

A Resource-Based Logic for Termination and Non-termination Proofs

Ton Chanh Le¹, Cristian Gherghina², Aquinas Hobor¹, and Wei-Ngan Chin¹

¹ Department of Computer Science, National University of Singapore, Singapore

² Singapore University of Technology and Design, Singapore

Abstract. We propose a unified logical framework for specifying and proving *both* termination and non-termination of various programs. Our framework is based on a resource logic which captures both upper *and* lower bounds on resources used by the programs. By an abstraction, we evolve this resource logic for execution length into a temporal logic with three predicates to reason about termination, non-termination or unknown. We introduce a new logical entailment system for temporal constraints and show how Hoare logic can be seamlessly used to prove termination and non-termination in our unified framework. Though this paper's focus is on the formal foundations for a new unified framework, we also report on the usability and practicality of our approach by specifying and verifying both termination and non-termination properties for about 300 programs, collected from a variety of sources. This adds a modest 5-10% verification overhead when compared to underlying partial-correctness verification system.

1 Introduction

Termination proving is an important part of correctness proofs for software systems as “so-called *partial correctness* is inadequate: if a program is intended to terminate, that fact must be part of its specification.” – Cliff Jones [26]. Thus, *total correctness* proofs, denoted by the Hoare triple $[P]c[Q]$, require the code fragment c to be shown terminating in addition to meeting the postcondition Q after execution. The termination of a loop or a recursive method is usually proven by a well-founded termination measure given to the specification. However, such a measure is not a component of the logical formulas for pre/post specifications. A reason for this distinction is that specification logic typically describes program states, while the termination proofs are concerned with the existence of well-founded measures to bound the execution length of loops/recursions, as argued by Hehner in [23]. Due to this distinction, we cannot automatically leverage richer logics that have been developed for safety properties to conduct more intricate termination and non-termination reasoning.

For illustration, let us use the Shuffle problem proposed in the Java Bytecode Recursive category of the annual Termination Competition [33]. In this problem, an acyclic linked list is shuffled by the `shuffle` method together with the auxiliary `reverse` method, whose source code is shown in Fig. 1. To prove that `shuffle` terminates, we need to firstly show that `reverse` also terminates. While the termination of `reverse` can be easily proved by current approaches, such as [30,7,11], proving `shuffle` terminates is harder because it requires a functional correctness related fact: the `reverse` method does not change the length of the list. Based on this fact, it is possible to show

```

public static List shuffle(List xs) { public static List reverse(final List l) {
  if (xs == null) return null;      if (l == null || l.next == null)
  else {                             return l;
    List next = xs.next;             final List nextItem = l.next;
    return new List(xs.value,        final List reverseRest =
      shuffle(reverse(next)));       reverse(nextItem);
  }                                  l.next = null; nextItem.next = l;
}                                    return reverseRest; }

```

Fig. 1. The Shuffle problem from the Termination Competition

that the linked list's length is also decreasing across the recursive method call `shuffle`; as a result, the method always terminates.

Therefore, without an integration of termination specification into logics for functional correctness, such as separation logic [36], the termination of `shuffle` is hardly specified and proved by verification systems based on the traditional Hoare logic for total correctness. Note that automated termination provers, such as AProVE [19] and COSTA [3], are not able to show that `shuffle` terminates, even after applying a numeric abstraction on the size property to `shuffle` [32], due to the lack of information flow between the correctness and the termination arguments. We believe that relatively complex problems, such as Shuffle, highlight the need of a more expressive logic with the ability of integration into various safety logics for termination reasoning.

Moreover, if the termination proof fails, *e.g.*, when the input list of `shuffle` is cyclic, the program will be *implicitly* assumed to be *possibly* non-terminating. That is, *definite* non-termination is neither *explicitly* stated nor proven by Hoare logic. Explicitly proving non-termination has two benefits. First, it allows more comprehensive specifications to be developed for better program understanding. Second, it allows a clearer distinction between expected non-termination (*e.g.*, reactive systems where loops are designed to be infinite) and failure of termination proofs, paving the way for focusing on real non-termination bugs that minimize on false positives.

