p' and $p \subseteq p'$, then C also has some divergent trace starting from any state in p. The Div-Local rule states that variable declaration diverges if its body does. The Div-Disj rule denotes that if the states in each of p_i lead to divergence, then so do the states in their union. The Div-Subst rule is the substitution rule for divergence and is as expected.

The remaining rules capture divergence for loops. Specifically, Div-LoopUnfold allows us to establish divergence after unrolling the loop once. This can be used for showing divergence in the case of nested loops, where the inner loop diverges. Specifically, using a combination of Div-Seq1 and Div-LoopUnfold we can derive Div-LoopNest as shown across, stating that if one iteration of the loop body (e.g. a nested loop) has a divergent trace, then the loop itself also has a divergent trace.

$$\cfrac{\cfrac{\cfrac{}{[p]\ \mathsf{C}\ [\infty]}\ \text{(given)}}{[p]\ \mathsf{C};\mathsf{C}^\star\ [\infty]}\ \text{(Div-Seq1)}}{[p]\ \mathsf{C}^\star\ [\infty]}\ \text{(Div-LoopUnfold)}$$

The Div-Loop rule states that if one iteration of a loop body terminates normally and transforms the states in p to ones in q (i.e. $\vdash_{\mathsf{B}} [p]\ \mathsf{C}\ [ok: q]$) and $q \subseteq p$, then $\mathsf{C}^\star$ has a divergent trace starting from p. Intuitively, the forward triple in the premise, $A \triangleq \vdash_{\mathsf{B}} [p]\ \mathsf{C}\ [ok: q]$, allows us to construct an infinite trace of $\mathsf{C}^\star$ from any state in p: given a state in $s_0 \in p$, (from A) executing C on s_0 results in a state $s_1 \in q \subseteq p$, and thus (from A) executing C on s_1 results in a state $s_2 \in q \subseteq p$, *ad infinitum*.

The Div-Subvar is the subvariant rule for divergence: if an iteration of the loop body terminates normally and transforms $p(n)$ to $p(n{+}1)$ for an arbitrary n, then $\mathsf{C}^\star$ has a divergent trace starting from the initial states $p(0)$. Note that given any loop body C, if C does not contain a conditional (if or while) statement and executing C does not encounter an error, then the non-deterministic loop $\mathsf{C}^\star$ always has a divergent trace. However, this is not necessarily the case with conditional if/while

while $(x = 0)$ skip (a)	while $(x \geq 0)$ $x := x{+}1$ (b)	$x := 1$ $y := 2$; while $(x{+}y > 1)$ $x := 3 - x$ $y := 3 - y$ (c)	while $(y < 100)$ if $(y \leq 50)$ $x := x{+}1$ else $y := y{+}1$ (d)	while $(y < 100)$ $x := 0$; while $(x \leq 100)$ if $(x = 100)$ $y := 0$ $x := x{+}1$ $y := y{+}1$ (e)	$x := 42$; $y := 1$; while $(y < 100)$ while $(x \leq 100)$ if $(x = 100)$ $x := 1$ $y := 2 \times y$ else $x := x{+}1$ $y := y{+}1$ (f)

Fig. 3. Several examples of programs with non-terminating behaviours where x, y initially hold 0

statements (encoded via assume). This is illustrated in Div-While, requiring that the loop condition B hold at the end of an iteration, which is not always the case. For instance, for while $(x = 0)\ x := 1$ we fail to establish $x = 0$ after an iteration of $x := 1$. The Div-WhileNest and Div-WhileSubvar rules are analogous to Div-LoopNest and Div-Subvar, respectively. As before, all highlighted rules in Fig. 2 can be derived from other rules (see the extended version [Raad et al. 2024a, §B]).

4 Examples

We present several simple examples of divergent programs (with divergent loops) and demonstrate how we can use our UNTer proof system to detect them. All divergent behaviours presented here, and many more, have also been detected using our Pulse$^\infty$ prototype (see §7).

Example 1 (Fig. 3a). Consider the simple example in Fig. 3a comprising a simple divergent loop. We can detect this using Div-While (with $p = q = \text{true}$) as shown below:

$$\dfrac{\dfrac{}{\vdash_B \left[x = 0\right] \text{skip} \left[ok\colon x = 0\right]}\ (\textsc{Skip})}{\left[x = 0\right] \text{while } (x = 0) \text{ skip } [\infty]}\ (\textsc{Div-While})$$

Example 2 (Fig. 3b). Consider the simple example in Fig. 3b comprising a simple while loop with a buggy check. We can detect this using Div-While (with $p = \text{true}$ and $q = x \geq 1$) as shown below:

$$\dfrac{\dfrac{\dfrac{}{\vdash_B \left[x \geq 0\right] x := x{+}1 \left[ok\colon \exists v.\ v \geq 0 \wedge x = v{+}1\right]}\ (\textsc{Assign})}{\vdash_B \left[x \geq 0\right] x := x{+}1 \left[ok\colon x \geq 1 \wedge x \geq 0\right]}\ (\textsc{ConsEq}) \quad \dfrac{}{x \geq 1 \subseteq x \geq 0}}{\left[x \geq 0\right] \text{while } (x \geq 0)\ x := x{+}1\ [\infty]}\ (\textsc{Div-While})$$

Example 3 (Fig. 3c). Consider the example in Fig. 3c. Prior to the first iteration of the loop $x{+}y = 3$ holds, and although the values of x and y are updated in each iteration, their sum remains unchanged after each iteration (i.e. $x{+}y = 3$) and thus the loop diverges. We present an UNTer proof outline of this divergent behaviour on the left of Fig. 4. For brevity, rather than giving full derivations, we follow the classical Hoare logic proof outline, annotating each line of the code with its pre- and post-states. We further commentate each proof step and write e.g. // Assign to denote an application of Assign. As in Hoare logic proof outlines, we assume that Seq is applied at every step; i.e. later instructions are executed only if the earlier ones execute normally (with *ok*).

Let $p \triangleq x{+}y = 3 \wedge x{+}y > 1$; after the initial assignment to x and y and applications of ConsEq and Div-Cons, we establish p (line 6). We then apply Div-While (lines 6–14) to show that the loop body leaves the set of states p unchanged (lines 8–13). The proof of lines 8–13 is then straightforward, and simply involves the applications of Assign and ConsEq.

Example 4 (Fig. 3d). Consider the example in Fig. 3d. At first glance it may seem that the loop terminates since the value of y is incremented in the else branch of each iteration. However, starting from $y = 0$, the then branch is taken in each iteration (since $y \leq 50$) and thus y is never incremented,

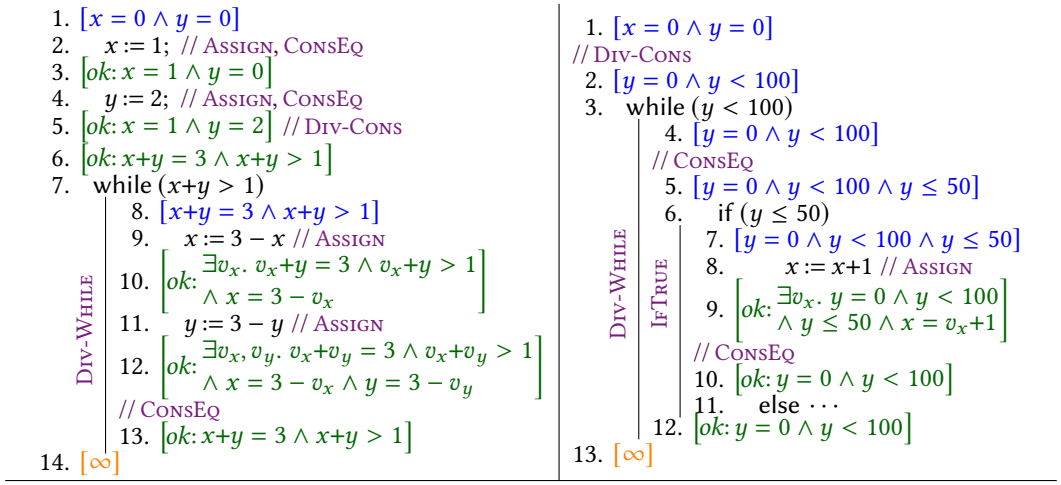


Fig. 4. Proof sketches of the divergence bugs in Fig. 3c (left) and Fig. 3d (right)

resulting in divergence. We present an UNTer proof outline of this divergent behaviour on the right of Fig. 4. After applying ConsEq to rewrite p equivalently as $p \wedge y \leq 50$ (line 5), we apply IfTrue to show we can take the then branch and arrive at p (lines 7–10).

Example 5 (Fig. 3e). Consider the example in Fig. 3e with nested loops. Note that the value of x is incremented at the end of each iteration of the inner loop and thus the inner loop terminates. By contrast, although y is incremented at the end of each iteration of the outer loop and thus it may seem at first glance that the outer loop terminates, on closer inspection the value of y us reset to 0 in the last iteration of the inner loop. As such, at the end of each iteration of the outer loop y is incremented and updated 1, and thus the outer loop diverges.

We present an UNTer proof outline of this at the top of Fig. 5. After applying Div-Cons to obtain $y < 100$, we apply Div-While (lines 2–23) to show that the loop body leaves $y < 100$ unchanged (lines 4–22). After the assignment on line 5, we apply ConsEq to rewrite the states as $p(0) \wedge x \leq 100$ (line 7), with $p(n)$ defined below the proof at the top of Fig. 5. We then apply WhileSubvar to show that at the end of the execution of the inner loop we arrive at $y = 0 \wedge x = 101 \wedge x \not\leq 100$ (lines 7–21). Note that WhileSubvar has two premises, which we establish in two columns on lines 9–14 and 15–20. On lines 9–14 we show that for $n < 100$, each iteration of the loop transforms $p(n) \wedge x \leq 100$ to $p(n{+}1) \wedge x \leq 100$; on lines 15–20 we show that in the final iteration of the loop with $p(100)$ (i.e. when $x = 100$), we reset y to 0 and increment x, arriving at $y = 0 \wedge x = 101 \wedge x \not\leq 100$ which is included in $y < 100$ (line 22), as per the second premise of Div-While.

Example 6 (Fig. 3f). Consider the nested loops in Fig. 3f. Note that starting with $x = 42$ (after the initial assignment), the else branch of the inner loop increments x in all but the last iteration of the inner loop (since $x = 100$), whereupon the value of x is reset to 1; i.e. the inner loop diverges.

We present an UNTer proof outline of this divergent behaviour at the bottom of Fig. 5. After the initial assignments (line 2) and applying Div-Cons to arrive at $x \leq 100 \wedge y < 100$ (line 4), we apply Div-WhileNest (lines 4–21) to show that the loop body diverges (lines 6–20). Once again, we apply Div-Cons to weaken the states to $x \leq 100$ (line 7) and subsequently apply Div-While (lines 7–20) to show that the body of the inner loop leaves the states $x \leq 100$ unchanged (lines 9–19). To do this, we first rewrite $x \leq 100$ equivalently as $x < 100 \vee x = 100$ (line 10), and then apply Disj to show that either disjunct results in $x \leq 100$ states (the two columns on lines 11–14 and 15–18). The proof of each disjunct is then straightforward and is obtained by reasoning about the associated branch.

1. $[x = 0 \wedge y = 0]$ // Div-Cons
2. $[y < 100]$
3. while $(y < 100)$
   4. $[y < 100]$
   5. $x := 0$ // Assign
   6. $[ok: y < 100 \wedge x = 0]$ // ConsEq
   7. $[ok: p(0) \wedge x \leq 100]$
   8. while $(x \leq 100)$
      9. $\forall n < 100.\ [p(n) \wedge n{<}100 \wedge x \leq 100]$ // WhileSubVar
      10. if $(x = 100)\ y := 0$
      11. else skip // IfFalse, Skip
      12. $[ok: p(n) \wedge n{<}100 \wedge x{\leq}100]$
      13. $x := x{+}1$ // Assign, ConsEq
      14. $[ok: p(n{+}1) \wedge x \leq 100]$
      15. $[p(100) \wedge x \leq 100]$
      16. if $(x = 100)\ y := 0$
      17. else skip // IfTrue, Assign
      18. $[ok: p(100) \wedge x \leq 100 \wedge y = 0]$
      19. $x := x{+}1$ // Assign, ConsEq
      20. $[ok: y = 0 \wedge x = 101 \wedge x \not\leq 100]$
   21. $[ok: y = 0 \wedge x = 101 \wedge x \not\leq 100]$
   22. $[ok: y < 100]$
23. $[\infty]$

Div-While

where for all $n \in \mathbb{N}$: $\quad p(n) \triangleq x = n \wedge y < 100$

1. $[x = 0 \wedge y = 0]$
2. $x := 42;\ y := 1;$ // Assign, ConsEq
3. $[ok: x = 42 \wedge y = 1]$ // Div-Cons
4. $[ok: x \leq 100 \wedge y < 100]$
5. while $(y < 100)$
   6. $[x \leq 100 \wedge y < 100]$ // Div-Cons
   7. $[x \leq 100]$
   8. while $(x \leq 100)$
      9. $[x \leq 100]$ // ConsEq
      10. $[x < 100 \vee x = 100]$ // Disj
      11. $[x < 100]$
      12. if $(x = 100)\ x := 1;\ y := 2{\times}y$
      13. else $x := x{+}1$
      // IfFalse, Assign, ConsEq
      14. $[ok: x \leq 100]$
      15. $[x = 100]$
      16. if $(x = 100)\ x := 1;\ y := 2{\times}y$
      17. else $x := x{+}1$
      // IfTrue, Assign, ConsB
      18. $[ok: x \leq 100]$
      19. $[ok: x \leq 100]$
   20. $[\infty]$
21. $[\infty]$

Div-WhileNest

Div-While

Disj

Fig. 5. Proof sketch of divergence in Fig. 3e (above), where the two columns on lines 9–14 and 15–20 denote the proof sketches of the two premises of WhileSubVar; proof sketch of divergence in Fig. 3f (below), where the two columns on lines 11–14 and 15–18 denote the proof sketches of the two premises of Disj.

5 The UNTer Model and Semantics

Instrumented Commands and Operational Semantics. Although in sequential settings the semantics is given in the big-step fashion [O'Hearn 2019; Raad et al. 2020], we opt for *small-step* semantics instead. This is because big-step semantics by definition describe *terminating* executions, while our aim is to formalise the semantics of divergent triples. Specifically, as we describe below, we formalise the semantics of a divergent triple as an *infinite*, non-terminating execution trace.

Note that local x in C declares a variable x whose scope is limited to C. To describe the semantics of local x in C in a small-step fashion, we introduce *instrumented commands*, defined by the grammar below (where C is as defined in §3), which additionally include the $\mathsf{end}(x, v)$ construct, recording the existing (old) value of x when redeclaring x in a new scope.

$$\mathbb{C} ::= \mathsf{C} \mid \mathsf{end}(x, v) \mid \mathbb{C}_1; \mathbb{C}_2$$

$$\textsc{S-Local}\ \frac{s' = s[x \mapsto v] \quad v \in \textsc{Val}}{\mathsf{local}\ x\ \mathsf{in}\ \mathsf{C}, s \longrightarrow \mathsf{C}; \mathsf{end}(x, s(x)), s', \mathit{ok}} \qquad \textsc{S-LocalEnd}\ \frac{s' = s[x \mapsto v]}{\mathsf{end}(x, v), s \longrightarrow \mathsf{skip}, s', \mathit{ok}} \qquad \textsc{S-Assign}\ \frac{s' = s[x \mapsto s(e)]}{x := e, s \longrightarrow \mathsf{skip}, s', \mathit{ok}}$$

$$\textsc{S-Assume}\ \frac{s(B) = \mathsf{true}}{\mathsf{assume}(B), s \longrightarrow \mathsf{skip}, s, \mathit{ok}} \qquad \textsc{S-Error}\ \mathsf{error}, s \longrightarrow \mathsf{skip}, s, \mathit{er} \qquad \textsc{S-Choice}\ \frac{i \in \{1, 2\}}{\mathsf{C}_1 + \mathsf{C}_2, s \longrightarrow \mathsf{C}_i, s, \mathit{ok}} \qquad \textsc{S-Seq1}\ \frac{\mathbb{C}_1, s \longrightarrow \mathbb{C}_1', s', \epsilon}{\mathbb{C}_1; \mathbb{C}_2, s \longrightarrow \mathbb{C}_1'; \mathbb{C}_2, s', \epsilon}$$

$$\textsc{S-SeqSkip}\ \mathsf{skip}; \mathbb{C}, s \longrightarrow \mathbb{C}, s, \mathit{ok} \qquad \textsc{S-Loop0}\ \mathsf{C}^\star, s \longrightarrow \mathsf{skip}, s, \mathit{ok} \qquad \textsc{S-Loop}\ \mathsf{C}^\star, s \longrightarrow \mathsf{C}; \mathsf{C}^\star, s, \mathit{ok}$$

$$\frac{\mathbb{C} \in \{\mathsf{local}\ x\ \mathsf{in}\ \mathsf{C}, x := e, \mathsf{assume}(B), \mathsf{error}, \mathsf{C}_1 + \mathsf{C}_2, \mathsf{C}^\star\}}{\mathbb{C}, s \rightsquigarrow_{\mathsf{er}} \mathsf{skip}, s} \qquad \frac{\mathbb{C}_1, s \rightsquigarrow_{\mathsf{er}} \mathbb{C}_1', s'}{\mathbb{C}_1; \mathbb{C}_2, s \rightsquigarrow_{\mathsf{er}} \mathbb{C}_1'; \mathbb{C}_2, s'}$$

$$\mathsf{end}(x, v), s \rightsquigarrow_{\mathsf{er}} \mathsf{skip}, s[x \mapsto v] \qquad \mathsf{skip}; \mathbb{C}, s \rightsquigarrow_{\mathsf{er}} \mathbb{C}, s$$

Fig. 6. The UNTer small-step transitions (above) and error transitions for restoring variables (below)

We present our small-step semantics in Fig. 6, with transitions of the form $\mathbb{C}, s \longrightarrow \mathbb{C}', s', \epsilon$, where $\mathbb{C}$ and s respectively denote the current (instrumented) command and store (state), $\mathbb{C}'$ and s' denote their continuations (what they reduce to) and ϵ denotes the exit condition, describing whether reducing $\mathbb{C}$ to $\mathbb{C}'$ took place normally (*ok*) or erroneously (*er*). As shown in S-Local, when evaluating local x in C under a state $s \in \textsc{Store}$, we assign an arbitrary value v to x in s, and continue with executing C followed by $\mathsf{end}(x, s(x))$. That is, we record the existing value of x, $s(x)$, so that we can restore it once the execution of C has ended, as reflected in the S-LocalEnd transition.

The remaining transition rules are standard: assigning e to x simply evaluates e in the current state (denoted by $s(e)$) and updates the value of x in the state, terminating normally; $\mathsf{assume}(B)$ reduces to skip normally when B evaluates to true in the current state; error reduces to skip erroneously; and $\mathsf{C}_1 + \mathsf{C}_2$ non-deterministically reduces to one of its branches (C_i with $i \in \{1, 2\}$). When reducing $\mathbb{C}_1; \mathbb{C}_2$, we either reduce the left-hand side until it reduces to skip (S-Seq1), or continue with the right-hand side when the left side is skip (S-SeqSkip). Finally, we either reduce a loop to skip, i.e. unroll it zero times (S-Loop0), or unroll it once and continue with $\mathsf{C}^\star$ (S-Loop).