Some specification languages, such as Dafny [31], ACSL [6] and JML [29], allow the specification of *possible* non-termination but their corresponding verifiers provide limited support for this feature. For example, the verifier of Dafny only allows such specification on loops or tail-recursive methods¹, while Frama-C verifier of ACSL [15] has not implemented it. On the other hand, we can use the `false` postcondition, which indicates that the method's exit is unreachable, to specify *definite* non-termination. However, such postcondition for partial correctness is not preferred as it is logically distinct from termination proofs. This distinction has been designed into Dafny, Frama-C and KeY with JML [2], that makes the tools fail to take into account non-terminating behavior when proving termination. For example, Dafny succeeds in proving the termination of a recursive method² though this method contains a call to a non-terminating method.³ In fact, for termination proofs, these tools simply check that there is a finite number of mutually recursive calls to the analyzed methods, rather than the methods' termination per se.

¹ <http://www.rise4fun.com/Dafny/PnRX>

² <http://www.rise4fun.com/Dafny/6FuR>

³ The examples in ACSL and JML are at

<http://loris-7.ddns.comp.nus.edu.sg/~project/hiptnt/others.zip>

Our Proposal. We propose integrating both termination and non-termination requirements directly into the specification logic for functional properties. Our work follows Hoare and He [24] and Hehner [22], in which the termination is reasoned together with partial correctness proof. In [22], the program is instrumented with a time variable t and the termination is proven by a finite bound on the *exact* execution time $t' - t$, where t, t' are the initial, resp. final time. In [24], a special ghost variable *ok* is used to signify termination. However, these approaches presently do not handle non-termination.

As a formal foundation to unify termination and non-termination reasoning and integrate them into functional correctness proofs, we introduce a new resource logic which captures the concept of resource capacity; tracking both *minimum and maximum* amounts of resources used by some given code. Our logic uses a primitive predicate $RC(l, u)$ with invariant $0 \leq l \leq u$ to capture a semantic notion of resource capacity (l, u) with the lower bound l and the upper bound u . Through this resource logic, we can specify a variety of complexity-related properties, including the notions of termination and non-termination, by tracking the number of calls (and loop iterations) executed by the given code. Termination is denoted by the presence of a *finite* upper bound, while non-termination is denoted an *infinite* lower bound on the execution length.

To support a more effective mechanism, we shall derive a simpler *temporal* logic from the richer resource logic itself. We define three temporal predicates, *Term* M , *Loop* and *MayLoop*, where M is a well-founded termination measure, and associate them with each method in a given program to denote the termination, definite non-termination and possible non-termination of these methods, respectively. In terms of resource reasoning, these predicates represent $RC(0, \text{embed}(M))$, $RC(\infty, \infty)$ and $RC(0, \infty)$, respectively, where $\text{embed}(M)$ is a finite bound obtained through an order-embedding of M into naturals. Using the enriched specification logic, functional correctness, termination and non-termination of methods can be verified under a single modular framework. With this unification, the predicate *Term* M denotes exactly definite termination, instead of just denoting the bound on the number of loop iterations or method recursions like the termination measures used in the traditional Hoare logic for total correctness.

Our research contributions can be summarized as follows:

- A new *resource logic* that can capture lower and upper bounds on resource usage via the concept of resource capacity, together with an entailment procedure to support correctness proofs with resource-related properties. (Sec. 3)
- A *temporal logic* that is abstracted from the resource logic to reason about both program termination and non-termination. We introduce three new temporal constraints, its *entailment* and *Hoare rules* lifted from the resource logic. (Sec. 4)
- A successful integration of both resource and temporal logics into a separation logic based verifier [35]. The new temporal logic is expressive enough to specify and successfully verify the (non-)termination behaviors for about 300 benchmark programs collected from a variety of sources, including the SIR/Siemens test suite [17] and problems from the Termination Competition (Sec. 5). The prototype implementation and benchmark are available for online use and download at:

<http://loris-7.ddns.comp.nus.edu.sg/~project/hiptnt/>

<pre> pred isEvenNat(int n) $\equiv n \geq 0 \wedge \exists m \cdot n = 2 * m$; int sumE (int n) requires isEvenNat(n) \wedge Term[n] \vee \negisEvenNat(n) \wedge Loop ensures true; { if (n==0) return 0; else return n + sumE(n-2); } (a) </pre>	<pre> while (x>y) requires x \leq y \wedge Term[] \vee x > y \wedge x < 0 \wedge Loop \vee x > y \wedge x \geq 0 \wedge MayLoop ensures x' \leq y'; { y=x+y; x=x-1; } (b) </pre>
---	--

Fig. 2. Examples on numerical programs

2 From Resource to Temporal Logic

We introduce a general resource predicate $RC\langle l, u \rangle$ where l is a lower bound and u is an upper bound on resource capacity, with invariant $0 \leq l \leq u$. This resource predicate can be specialized to execution capacity to capture a variety of complexity-related properties, via lower and upper bounds on the total number of method calls during the execution of a given piece of code. We shall give an instrumented semantics for this specific resource logic, and also specialize it for reasoning about termination and non-termination. To prove termination, we simply use the predicate $RC\langle 0, u \rangle$ where u is some finite value, namely $u < \infty$. To prove non-termination, we can use the predicate $RC\langle \infty, \infty \rangle$ which signifies an infinite lower bound. Lastly, if we cannot prove either termination or non-termination, we use the predicate $RC\langle 0, \infty \rangle$ which covers all possibilities.

The resource logic we have outlined is quite expressive, and could moreover be specialized for reasoning on just termination and non-termination with the direct handling of infinity ∞ value. In order to design a simpler logic, we introduce a temporal logic with three distinct predicates, as follows: (i) $\text{Term } M$ to denote $RC\langle 0, \text{embed}(M) \rangle$, (ii) Loop to denote $RC\langle \infty, \infty \rangle$ and (iii) MayLoop to denote $RC\langle 0, \infty \rangle$. Such a temporal logic is considerably simpler than the more expressive resource logic, since we can omit reasoning with ∞ . We can also use a simpler termination measure M , based on depth of recursion rather than number of calls, but relate to the latter using $\text{embed}(M)$. Moreover, these temporal predicates can be made flow-insensitive, and thus need only appear in each method's precondition where they describe execution capacity required for the method's execution. This two-level approach simplifies both the design of a formal semantics, and the development of a verification framework for (non-)termination.

For illustration, let us look at some numerical examples, starting with the method `sumE` in Fig. 2(a). This method is required to return the sum of all even natural numbers that are less than or equal to the input `n`. However, the implementation satisfies this requirement only when `n` is an even natural number, denoted by the predicate `isEvenNat(n)`; otherwise, the method does not terminate⁴. In our approach, these distinct scenarios can be described in a termination-enriched specification by seamlessly integrating the temporal constraints `Term[n]` and `Loop` into a logic with disjunctions.

⁴ The verification system assumes the use of arbitrary precision integers. When finite integers are used, we may give a different temporal specification for those prestates.

JML and ACSL also support the specification of several method behaviors. However, the current ACSL implementation in Frama-C does not allow fine-grained termination related specification of each behavior and ignores conditional termination clauses. As a result, it cannot verify all the (non-)terminating behaviors of `sumE` together. KeY allows the specification of termination for each individual method behavior but it cannot disprove the termination of `sumE` when `n` is an odd positive number, because the variant `n` is still valid under this precondition.³ In contrast, our unified termination and non-termination reasoning does not accept the temporal constraint $\text{Term}[n]$ in these prestates because the execution starting from them will eventually reach a non-terminating execution when $n < 0$. In terms of resource reasoning, $\text{Term}[n]$, denoting a finite resource, is invalid as it cannot satisfy the infinite resource required by the non-termination.