Semantic BUA and FUA Triples. Recall that intuitively a BUA triple $\vdash_{\mathsf{B}} [p]\ \mathsf{C}\ [\epsilon : q]$ states that every pre-state s_p in p can reach some post-state s_q in q under ϵ by executing C. Analogously, a FUA triple $\vdash_{\mathsf{F}} [p]\ \mathsf{C}\ [\epsilon : q]$ states that every post-state s_q in q can be reached from some pre-state s_p in p under ϵ by executing C. Put formally, in both cases we must have $\mathsf{C}, s_p \xrightarrow{n} -, s_q, \epsilon$, denoting that executing C *terminates* after n steps under ϵ and transforms s_p to s_q (see Def. 1 below).

Definition 1 (Semantic BUA and FUA triples)**.** A BUA triple is *valid*, written $\models_{\mathsf{B}} [p]\ \mathsf{C}\ [\epsilon : q]$, iff for all $s_p \in p$, there exists $s_q \in q$ and n such that $\mathsf{C}, s_p \xrightarrow{n} -, s_q, \epsilon$, where:

$$\begin{aligned}\mathbb{C}, s \xrightarrow{n} \mathbb{C}', s', \epsilon \ \overset{\text{def}}{\iff}\ & (n{=}0 \wedge \mathbb{C}{=}\mathbb{C}'{=}\mathsf{skip} \wedge s{=}s' \wedge \epsilon{=}\mathit{ok}) \\ & \vee\ (n{=}1 \wedge \epsilon \in \textsc{ErExit} \wedge \exists s''.\ \mathbb{C}, s \longrightarrow \mathbb{C}', s'', \epsilon \wedge \mathbb{C}', s'' \rightsquigarrow_{\mathsf{er}}^{+} \mathsf{skip}, s') \\ & \vee\ (\exists k, \mathbb{C}'', s''.\ n{=}k{+}1 \wedge \mathbb{C}, s \longrightarrow \mathbb{C}'', s'', \mathit{ok} \wedge \mathbb{C}'', s'' \xrightarrow{k} \mathbb{C}', s', \epsilon)\end{aligned}$$

and $\mathbb{C}, s \longrightarrow \mathbb{C}', s', \epsilon$ is the UNTer small-step transitions given at the top of Fig. 6, while $\rightsquigarrow_{\mathsf{er}}^{+}$ denotes the transitive closure of the error transitions $\rightsquigarrow_{\mathsf{er}}$ as defined at the bottom of Fig. 6 (described shortly). A FUA triple is *valid*, written $\models_{\mathsf{F}} [p]\ \mathsf{C}\ [\epsilon : q]$, iff for all $s_q \in q$, there exists $s_p \in p$ and n such that $\mathsf{C}, s_p \xrightarrow{n} -, s_q, \epsilon$.

The first disjunct in $\mathbb{C}, s \xrightarrow{n} \mathbb{C}', s', \epsilon$ denotes that a state is reached under *ok* in zero steps without changing the underlying state, provided that $\mathbb{C}$ is simply skip. The last disjunct captures the inductive case (n=k+1), where $\mathbb{C}$ takes an *ok* step, and s' is subsequently reached in k steps under ϵ.

Finally, the second disjunct captures the short-circuiting semantics of errors: a state s' is reached in one step under *er* when $\mathbb{C}$ takes an erroneous step, whereupon locally declared variables (through local) are restored to their oldest values (outer-most scope) via $\rightsquigarrow_{\text{er}}$ transitions (defined in Fig. 6). The $\rightsquigarrow_{\text{er}}$ transitions 'skip over' the execution of most commands and restore the value of a variable x when encountering $\text{end}(x, v)$. Specifically, the $\rightsquigarrow_{\text{er}}$ transitions of all commands, except those of $\text{end}(x, v)$ and sequential composition, do not change the underlying state and simply reduce to skip (i.e. ignore their effects), while the $\rightsquigarrow_{\text{er}}$ transition for $\text{end}(x, v)$ restores the value of x to v. The $\rightsquigarrow_{\text{er}}$ transitions for sequential composition are defined inductively as expected.

We next show that the BUA and FUA proof systems presented in Fig. 1 are *both sound and complete*, with the full proof given in the extended version [Raad et al. 2024a, §C.1, §D.1]).

Theorem 7 (BUA and FUA soundness). *For all p, q, C and ϵ:*

1) if $\vdash_{\text{B}} [p]\ \text{C}\ [\epsilon : q]$ is derivable using the rules in Fig. 1, then $\models_{\text{B}} [p]\ \text{C}\ [\epsilon : q]$ holds; and
2) if $\vdash_{\text{F}} [p]\ \text{C}\ [\epsilon : q]$ is derivable using the rules in Fig. 1, then $\models_{\text{F}} [p]\ \text{C}\ [\epsilon : q]$ holds.

Theorem 8 (BUA and FUA completeness). *For all p, q, C and ϵ:*

1) if $\models_{\text{B}} [p]\ \text{C}\ [\epsilon : q]$ holds, then $\vdash_{\text{B}} [p]\ \text{C}\ [\epsilon : q]$ is derivable using the rules in Fig. 1; and
2) if $\models_{\text{F}} [p]\ \text{C}\ [\epsilon : q]$ holds, then $\vdash_{\text{F}} [p]\ \text{C}\ [\epsilon : q]$ is derivable using the rules in Fig. 1.

Definition 2 (Semantic divergent triples). A divergent triple is *valid*, written $\models \lfloor p \rfloor\ \text{C}\ [\infty]$, iff for all $s \in p$, there exists an infinite series of $\mathbb{C}_1, \mathbb{C}_2, \cdots$, $s_1, s_2, \cdots$ and $n_1, n_2, \cdots$ such that $\text{C}, s \rightsquigarrow^{n_1} \mathbb{C}_1, s_1, ok \rightsquigarrow^{n_2} \mathbb{C}_2, s_2, ok \rightsquigarrow^{n_3} \cdots$, where the chain $\text{C}, s \rightsquigarrow^{n_1} \mathbb{C}_1, s_1, ok \rightsquigarrow^{n_2} \mathbb{C}_2, s_2, ok \rightsquigarrow^{n_3} \cdots$ is a shorthand for $\text{C}, s \rightsquigarrow^{n_1} \mathbb{C}_1, s_1, ok \wedge \mathbb{C}_1, s_1 \rightsquigarrow^{n_2} \mathbb{C}_2, s_2, ok \wedge \cdots$, and $\rightsquigarrow^n$ is defined as follows:

$$\begin{aligned}\mathbb{C}, s \rightsquigarrow^n \mathbb{C}', s', \epsilon \overset{\text{def}}{\iff}\ & (n = 1 \wedge \epsilon = ok \wedge \mathbb{C}, s \rightarrow \mathbb{C}', s', \epsilon) \\ & \vee (n = 1 \wedge \epsilon \in \textsc{ErExit} \wedge \exists s''.\ \mathbb{C}, s \rightarrow \mathbb{C}', s'', \epsilon \wedge \mathbb{C}', s'' \rightsquigarrow^{+}_{\text{er}} \text{skip}, s') \\ & \vee (\exists k, s'', \mathbb{C}''.\ n{=}k{+}1 \wedge \mathbb{C}, s \rightarrow \mathbb{C}'', s'', ok \wedge \mathbb{C}'', s'' \rightsquigarrow^{k} \mathbb{C}', s', \epsilon)\end{aligned}$$

Note that unlike the $\mathbb{C}, s \xrightarrow{n} \mathbb{C}', s'$ transitions in Def. 1 which describe *terminating* traces (by reduction to skip), the $\mathbb{C}, s \rightsquigarrow^n \mathbb{C}', s'$ transitions do not stipulate termination and simply state that executing $\mathbb{C}$ from s for n steps reduces to $\mathbb{C}'$ and results in s'.

We next show that the divergence proof system presented in Fig. 2 is *both sound and complete*, with the full proof given in the extended version [Raad et al. 2024a, §C.2 and §D.2].

Theorem 9 (Divergence soundness and completeness). *For all p and C, if $\vdash \lfloor p \rfloor\ \text{C}\ [\infty]$ is derivable using the rules in Fig. 2, then $\models \lfloor p \rfloor\ \text{C}\ [\infty]$ holds.*
For all p and C, if $\models \lfloor p \rfloor\ \text{C}\ [\infty]$ holds, then $\vdash \lfloor p \rfloor\ \text{C}\ [\infty]$ is derivable using the rules in Fig. 2.

Finally, we formalise the relationship between FUA and BUA triples (see p. 8), with the proof in the extended version [Raad et al. 2024a, §E].

Theorem 10. *For all p, C, q, ϵ:*

1) if $\models_{\text{F}} [p]\ \text{C}\ [\epsilon : q]$ and $\text{min}_{\text{pre}}(p, \text{C}, q)$ hold, then $\models_{\text{B}} [p]\ \text{C}\ [\epsilon : q]$ also holds; and
2) if $\models_{\text{B}} [p]\ \text{C}\ [\epsilon : q]$ and $\text{min}_{\text{post}}(p, \text{C}, q)$ hold, then $\models_{\text{F}} [p]\ \text{C}\ [\epsilon : q]$ also holds, where:

$$\text{min}_{\text{pre}}(p, \text{C}, q) \overset{def}{\iff} \forall p'.\ p' \subset p \Rightarrow \not\models_{\text{F}} \lfloor p' \rfloor\, \text{C}\, [\epsilon : q] \qquad \text{min}_{\text{post}}(p, \text{C}, q) \overset{def}{\iff} \forall q'.\ q' \subset q \Rightarrow \not\models_{\text{B}} [p]\, \text{C}\, [\epsilon : q']$$

AssignSL
$\vdash_\dagger [x{=}x']\ x := e\ [ok : x{=}e[x'/x]]$

Store
$\vdash_\dagger [x \mapsto e]\ [x] := y\ [ok : x \mapsto y]$

StoreEr
$\vdash_\dagger [x \not\mapsto]\ [x] := y\ [er : x \not\mapsto]$

StoreNull
$\vdash_\dagger [x{=}\mathsf{null}]\ [x] := y\ [er : x{=}\mathsf{null}]$

Frame
$$\frac{\vdash_\dagger [p]\ \mathsf{C}\ [\epsilon : q] \quad \mathsf{mod}(\mathsf{C}) \cap \mathsf{fv}(r){=}\emptyset}{\vdash_\dagger [p * r]\ \mathsf{C}\ [\epsilon : q * r]}$$

Div-Frame
$$\frac{\vdash [p]\ \mathsf{C}\ [\infty]}{\vdash [p * r]\ \mathsf{C}\ [\infty]}$$

Fig. 7. UNTer$^{\mathsf{SL}}$ proof rules (excerpt), where x and x' are distinct variables and $\dagger$ in each rule can be instantiated as F or B; see the extended version for the full set of UNTer$^{\mathsf{SL}}$ rules [Raad et al. 2024a, §F].

Note that while the theoretical result in Theorem 10 does not have an immediate practical impact, it nevertheless reconciles FUA and BUA reasoning and shows how we can use tools such as Pulse that are underpinned by FUA to detect non-termination. Specifically, $\mathsf{min}_{\mathsf{pre}}$ describes a minimal precondition that has not been arbitrarily weakened (grown) using the ConsF rule. Similarly, $\mathsf{min}_{\mathsf{post}}$ describes a minimal postcondition that has not been arbitrarily weakened using ConsB. As such, given a set S of FUA triples inferred by a FUA-based analysis tool such as Pulse, the triples in S can be soundly interpreted as BUA ones (and therefore used to prove divergence), provided that their preconditions have not been weakened using ConsF, which is indeed the case in Pulse.

6 Extension to Separation Logic

We describe how we develop UNTer$^{\mathsf{SL}}$ by extending UNTer with the compositional reasoning principles of separation logic (SL) [Ishtiaq and O'Hearn 2001]. Raad et al. [2020] have developed incorrectness separation logic (ISL) by extending the FUA-based incorrectness logic (IL) [O'Hearn 2019] with separation logic. We adopt the model of Raad et al. [2020] and show that it is also sound for BUA reasoning.

UNTer$^{\mathsf{SL}}$ Programming Language and Assertions. To account for operations that access the heap, in UNTer$^{\mathsf{SL}}$ we extend our programming language from §3 with the following heap-manipulating operations (below, left) for allocation ($x := \mathsf{alloc}()$), deallocation ($\mathsf{free}(x)$), reading from the heap (lookup, $x := [y]$) and writing to the heap (mutation, $[x] := y$). We similarly extend the UNTer assertions as follows (below, right) by adding structural assertions to describe heaps.

$$\begin{aligned} \mathsf{Comm} \ni \mathsf{C} ::= {} & \cdots \mid x := \mathsf{alloc}() \mid \mathsf{free}(x) \\ & \mid x := [y] \mid [x] := y \end{aligned} \qquad \begin{aligned} \mathsf{Ast} \ni p, q, r ::= {} & \cdots \mid \mathsf{emp} \mid e \mapsto e' \\ & \mid e \not\mapsto \mid p * q \end{aligned}$$

The UNTer$^{\mathsf{SL}}$ assertions describe sets of *states*, where a state comprises a (variable) store and a heap. Existing UNTer assertions from §3 then simply describe states where the heap is empty and the store satisfies the assertion. The structural assertions above are those of ISL [Raad et al. 2020], which themselves are standard SL assertions [Ishtiaq and O'Hearn 2001] extended with $e \not\mapsto$. The $e \not\mapsto$ describes states where the heap comprises a single location at e containing the designated value $\bot$. In particular, whilst $e \mapsto e'$ states that location e is allocated (and contains value e'), $e \not\mapsto$ states that location e is *deallocated*.

UNTer$^{\mathsf{SL}}$ Proof Rules (Syntactic UNTer$^{\mathsf{SL}}$ Triples). We present an excerpt of the UNTer$^{\mathsf{SL}}$ proof rules in Fig. 7; see the extended version for the full set of rules [Raad et al. 2024a, §F]. Note that all UNTer rules (both BUA and FUA) in Fig. 1, except Constancy and Assign, are also UNTer$^{\mathsf{SL}}$ rules and are omitted from Fig. 7. In particular, we replace Constancy with the more powerful Frame rule and give a *local* rule for assignment (see below). As with ISL (and in contrast to UNTer), UNTer$^{\mathsf{SL}}$ triples are *local* in that their pre-states only contain the resources needed by the program. For instance, as assignment requires no heap resources, as shown in AssignSL the pre-state of skip

1. $[input[0] \mapsto 0 * size > 0 * off = v * newoff = - * i = -]$
2. $off := 0$; // AssignSL, Frame
3. $\lceil ok{:}\ input[0] \mapsto 0 * size > 0 * off = 0 * newoff = - * i = - \rceil$ // ConsEq
4. $\lceil ok{:}\ input[0] \mapsto 0 * size > 0 * off = 0 * newoff = - * i = - * off < size \rceil$
5. while $(off < size)$ — Div-While
   6. $[input[0] \mapsto 0 * size > 0 * off = 0 * newoff = - * i = - * off < size]$
   7. $newoff := [input[off]]$ // Load, Frame
   8. $\lceil ok{:}\ input[0] \mapsto 0 * size > 0 * off = 0 * newoff = 0 * i = - * off < size \rceil$
   9. $i := off$ // AssignSL, Frame
   10. $\lceil ok{:}\ input[0] \mapsto 0 * size > 0 * off = 0 * newoff = 0 * i = 0 * off < size \rceil$ // ConsEq
   11. $\lceil ok{:}\ input[0] \mapsto 0 * size > 0 * off = 0 * newoff = 0 * i = 0 * off < size * \neg(i < newoff) \rceil$
   12. while $(i < newoff)\ \{\cdots; i{+}{+}\}$ // WhileFalse
   13. $\lceil ok{:}\ input[0] \mapsto 0 * size > 0 * off = 0 * newoff = 0 * i = 0 * off < size * \neg(i < newoff) \rceil$
   14. $off := off + newoff$ // AssignSL, Frame, ConsEq
   15. $\lceil ok{:}\ input[0] \mapsto 0 * size > 0 * off = 0 * newoff = 0 * i = 0 * off < size \rceil$
16. $[\infty]$

Fig. 8. UNTer^{SL} proof sketch of CVE-2023-34966 in the Samba library (see Example 11)

is simply given by the pure (non-heap) assertion $x{=}x'$, recording the old value of x which can be used in the post-state.

As in SL and ISL, the crux of UNTer^{SL} lies in the Frame rule, allowing one to extend the pre- and post-states with disjoint resources in r, where $\mathsf{fv}(r)$ returns the set of free variables in r, and $\mathsf{mod}(C)$ returns the set of (program) variables modified by C (i.e. those on the left-hand of ':=' in assignment, lookup and allocation). These definitions are standard and elided. Heap manipulation rule are identical to those of ISL. For instance, Store describes a successful heap mutation, while StoreEr and StoreNull state that mutating x causes an error when x is deallocated or null, respectively.

The UNTer^{SL} divergent rules are those of UNTer in Fig. 2, except that the BUA UNTer triples in the premises (e.g. the first premise of Div-Seq2) are replaced with their UNTer^{SL} counterparts. Additionally, we extend the framing principle to divergent triples as shown in Div-Frame: if C diverges starting from the states in p, then it also diverges starting from the states in $p * r$.

We next use UNTer^{SL} to detect a known divergence bug in the Samba library, which has already been reported to the Common Vulnerabilities and Exposures (CVE) database as CVE-2023-34966.

Example 11 (Samba). The example in Fig. 8 is a stylised excerpt from the Samba library, where the body of `sl_unpack_loop` is repeated below. The excerpt shown reads chunks of data, where the size of each chunk is given by the corresponding entry in the *input* array (the size of the first chunk is stored in $input[0]$, the size of the second in $input[1]$ and so forth). The *off* records the offset at which next chunk to be read is stored and is initially set to zero (line 2). At each iteration of the outer while loop (lines 5–16), the size of the next chunk is read from *input* into *newoff* (line 7), and subsequently the offset is incremented by *newoff* (line 14). The inner while loop (line 12) then proceeds to read the data between *off* and *newoff* (elided here as $\cdots$) one unit at a time (incrementing i each time). Note that when $i = newoff = 0$, then this inner loop is never entered. Moreover, if $newoff = 0$, then the old offset is never incremented (i.e. the increment at line 14 is idempotent), and thus the loop never terminates. This is indeed the cause of divergence in this example, which has since been patched by simply adding a check at the beginning of the outer loop, ensuring that *newoff* is non-zero and returning an error value when that is the case. Using UNTer^{SL} we can detect this bug as shown in Fig. 8.