The next example in Fig. 2(b) illustrates a usage of `MayLoop` constraint. Starting from any prestate satisfying $x > y \wedge x \geq 0$, the execution of the given loop may reach either the base case (when $x \leq y$, indicated by $\text{Term}[]$) or the non-terminating case (when $x > y \wedge x < 0$, indicated by `Loop`). We observe that this `MayLoop` precondition can be strengthened to the *non-linear* constraint $4x^2 + 4x + 8y + 9 \geq 0$ for non-termination, but this requires stronger arithmetic solvers.

Though our proposal is independent of the underlying logics on functional properties, it can leverage infrastructures of richer logics⁵ to conduct termination and non-termination reasoning for more complex domains. For example, our proposed temporal constraints are easily integrated into formulas of separation logic to reason about the termination and non-termination of heap-based programs. We choose a fragment of separation logic with the separating conjunction $*$ and the points-to operator \mapsto to specify the heap assertions. These operators are used to describe several data structures, such as linked list and tree. For example, the inductive predicate $\text{lseg}(\text{root}, \text{p}, \text{n})$ declared in Fig. 3(a) describes a list segment size of `n` from `root` to `p` with an invariant property stated that the list’s size is non-negative. This predicate can be used to specify either null-terminating lists (when `p = null`) or circular lists (when `p = root`).

<pre>data List { int value; List next; } pred lseg(root, p, n) ≡ root=p ∧ n=0 ∨ ∃v, q · root↦List(v, q)*lseg(q, p, n-1) inv n ≥ 0;</pre> <p style="text-align: center;">(a)</p>	<pre>List reverse (List l) requires lseg(l, null, n) ∧ Term[n] ensures lseg(res, null, n); List shuffle (List xs) requires lseg(xs, null, n) ∧ Term[n] ensures lseg(res, null, n);</pre> <p style="text-align: center;">(b)</p>
--	--

Fig. 3. A specification in separation logic to verify the correctness of `Shuffle`’s methods

We then use the predicate lseg for the pre and postconditions of two methods `reverse` and `shuffle` in the `Shuffle` problem. The specification of each method indicates that the method’s result `res` is a linked list with the same size `n` as the input list. From these safety specifications, the temporal constraint $\text{Term}[n]$ integrated into the precondition of each method is able to specify that the depth of recursion is bounded by the size of the input list, thus indicating the method’s termination.

⁵ In comparison with the first-order logic with linear arithmetic for numerical programs.

From the perspective of resource reasoning, a temporal constraint in the precondition of a method defines the bounds of available resource allowed for program executions from prestates satisfying (safety part of) this precondition. This idea is similar to Atkey’s logic [5], a type-based amortized resource analysis for imperative programs, which associates a piece of resource with each element of the data structures prior program execution. However, Atkey’s approach only tracks the upper bound of resource usage, so that it cannot reason about non-termination. This shortcoming also applies to other type-based approaches for termination reasoning, such as [1,38]. In addition, while the amortized resource analysis accounts for individual time-step (or heap chunk), we use termination measures, which are much simpler, to facilitate termination proofs. For example, to analyze `shuffle`, Atkey’s logic requires the global length property to present the polynomial resource associated with the input list using the technique of Hoffmann and Hofmann [25], which is much harder than locally reasoning about each node of the list as stated in his paper. Finally, this logic is built on top of just separation logic, rather than being generic as our proposal.

3 A Logic for Resource Reasoning

In proving termination and non-termination, our goal is to use resource reasoning based on execution capacity to provide a means for quantitatively assessing the execution length of a program. For this purpose, we introduce a resource logic to formally assess the minimum and a maximum bounds on a program’s resource consumption. We first extend the program state model with a mechanism to track resource capacities of the underlying machine. Since the particular consumed resource is countable and possibly infinite, we use the set \mathbb{N}^∞ , short for $\mathbb{N} \cup \{\infty\}$, as its domain.

3.1 Resource Capacity

Definition 1 (Program states). *A program state σ is a triple (s, h, r) of stack $s \in \mathcal{S}$ (locals), heap $h \in \mathcal{H}$ (memory) and $r \in \mathcal{R}$, resource capacity where r is a pair (r_l, r_u) of bounds in \mathbb{N}^∞ , with $0 \leq r_l \leq r_u$, denoting the allowed minimum and maximum resource consumption for executions starting from the current program state.*

Intuitively, a program state’s resource capacity (r_l, r_u) ensures that any execution starting from this state must consume *at least* r_l and *at most* r_u of the tracked resource.