Table 1. Positioning Pulse$^\infty$ amongst (non)-termination analysis tools in the literature

Tool	Term.	Non-term.	Non-det.	Heap	Auto.	UA/OA	Large Code / Libraries
Terminator [Cook et al. 2006a,b]	✓	✗	✓	✗	✓	OA	✗
Mutant [Berdine et al. 2006]	✓	✗	✗	✓	✓	OA	✗
TNT [Gupta et al. 2008]	✗	✓	✗	✗	✓	OA	✗
KEY [Velroyen and Rümmer 2008]	✗	✓	✗	✗	✓	OA	✗
CPROVER [Kroening et al. 2010]	✓	✓	✓	✗	✓	UA-OA	✗
TRex [Harris et al. 2010]	✓	✓	✓	✗	✓	UA-OA	✗
T2 [Cook et al. 2013]	✓	✗	✓	✗	✓	OA	✗
Coop-T2 [Brockschmidt et al. 2013]	✓	✓	✓	✗	✓	UA-OA	✗
Caber [Brotherston and Gorogiannis 2014]	✓	✗	✓	✓	✓	OA	✗
CPP-INV [Larraz et al. 2014]	✗	✓	✓	✗	✓	OA	✗
HipTNT [Le et al. 2014]	✓	✓	✓	Partial	✗	OA	✗
HipTNT+ [Le et al. 2015]	✓	✓	✓	Partial	✓	OA	✗
DynamiTe [Le et al. 2020]	✓	✓	✓	Partial	✓	OA	✗
RevTerm [Chatterjee et al. 2021]	✗	✓	✓	✗	✓	OA	✗
AProVE [Hensel et al. 2022]	✗	✓	✓	✗	✓	OA	✗
Ultimate [Heizmann et al. 2014]	✓	✗	✓	✗	✓	OA	✗
Pulse$^\infty$	✗	✓	✓	✓	✓	UA	✓

UNTer^{SL} Semantics and Soundness. We present the UNTer^{SL} model in the extended version [Raad et al. 2024a, §F]. The formal interpretations of BUA, FUA and divergent triples in UNTer^{SL} are identical to their UNTer counterparts, except that the UNTer states are replaced with corresponding UNTer^{SL} states to account for heaps. We show that the BUA, FUA and divergent proof system of UNTer^{SL} are sound, with the full proof given in the extended version [Raad et al. 2024a, §G].

Theorem 12 (UNTer^{SL} soundness)**.** *The BUA, FUA and divergent proof system of UNTer^{SL} are sound.*

7 Evaluation

To demonstrate the feasibility of UNTer^{SL} for detecting divergence bugs, we have developed Pulse$^\infty$ as an extension of the existing Pulse program analyser (underpinned by the ISL [Raad et al. 2020] theory and compatible with both FUA and BUA reasoning). The most fundamental extensions to Pulse are: 1) addition of the divergent triple; 2) symbolic execution steps corresponding to proof rules for divergence; and 3) a stopping condition corresponding to fixpoints (which is unusual for under-approximation) and the associated "test oracle" for recognising divergence. In addition to these fundamental changes, we needed to make some detailed but conceptually minor alterations to the treatment of Booleans in Pulse: it had optimisations which were sound for proving violation of safety properties, but which interfered with our oracle for divergence.

7.1 Pulse$^\infty$ in Context

Table 1 places Pulse$^\infty$ in the context of other termination and non-termination tools within the last two decades, where **Non-det.** denotes whether the tool supports non-deterministic programming constructs such as `rand()`, **Auto.** denotes whether it is fully automated, and **UA/OA** denotes whether it performs under-approximate (UA) or over-approximate (OA) analysis. To our knowledge, Pulse$^\infty$ is the first tool for non-termination analysis with full support for heap reasoning. Tools such as Mutant [Berdine et al. 2006] and Caber [Brotherston and Gorogiannis 2014] are also capable of heap reasoning using separation logic, however these tools were specifically developed to perform over-approximate (OA) termination analysis, and do not support non-termination analysis. Other tools such as TNT [Gupta et al. 2008] and HipTNT [Le et al. 2014, 2015] have limited support for heap reasoning which does not include array or string capabilities.

Most importantly, as we discuss below, to our knowledge Pulse$^\infty$ is *the first tool* that can automatically prove non-termination on *large code bases and libraries* such as OpenSSL. As we discuss in §8, several existing tools [Brockschmidt et al. 2013; Chatterjee et al. 2021; Hensel et al. 2022; Kroening et al. 2010; Larraz et al. 2014; Le et al. 2020; Velroyen and Rümmer 2008], can only handle *integer C programs* and thus do not support programs *with function calls*. In other words, unlike Pulse$^\infty$, these tools cannot be run on large C codebases or libraries such as OpenSSL. In particular, Pulse$^\infty$ inherits all the *incremental* capabilities of Pulse (as documented on https://fbinfer.com/docs/next/steps-for-ci#differential-workflow). As a point of reference, analysing OpenSSL on 30 cores with Pulse takes 1m26s the first time. After modifying a file in the project, re-analysing OpenSSL again in incremental mode takes only 6s.

On this point we consulted with authors of the tools referenced in the table. None responded in the affirmative that they could run on large projects. Several tools [Le et al. 2014, 2015] stipulate the existence of a `main` procedure (thus excluding libraries), while others [Heizmann et al. 2014] require that the input to the tool be *a single C file*, and thus one must put the target program and all its dependencies into a single file (thus precluding large code bases and libraries). One theoretical possibility is to automatically construct a fake `main` procedure which calls methods in a representative fashion, but this would get us into a problem similar to *harness generation* in fuzzing, itself a challenging problem and a source of false positives. Our impression overall is that, far from being a simple matter, each of the other tools is one or several research projects away from automatic application at scale.

We evaluate Pulse$^\infty$ in two ways. First, in §7.2 we compare its ability to detect non-termination bugs on small examples and compare it against the state-of-the-art tools. Second, in §7.3 we run Pulse$^\infty$ on several large projects and libraries, comprising over 1.3 million lines of code. Note that as discussed above (and detailed later in §8), no existing automated tool for divergence analysis can be applied to large code bases or libraries *out of the box*, and Pulse$^\infty$ is the only such tool. As such, we could not compare the Pulse$^\infty$ performance against the state of the art for analysing libraries. We refer the reader to our project page for more detailed information about Pulse$^\infty$ and our benchmarks for evaluating it [Raad et al. 2024a].

Experimental Setup. We ran all our experiments below on a single server equipped with an AMD EPYC 7543P processor at 3.4 Ghz clock on 30 active cores (make -j 30) and 512GB of RAM.

7.2 Evaluating Pulse$^\infty$ on Small Examples

SV-COMP Benchmarks. To evaluate Pulse$^\infty$ against existing tools for detecting divergence, we focused on the state-of-the-art termination and divergence non-linear arithmetic benchmarks extending the Competition on Software Verification (SV-COMP) initiative. DynamiTe [Le et al. 2020] outperforms all other tools listed in Table 1. Specifically, at the time of publication, Le et al. [2020] demonstrated that DynamiTe outperformed AProVE [Hensel et al. 2022] for non-termination analysis, and AProVE was in turn shown to outperform all other pre-existing tools in Table 1. As such, rather than comparing Pulse$^\infty$ against all tools listed in Table 1, we compare it directly to DynamiTe and its non-linear arithmetic (NLA) benchmark extension distributed with SV-COMP.

We present our comparison result against DynamiTe in Table 2. The tests listed are rather small and contain at most three loops, with the majority of them comprising a single loop. Both DynamiTe and Pulse$^\infty$ analyse these tests within a few seconds, with the time difference being insignificant. As such, in Table 2 we do not compare the time-performance of the two tools, and focus instead on their reported outcomes. As shown, Pulse$^\infty$ performs competitively against DynamiTe, though the results are not directly comparable. Specifically, Pulse$^\infty$ successfully found divergence bugs that were undetected by DynamiTe (e.g. the dijkstraX-both-nt tests), while it missed others found by

Table 2. Comparing Pulse$^\infty$ against DynamiTe on SV-COMP non-linear arithmetic (NLA) benchmarks for termination (T) and non-termination (NT). For NT cases, ✓ denotes a true positive (i.e. the tool correctly detected non-termination) and FN a false negative (i.e. the tool failed to detect non-termination. For T cases, ? under Pulse$^\infty$ denotes an unknown result (Pulse$^\infty$ proves non-termination and not termination), and FP under DynamiTe denotes a false positive (i.e. DynamiTe incorrectly reported that the program terminates).

Test	T/NT	Pulse$^\infty$	DynamiTe	Test	T/NT	Pulse$^\infty$	DynamiTe
bresenham1-both-nt	NT	FN	✓	freire1-both-nt	NT	✓	FN
cohencu1-both-nt	NT	FN	✓	geo1-both-nt	NT	FN	✓
cohencu2-both-nt	NT	FN	FN	geo2-both-nt	NT	FN	✓
cohencu3-both-nt	NT	FN	FN	geo3-both-nt	NT	FN	✓
cohencu4-both-nt	NT	FN	FN	hard2-both-nt	NT	FN	✓
cohencu5-both-nt	NT	FN	✓	hard-both-nt	NT	✓	✓
dijkstra1-both-nt-2	NT	FN	FN	hard-both-t	T	?	FP
dijkstra1-both-nt	NT	✓	FN	knuth-both-nt	NT	FN	FN
dijkstra2-both-nt	NT	✓	FN	knuth-nosqrt-both-nt	NT	FN	✓
dijkstra3-both-nt	NT	✓	FN	lcm1-both-t	T	?	FP
dijkstra4-both-nt	NT	✓	FN	lcm1-both-nt	NT	✓	✓
dijkstra5-both-nt	NT	✓	FN	lcm2-both-nt	NT	✓	✓
dijkstra6-both-nt	NT	✓	FN	mannadiv-both-nt	NT	✓	FN
divbin1-both-nt	NT	FN	FN	prod4br-both-nt	NT	FN	FN
egcd2-both-nt	NT	✓	✓	prodbin-both-nt	NT	FN	FN
egcd3-both-t	T	?	FP	ps2-both-nt	NT	FN	✓
egcd3-both-nt	NT	✓	✓	ps3-both-nt	NT	FN	FN
egcd-both-nt	NT	✓	✓	ps4-both-nt	NT	FN	✓
fermat1-both-t	T	?	FP	ps5-both-nt	NT	FN	FN
fermat1-both-nt	NT	FN	FN	ps6-both-nt	NT	FN	FN
fermat2-both-nt	NT	✓	✓	sqrt1-both-nt	NT	FN	✓
fermat3-both-nt	NT	FN	✓	sqrt2-both-nt	NT	FN	FN
Total		**Pulse$^\infty$:** 15 ✓, 25 FN, 0 FP				**DynamiTe:** 19 ✓, 21 FN, 4 FP	

DynamiTe (e.g. the geoX-both-nt tests). Notably, thanks to its under-approximate nature, Pulse$^\infty$ reported *zero false positives* (FP), in contrast to DynamiTe which suffered several false positives.

From this evaluation we cannot conclude that Pulse$^\infty$ has superior precision on small examples compared to the state of the art, but neither can we conclude that it is inferior. Note that, because Pulse$^\infty$'s theory is sound and complete, and it directly represents the fundamental recurrence set idea of Gupta et al. [2008] (see 8), we *could* in principle import any algorithmic techniques that appear in other papers, for the purpose of establishing divergence triples. So the precision here is an implementation rather than a fundamental matter, and the purpose of the evaluation on the small benchmarks was just to confirm that the Pulse$^\infty$ is not too far off the state of the art, and having done so this sets us up for the more significant evaluation on larger projects.

7.3 Running Pulse$^\infty$ on Large Projects and Libraries

We evaluated Pulse$^\infty$ on a number of large open source projects including OpenSSL, libxml2, CryptoPP and libxpm. We present the result of our analysis in Table 3. As shown, each of the libraries analysed comprises thousands of lines of code (LOC) and were each analysed within minutes. In total, we identified and reported eleven previously unknown divergence bugs in OpenSSL, libxml2, CryptoPP and libxpm. Several of these bugs have been acknowledged with fixes waiting to be merged, while others are under discussion in the bug tracking system. Some of the issues we reported in OpenSSL are instances of *latent* (non-manifest) errors according to the

Table 3. Evaluating Pulse$^{\infty}$ on large projects (1.3 million lines of code analysed and 11 new bugs found)

Library	Language	#LOC analysed	Time	# Bugs reported
OpenSSL	C	804 K	1m, 26s	4
libxml2	C	300 K	1m, 3s	4
CryptoPP	C++	51 K	2m, 40s	2
libxpm	C	11 K	2s	1
libpng	C	96 K	6s	0
zlib	C	41 K	1m, 7s	0
ngiflib	C	1.7 K	1s	0
Total		**1.3 M**	**14m, 5s**	**11**

classification by Le et al. [2022]: they are only reachable when the culprit function is called with specific parameter values. Until these latent conditions are removed, it remains up to the caller to enforce well-behaved input to the problematic callee functions.

In the cases of CryptoPP and libxpm we have submitted pull requests with patches to be merged in due course. We also reported our findings in libxml2 on its project bug tracking system, where our discussion with project maintainers suggests that while the vulnerable functions are active, the specific divergent conditional branches in them are never executed and can be safely removed.

We present a bug we found in CryptoPP (a popular cryptographic toolkit in C++) in Listing 2. The procedure shown attempts to allocate memory in a loop (lines 4–9). However, as there is no guarantee that `malloc` will succeed and return a non-null value, the loop may not terminate. Our proposed fix is shown at lines 2, 5 and 7 (prefixed with '+'), where we record the number of unsuccessful allocations in `cnt` and throw an exception after 10 attempts.

```
1  void* AlignedAllocate(size_t size) {
2  +  unsigned int cnt = 0;
3     byte *p;
4     while ((p = (byte *)malloc(size+16)) == NULLPTR) {
5  +     if (cnt >= 10) { throw std::bad_alloc(); }
6        CallNewHandler();
7  +     cnt++;
8     }
9     CRYPTOPP_ASSERT(IsAlignedOn(p, 16));
10    return p;  }
```

Listing 2. A divergence bug found by Pulse$^{\infty}$ in CryptoPP, with our proposed fix given by adding the '+'-prefixed lines.

8 Related Work

Termination and Non-Termination Tools. There are many individual reports of divergence bugs which many readers will no doubt relate to. Notably, a recent empirical study on OSS projects found 445 non-termination bugs from 3,142 GitHub commits [Shi et al. 2022].

There has been significant work on automated methods for proving termination; see the survey by Cook et al. [2011]. When a termination prover fails, the question of whether the failed proof identifies a termination bug or if it is a false positive is more difficult than proving safety: termination bugs cannot be generally witnessed with finite traces (assuming unbounded resources in the computation model, that is). However, as Godefroid [2005] argues, the main value of analysis tools lies in the discovery of bugs, not in the proof of program correctness. Thus, it is valuable to consider proving non-termination, even without waiting for the wide deployment of termination verifiers.

The fundamental work of Gupta et al. [2008] uses proof to find divergence bugs. They use a transition system with initial/final states and a transition relation, and they identify the notion of a *recurrence set* R as (i) a non-empty intersection with the initial set of states; and (ii) reachability of R from every state satisfying R. Reachability in (ii) corresponds to $\vdash_{\mathrm{B}} [R]\ \mathrm{C}\ [ok\colon R]$. One might argue

that the relation between the UNTER proof system for $\vdash_{\mathsf{B}} [p]\ \mathsf{C}\ [ok: q]$ and the model of Gupta et al. [2008] is analogous to the relation between Hoare's logic and Floyd's proof method [Apt and Olderog 2019]: using under-approximate triples enables compositional reasoning. There are many detailed differences beyond these. They first run a concolic executor to gather assertions at program points, especially loop entry, but then employ an arithmetic encoding to derive reachability facts for loop bodies, and they treat the heap concretely (as the encoding is difficult otherwise). By contrast, we reason about reachability both of the loop stems and bodies within UNTER$^{\mathsf{SL}}$, and we harness separation logic (SL) to reason abstractly about heaps (SL-based analyses were not available at the time of the work by Gupta et al. [2008]).

Our prototype, Pulse$^{\infty}$, inherits the strengths and weaknesses of Pulse. Specifically, it is easy to run Pulse on program snippets, to scale it to large programs, and to incorporate it in a CI-based deployment on pull requests. On the other hand, Pulse has a weak treatment of arithmetic, meaning that some complex examples (as in the work of Gupta et al. [2008]) may not be provable. The strengths and weaknesses of Gupta et al. [2008] are the converse. We do not believe the weaknesses of either are inevitable; e.g. by adding a stronger arithmetic solver to Pulse$^{\infty}$ it would be possible to prove the more complex examples; the question is the effect this would have on performance. Upon contacting Gupta et al. [2008], we were informed that *their tool is no longer available*; as such, we were unable to compare Pulse$^{\infty}$ against it.

After Gupta et al. [2008], there have been many further papers on automatically proving/checking termination/non-termination. Cook et al. [2015] and Chen et al. [2014] introduce novel ideas on the use of over-approximation, going beyond the under-approximate logics here.

Many existing tools [Brockschmidt et al. 2013; Chatterjee et al. 2021; Kroening et al. 2010; Larraz et al. 2014; Velroyen and Rümmer 2008] focus on the syntax of termination problems defined by the Termination Competition and can *only* handle *integer C programs*, i.e. programs 1) with only integer datatypes; and 2) *without function calls*. As a result, such tools *cannot* run on programs that do not conform to this syntax, unless they are first pre-processed into integer C programs. In other words, unlike Pulse$^{\infty}$, these tools cannot run on existing C codebases or libraries such as OpenSSL.

The T2 tool by Cook et al. [2013] can be used to prove termination (and not divergence). However, T2 does not support heaps, and thus (unlike Pulse$^{\infty}$) cannot handle examples where termination is due to e.g. pointer arithmetic. Moreover, as our direct communications with the authors have revealed, T2 requires a C front-end, and the front-end the authors used bit-rotted' a while ago. The AProVE tool by Hensel et al. [2022] uses T2 as one of its back-ends for analysing termination; however, AProVE only supports integer C programs (as discussed above). Furthermore, upon contacting the authors about running AProVE on large libraries, we were informed that their techniques "are quite precise and do not have sufficient abstraction methods for handling very large programs within reasonable time". Most significantly, however, AProVE only supports programs that contain a `main()` method, and thus *cannot be used to analyse libraries* such as OpenSSL. This is in contrast to Pulse$^{\infty}$, where we successfully analysed hundreds of thousands of lines of code within minutes, and could effortlessly analyse libraries such as OpenSSL and libxml2. Heizmann et al. [2014] have extended the ULTIMATE framework to detect divergence using Büchi automata. However, as confirmed by the authors, a key technical limitation of their tool is that the input must be *a single C file* (and thus the program being analysed and all its dependencies must be included in one file), thus precluding its application to large projects and libraries comprising numerous modules and dependencies spread across several files.

Brotherston and Gorogiannis [2014] present CABER for proving termination (not divergence). While CABER supports heaps, it has only been applied to a handful of small programs, and not large code bases or libraries. Le et al. [2014, 2015] present the HipTNT and HipTNT+ tools for proving termination and non-termination. However, as we were informed in the course of our

direct communication with the authors, these tools can only handle small programs such as those in the SV-COMP benchmarks. Moreover, these tools require the user to supply certain annotations and are thus not fully automated. Le et al. [2020] later adapted these tools to develop DynamiTe, a dynamic termination analyser for non-linear program. However, DynamiTe can only handle integer programs (as discussed above.)

The idea of detecting divergence using proof is appealing and intuitively not too complicated. Although our work is but a step on the way, it is reasonable to hope that divergence proof techniques may mature to a degree where they can be routinely deployed in engineering practice.

Under-Approximation and Incorrectness. This paper follows a line of work on under-approximate reasoning following incorrectness logic [O'Hearn 2019], but is part of a more extensive history. O'Hearn [2019] used FUA triples to reason about incorrectness and to avoid false positives. FUA triples were studied previously by de Vries and Koutavas [2011], but they did not make the connection to incorrectness or absence of false positives. Further, as O'Hearn [2019] remarked, FUA triples could be expressed in dynamic logic [Harel 1979] with a backwards diamond modality (or forwards with a transition reversal operator). Moreover, they are similar to the must$^-$ transitions used by e.g. [Ball et al. 2005]. As such, the FUA triple is not itself novel, but its significance has been uncovered gradually.