Definition 2 (Resource Capacity Ordering). *Let $(\leq_c) \subset \mathbb{N}^\infty \times \mathbb{N}^\infty$ be the resource capacity ordering, such that $(b_l, b_u) \leq_c (a_l, a_u)$ iff $a_l \leq b_l$ and $b_u \leq a_u$.*

The resource capacity (a_l, a_u) is considered larger (or more general) than (b_l, b_u) if $a_l \leq b_l$ and $b_u \leq a_u$. The intuition is that under this condition, any execution which guarantees the capacity (b_l, b_u) also guarantees the capacity (a_l, a_u) . Based on this observation, $(0, \infty)$ is the largest resource capacity. In fact, it indicates an unconstrained resource consumption.

In order to properly define an operational semantics in terms of the proposed program state model, we also need to be able to express resource consumption. To this end we define a splitting operation over the resource capacity. We will say that a capacity (a_l, a_u) can be split into capacities (b_l, b_u) and (c_l, c_u) , written $(a_l, a_u) \ominus (b_l, b_u) = (c_l, c_u)$, if whenever an execution that guarantees the capacity (b_l, b_u) starts from a state

$(s, h, r) \models \Psi_1 \vee \Psi_2$	\equiv	$(s, h, r) \models \Psi_1$ or $(s, h, r) \models \Psi_2$
$(s, h, r) \models \Psi_1 \wedge \Psi_2$	\equiv	$(s, h, r) \models \Psi_1$ and $(s, h, r) \models \Psi_2$
$(s, h, r) \models \exists x_i^* \cdot \Psi$	\equiv	$\exists \nu_i^* \cdot (s[(x_i \mapsto \nu_i)^*], h, r) \models \Psi$
$(s, h, r) \models \mu$	\equiv	$(s, h) \models \mu$
$(s, h, r) \models \text{RC}\langle a_l, a_u \rangle$	\equiv	$(s, h) \models r_l = a_l \wedge r_u = a_u$ where $r = (r_l, r_u)$
$(s, h, r) \models \rho_1 \blacktriangleright \rho_2$	\equiv	$\forall r' \cdot \text{if } (s, h, r') \models \rho_1 \text{ then } (s, h, r \oplus r') \models \rho_2$

Fig. 5. Semantics of Assertions in the Resource-Aware Logic

with the capacity (a_l, a_u) then the remaining capacity is (c_l, c_u) . In other words, the executions allowed by (a_l, a_u) can be decomposed into executions required by (b_l, b_u) followed by executions required by (c_l, c_u) .

Definition 3 (Resource Capacity Splitting). *Given resource capacities $(a_l, a_u), (b_l, b_u)$ with $b_u \leq a_u$ and $a_l + b_u \leq a_u + b_l$ then $(a_l, a_u) \ominus (b_l, b_u) = (c_l, c_u)$ where*

$$c_l = \min\{x_l \in \mathbb{N}^\infty \mid x_l + b_l \geq a_l\} \text{ and } c_u = \max\{x_u \in \mathbb{N}^\infty \mid x_u + b_u \leq a_u\}.$$

From Defn. 3, (c_l, c_u) is the *largest* resource consumption allowed for any execution following executions satisfying (b_l, b_u) such that the overall resource consumption is described by (a_l, a_u) . Under this interpretation it follows naturally that when $b_u > a_u$ the splitting operation is undefined as c_u does not exist. In addition, when $a_l + b_u > a_u + b_l$, the splitting operation is also undefined as it would lead to $c_l > c_u$.

3.2 Assertion Language and Semantics for a Resource-Aware Logic

To support resource reasoning, we extend a minimalistic assertion language with two resource assertions, as shown in Fig. 4. We use v and v^* for variables and sequences of variables, $f(v^*)$ for functions from variables to \mathbb{N}^∞ , μ and Φ to represent resource-free formulas and ρ for resource assertions.