Here we study both BUA and FUA triples. BUA triples were mentioned by de Vries and Koutavas [2011] under the name "total Hoare triple", but were not studied by them. These triples can also be expressed immediately in dynamic logic, without resorting to backwards modalities or reversal, and they are similar to the must$^+$ transitions used of [Ball et al. 2005]. More recently BUA triples were suggested by Derek Dreyer and Ralf Jung just before IL was formulated, but they remained unexplored. (This was during a discussion with Peter O'Hearn and Jules Villard at POPL 2019 in Lisbon, and thus BUA triples are also informally referred to as 'Lisbon' triples in the literature.) BUA triples are also studied by Möller et al. [2021], but only their metatheory and not their applications. Zilberstein et al. [2023] developed Outcome Logic (OL) where they make meta-theoretic remarks on how BUA could serve as a foundation for incorrectness, but they do not demonstrate the practical advantages or disadvantages of such reasoning. Specifically, while they discuss the merits of BUA for identifying manifest errors, they do not demonstrate the practical impact of this e.g. in a scalable analysis tool. Nor do they explore the advantage of BUA for non-termination analysis. Zilberstein et al. [2024] later extend OL with separation logic. Ascari et al. [2024] also study BUA triples in sufficient incorrectness logic (SIL). As with FUA, the notion of BUA is not itself novel, but its significance is emerging gradually. This paper adds two new insights about BUA.

The first is that abducing preconditions is, in a sense, forced on us if we are to do forward reasoning with BUA. This is because BUA triples are not closed under postconditions: given a precondition p and a program C, there need not exist a corresponding postcondition making the triple valid. This is the case for any p which has a state on which C always diverges. As a result, there is no analogue for BUA of Dijkstra's strongest postcondition predicate transformer (where this transformer works for FUA). This would, at first glance, make BUA appear problematic for forward reasoning: forward reasoning in abstract interpretation uses over-approximations of Dijkstra's transformer, reasoning with FUA can use under-approximation of it, and this abstraction-of-post tactic is not available for BUA. It might be possible to automate backward reasoning instead (find preconditions given a program and a post), but there is another possibility: abduction. Given a precondition and a program, we can try to abduce an addition $?A$ to the precondition to guarantee the existence of a postcondition. For example, for $\vdash_{\mathsf{B}} \left[true \wedge ?A\right]$ `if(even(x)) diverge else x:=x+1` $\left[ok\colon ?B\right]$, we can abduce $?A$ = `odd(x)` to give us a precondition to establish the postcondition $?B$ = `even(x)`.

Abduction is used to reason forwards with BUA triples to skirt the absence of general postconditions. This is not unlike the case in classic separation logic (SL), where the absence of general postconditions for the fault-avoiding interpretation of SL triples is circumvented by abducing safe preconditions [Calcagno et al. 2011]. We emphasise that the point we are making here is stronger than the mere *compatibility* of BUA with abduction (which has been recognised by Zilberstein et al. [2024]): we cannot reason forward from an arbitrary preconditon, without changing it. For comparison, consider the situation with FUA. In FUA we can do abduction, it is compatible with it, but since FUA is closed under postconditions (there is a post for any given precondition) we could in principle reason forwards if we so desired, without using abduction to shrink the precondition.

The second BUA insight in our work is that (unlike in other works [Ascari et al. 2024; de Vries and Koutavas 2011; Möller et al. 2021; Zilberstein et al. 2023, 2024]) we demonstrate its suitability for reasoning about non-termination, where pure FUA is unsound for non-termination proof rules. This goes together with the ability of BUA triples to *weaken a postcondition*, which opens up the possibility of iterating to a fixed-point as suggested in the main proof rule for diverging loops. Thus, BUA tools can share some of the iteration strategies with their over-approximate cousins, and in contrast to FUA-only tools. A crucial difference is that the fixpoint then implies divergence.

It is important to note that there are disadvantages to BUA-only approaches. For example, BUA does not support shrinking a postcondition, which would block the application of partial concretisation as in Klee, DART and similar tools (see the work of [O'Hearn 2019]). Indeed, must$^-$ transitions, a relative of FUA triples, have been used to formalise the reasoning in such tools [Godefroid et al. 2010]. Note that there is a more basic under-approximate triple, let us call it the existential (EUA) triple written as $\vdash_{\mathsf{E}} [P]\ \mathsf{C}\ [\epsilon : Q]$, stating that *some* state in P reaches *some* state in Q by executing C. EUA is sufficient for proving incorrectness and avoiding false positives, so why is it not the basis of a program logic? The problem is that EUA is not closed under sequential composition (the Seq rule fails), which makes reasoning about paths challenging. Both BUA and FUA triples are closed under sequential composition, and this is (we presume) why they have received more attention. But, the composition of a BUA followed by a FUA triple establishes an EUA triple; BUA and FUA can be used together [Ball et al. 2005].

Due to the reasons discussed above, it is better for a tool, or a portion of a tool, to be compatible with both BUA and FUA, rather than one or the other. As observed here, the Pulse framework is compatible with both. Extension to non-termination or concretisation should take into account considerations as above. While the BUA and FUA metatheory seems mostly settled, we do not claim that the above remarks fully account for their strengths, weaknesses or overall significance for reasoning: we are still learning about them.

Data Availability Statement

The proofs of all stated theorems are given in the extended version [Raad et al. 2024a]. Our prototype tool Pulse$^\infty$ is open-source and available as an artifact online [Raad et al. 2024b].

Acknowledgments

We thank the anonymous OOPSLA 2024 reviewers for their constructive feedback, and thank Viktor Vafeiadis and Jules Villard for many fruitful discussions. Azalea Raad is supported by the UKRI Future Leaders Fellowship MR/V024299/1, by the EPSRC grant EP/X037029/1 and by VeTSS.

References

Krzysztof R. Apt and Ernst-Rüdiger Olderog. 2019. Fifty years of Hoare's logic. *Formal Aspects Comput.* 31, 6 (2019), 751–807. https://doi.org/10.1007/s00165-019-00501-3

Flavio Ascari, Roberto Bruni, Roberta Gori, and Francesco Logozzo. 2024. Sufficient Incorrectness Logic: SIL and Separation SIL. arXiv:2310.18156 [cs.LO] https://arxiv.org/abs/2310.18156

Thomas Ball, Orna Kupferman, and Greta Yorsh. 2005. Abstraction for Falsification. In *Computer Aided Verification, 17th International Conference, CAV 2005, Edinburgh, Scotland, UK, July 6-10, 2005, Proceedings (Lecture Notes in Computer Science)*, Kousha Etessami and Sriram K. Rajamani (Eds.), Vol. 3576. Springer, 67–81. https://doi.org/10.1007/11513988_8

Josh Berdine, Aziem Chawdhary, Byron Cook, Dino Distefano, and Peter O'Hearn. 2007. Variance Analyses from Invariance Analyses. *SIGPLAN Not.* 42, 1 (jan 2007), 211–224. https://doi.org/10.1145/1190215.1190249

Josh Berdine, Byron Cook, Dino Distefano, and Peter W. O'Hearn. 2006. Automatic Termination Proofs for Programs with Shape-Shifting Heaps. In *Computer Aided Verification*, Thomas Ball and Robert B. Jones (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 386–400.

Sam Blackshear, Nikos Gorogiannis, Peter W. O'Hearn, and Ilya Sergey. 2018. RacerD: Compositional Static Race Detection. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 144 (Oct. 2018), 28 pages. https://doi.org/10.1145/3276514

Marc Brockschmidt, Byron Cook, and Carsten Fuhs. 2013. Better termination proving through cooperation. In *Computer Aided Verification: 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings 25*. Springer, 413–429.

James Brotherston and Nikos Gorogiannis. 2014. Cyclic Abduction of Inductively Defined Safety and Termination Preconditions. In *Static Analysis*, Markus Müller-Olm and Helmut Seidl (Eds.). Springer International Publishing, Cham, 68–84.

Cristiano Calcagno, Dino Distefano, Peter W. O'Hearn, and Hongseok Yang. 2011. Compositional Shape Analysis by Means of Bi-Abduction. *J. ACM* 58, 6, Article 26 (Dec. 2011), 66 pages. http://doi.acm.org/10.1145/2049697.2049700

Krishnendu Chatterjee, Ehsan Kafshdar Goharshady, Petr Novotnỳ, and Đorđe Žikelić. 2021. Proving non-termination by program reversal. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. 1033–1048.

Aziem Chawdhary, Byron Cook, Sumit Gulwani, Mooly Sagiv, and Hongseok Yang. 2008. Ranking Abstractions. In *Programming Languages and Systems*, Sophia Drossopoulou (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 148–162.

Hong Yi Chen, Byron Cook, Carsten Fuhs, Kaustubh Nimkar, and Peter W. O'Hearn. 2014. Proving Nontermination via Safety. In *Tools and Algorithms for the Construction and Analysis of Systems - 20th International Conference, TACAS 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014. Proceedings (Lecture Notes in Computer Science)*, Erika Ábrahám and Klaus Havelund (Eds.), Vol. 8413. Springer, 156–171. https://doi.org/10.1007/978-3-642-54862-8_11

Byron Cook, Carsten Fuhs, Kaustubh Nimkar, and Peter W. O'Hearn. 2015. Embracing Overapproximation for Proving Nontermination. *Tiny Trans. Comput. Sci.* 3 (2015). http://tinytocs.org/vol3/papers/TinyToCS_3_cook.pdf

Byron Cook, Andreas Podelski, and Andrey Rybalchenko. 2006a. Termination Proofs for Systems Code. *SIGPLAN Not.* 41, 6 (jun 2006), 415–426. https://doi.org/10.1145/1133255.1134029

Byron Cook, Andreas Podelski, and Andrey Rybalchenko. 2006b. Terminator: Beyond Safety. In *Computer Aided Verification*, Thomas Ball and Robert B. Jones (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 415–418.

Byron Cook, Andreas Podelski, and Andrey Rybalchenko. 2011. Proving program termination. *Commun. ACM* 54, 5 (2011), 88–98. https://doi.org/10.1145/1941487.1941509

Byron Cook, Abigail See, and Florian Zuleger. 2013. Ramsey vs. lexicographic termination proving. In *Tools and Algorithms for the Construction and Analysis of Systems: 19th International Conference, TACAS 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings 19*. Springer, 47–61.

Pedro da Rocha Pinto, Thomas Dinsdale-Young, Philippa Gardner, and Julian Sutherland. 2016. Modular Termination Verification for Non-blocking Concurrency. In *Programming Languages and Systems*, Peter Thiemann (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 176–201.

Edsko de Vries and Vasileios Koutavas. 2011. Reverse Hoare Logic. In *Software Engineering and Formal Methods - 9th International Conference, SEFM 2011, Montevideo, Uruguay, November 14-18, 2011. Proceedings*. 155–171. https://doi.org/10.1007/978-3-642-24690-6_12

Dino Distefano, Manuel Fähndrich, Francesco Logozzo, and Peter W. O'Hearn. 2019. Scaling static analyses at Facebook. *Commun. ACM* 62, 8 (2019), 62–70. https://doi.org/10.1145/3338112

Emanuele D'Osualdo, Julian Sutherland, Azadeh Farzan, and Philippa Gardner. 2021. TaDA Live: Compositional Reasoning for Termination of Fine-Grained Concurrent Programs. *ACM Trans. Program. Lang. Syst.* 43, 4, Article 16 (nov 2021), 134 pages. https://doi.org/10.1145/3477082

Patrice Godefroid. 2005. The soundness of bugs is what matters (position statement). https://www.cs.umd.edu/~pugh/BugWorkshop05/papers/11-godefroid.pdf

Patrice Godefroid, Aditya V. Nori, Sriram K. Rajamani, and SaiDeep Tetali. 2010. Compositional may-must program analysis: unleashing the power of alternation. In *Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, January 17-23, 2010*, Manuel V. Hermenegildo and Jens Palsberg (Eds.). ACM, 43–56. https://doi.org/10.1145/1706299.1706307

Ashutosh Gupta, Thomas A. Henzinger, Rupak Majumdar, Andrey Rybalchenko, and Ru-Gang Xu. 2008. Proving non-termination. In *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008*, George C. Necula and Philip Wadler (Eds.). ACM, 147–158. https://doi.org/10.1145/1328438.1328459

David Harel. 1979. *First-Order Dynamic Logic.* Springer-Verlag, Berlin, Heidelberg.

William R Harris, Akash Lal, Aditya V Nori, and Sriram K Rajamani. 2010. Alternation for termination. In *Static Analysis: 17th International Symposium, SAS 2010, Perpignan, France, September 14-16, 2010. Proceedings 17.* Springer, 304–319.

Matthias Heizmann, Jochen Hoenicke, and Andreas Podelski. 2014. Termination Analysis by Learning Terminating Programs. In *Computer Aided Verification*, Armin Biere and Roderick Bloem (Eds.). Springer International Publishing, Cham, 797–813.

Jera Hensel, Constantin Mensendiek, and Jürgen Giesl. 2022. AProVE: Non-Termination Witnesses for C Programs: (Competition Contribution). In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 403–407.

Samin S. Ishtiaq and Peter W. O'Hearn. 2001. BI as an Assertion Language for Mutable Data Structures. In *Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (London, United Kingdom) *(POPL)*. Association for Computing Machinery, New York, NY, USA, 14–26. https://doi.org/10.1145/360204.375719

Daniel Kroening, Natasha Sharygina, Aliaksei Tsitovich, and Christoph M Wintersteiger. 2010. Termination analysis with compositional transition invariants. In *International Conference on Computer Aided Verification*. Springer, 89–103.

Daniel Larraz, Kaustubh Nimkar, Albert Oliveras, Enric Rodríguez-Carbonell, and Albert Rubio. 2014. Proving non-termination using Max-SMT. In *Computer Aided Verification: 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18-22, 2014. Proceedings 26*. Springer, 779–796.

Quang Loc Le, Azalea Raad, Jules Villard, Josh Berdine, Derek Dreyer, and Peter W. O'Hearn. 2022. Finding Real Bugs in Big Programs with Incorrectness Logic. *Proc. ACM Program. Lang.* 6, OOPSLA1, Article 81 (apr 2022), 27 pages. https://doi.org/10.1145/3527325

Ton Chanh Le, Timos Antonopoulos, Parisa Fathololumi, Eric Koskinen, and ThanhVu Nguyen. 2020. DynamiTe: dynamic termination and non-termination proofs. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 189 (nov 2020), 30 pages. https://doi.org/10.1145/3428257

Ton Chanh Le, Cristian Gherghina, Aquinas Hobor, and Wei-Ngan Chin. 2014. A resource-based logic for termination and non-termination proofs. In *Formal Methods and Software Engineering: 16th International Conference on Formal Engineering Methods, ICFEM 2014, Luxembourg, Luxembourg, November 3-5, 2014. Proceedings 16*. Springer, 267–283.

Ton Chanh Le, Shengchao Qin, and Wei-Ngan Chin. 2015. Termination and non-termination specification inference. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation, Portland, OR, USA, June 15-17, 2015*, David Grove and Stephen M. Blackburn (Eds.). ACM, 489–498. https://doi.org/10.1145/2737924.2737993

Hongjin Liang and Xinyu Feng. 2016. A Program Logic for Concurrent Objects under Fair Scheduling. *SIGPLAN Not.* 51, 1 (jan 2016), 385–399. https://doi.org/10.1145/2914770.2837635

Bernhard Möller, Peter W. O'Hearn, and Tony Hoare. 2021. On Algebra of Program Correctness and Incorrectness. In *Relational and Algebraic Methods in Computer Science - 19th International Conference, RAMiCS 2021, Marseille, France, November 2-5, 2021, Proceedings (Lecture Notes in Computer Science)*, Uli Fahrenberg, Mai Gehrke, Luigi Santocanale, and Michael Winter (Eds.), Vol. 13027. Springer, 325–343. https://doi.org/10.1007/978-3-030-88701-8_20

Peter W. O'Hearn. 2019. Incorrectness Logic. *Proc. ACM Program. Lang.* 4, POPL, Article 10 (Dec. 2019), 32 pages. http://doi.acm.org/10.1145/3371078

Azalea Raad, Josh Berdine, Hoang-Hai Dang, Derek Dreyer, Peter O'Hearn, and Jules Villard. 2020. Local Reasoning About the Presence of Bugs: Incorrectness Separation Logic. In *Computer Aided Verification*, Shuvendu K. Lahiri and Chao Wang (Eds.). Springer International Publishing, Cham, 225–252.

Azalea Raad, Josh Berdine, Derek Dreyer, and Peter W. O'Hearn. 2022. Concurrent Incorrectness Separation Logic. *Proc. ACM Program. Lang.* 6, POPL, Article 34 (jan 2022), 29 pages. https://doi.org/10.1145/3498695

Azalea Raad, Julien Vanegue, Josh Berdine, and Peter O'Hearn. 2023. A General Approach to Under-Approximate Reasoning About Concurrent Programs. In *34th International Conference on Concurrency Theory (CONCUR 2023) (Leibniz International Proceedings in Informatics (LIPIcs))*, Guillermo A. Pérez and Jean-François Raskin (Eds.), Vol. 279. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 25:1–25:17. https://doi.org/10.4230/LIPIcs.CONCUR.2023.25

Azalea Raad, Julien Vanegue, and Peter O'Hearn. 2024a. Extended Version. https://www.soundandcomplete.org/papers/OOPSLA2024/Unter/

Azalea Raad, Julien Vanegue, and Peter O'Hearn. 2024b. The Pulse$^\infty$ prototype tool. https://doi.org/10.5281/zenodo.12637589

Caitlin Sadowski, Edward Aftandilian, Alex Eagle, Liam Miller-Cushon, and Ciera Jaspan. 2018. Lessons from Building Static Analysis Tools at Google. *Commun. ACM* 61, 4 (March 2018), 58–66. https://doi.org/10.1145/3188720

Xiuhan Shi, Xiaofei Xie, Yi Li, Yao Zhang, Sen Chen, and Xiaohong Li. 2022. Large-scale analysis of non-termination bugs in real-world OSS projects. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2022, Singapore, Singapore, November 14-18, 2022*, Abhik Roychoudhury, Cristian Cadar, and Miryung Kim (Eds.). ACM, 256–268. https://doi.org/10.1145/3540250.3549129

Helga Velroyen and Philipp Rümmer. 2008. Non-termination checking for imperative programs. In *International Conference on Tests and Proofs*. Springer, 154–170.

Noam Zilberstein, Derek Dreyer, and Alexandra Silva. 2023. Outcome Logic: A Unifying Foundation for Correctness and Incorrectness Reasoning. *Proc. ACM Program. Lang.* 7, OOPSLA1, Article 93 (apr 2023), 29 pages. https://doi.org/10.1145/3586045

Noam Zilberstein, Angelina Saliling, and Alexandra Silva. 2024. Outcome Separation Logic: Local Reasoning for Correctness and Incorrectness with Computational Effects. *Proc. ACM Program. Lang.* 8, OOPSLA1, Article 104 (apr 2024), 29 pages. https://doi.org/10.1145/3649821