$\Psi ::= \bigvee (\exists v^* \cdot \mu \wedge \rho)^*$
$\Phi ::= \bigvee (\exists v^* \cdot \mu)^*$
$\rho ::= \text{RC}\langle a_l, a_u \rangle \mid \rho_1 \blacktriangleright \rho_2$
$a ::= f(v^*)$

Fig. 4. The Assertion Language

The resource assertion ρ ranges over (i) *atomic* resource assertions $\text{RC}\langle a_l, a_u \rangle$, where a_l, a_u are functions from variables to \mathbb{N}^∞ ; and (ii) *splitting* resource assertions $\rho_1 \blacktriangleright \rho_2$, which holds for states that allow executions to be split into two execution fragments, on which ρ_1 and ρ_2 hold respectively.

We concisely list in Fig. 5 the semantic model for the assertion language. We observe that the usual semantics of the logical connectives, e.g., conjunctions and disjunctions, lifts naturally over resource assertions. The semantics of the resource-free assertions is straightforward: a resource-free formula μ holds for all states (s, h, r) such that $(s, h) \models \mu$ with respect to the semantics of the corresponding underlying logic.

We point out that we have chosen to model the $\text{RC}\langle a_l, a_u \rangle$ assertion as a precise predicate. That is, a program state σ satisfies a resource constraint ρ if the resource capacity in σ is equal to the evaluation, in the context of σ , of the upper and lower functions associated with ρ . This modeling relation ensures that the resource assertion ρ is *precise* with regards to the resource capacity, where $(s, h, r) \models \rho$ does not imply

$(s, h, r') \models \rho$ whenever r' is larger than r , *i.e.*, $r' \geq_c r$. Consequently, $\text{RC}\langle a_l, a_u \rangle \vdash \text{RC}\langle b_l, b_u \rangle$ iff $(s, h) \models a_l=b_l \wedge a_u=b_u$. Additionally, $\text{RC}\langle a_l, a_u \rangle \wedge \text{RC}\langle b_l, b_u \rangle \equiv \text{RC}\langle a_l, a_u \rangle$ iff $a_l=b_l \wedge a_u=b_u$; otherwise, $\text{RC}\langle a_l, a_u \rangle \wedge \text{RC}\langle b_l, b_u \rangle \equiv \text{false}$.

To provide a precise modular resource reasoning, we lift the semantic split operation into a resource splitting assertion $\rho_1 \blacktriangleright \rho_2$. This enables our proof construction to follow the same style of other resource manipulating logics, such as separation logic. The intuition behind the splitting resource assertions is that $\rho_1 \blacktriangleright \rho_2$ holds for any program state from which it is possible to consume as many resources as ρ_1 requires and end in a state that satisfies ρ_2 . Or equivalently, $\rho_1 \blacktriangleright \rho_2$ holds for all states whose resource capacity can be split into two portions, such that the resulting capacities satisfy ρ_1 and ρ_2 , respectively. In addition, we can use \blacktriangleright to add a resource capacity ρ_1 into the current available resource capacity ρ , resulting in $\rho \blacktriangleright \rho_1$. The semantics of $\rho_1 \blacktriangleright \rho_2$ is also given in Fig. 5.

3.3 Resource-Enhanced Entailment with Frame Inference

Based on the semantics of resource assertions and the standard definition of the logical entailment relation (*i.e.*, $\Psi_1 \vdash \Psi_2$ iff $\forall \sigma \cdot$ if $\sigma \models \Psi_1$ then $\sigma \models \Psi_2$), it is possible to define an entailment for resource constraints of the form $\rho \vdash \rho_1 \blacktriangleright \rho_2$ as follows:

Lemma 1 (Resource Entailments). *Given resource assertions ρ , ρ_1 and ρ_2 , $\rho \vdash \rho_1 \blacktriangleright \rho_2$ iff $\forall s, h, r, r_1 \cdot$ if $(s, h, r) \models \rho$ and $(s, h, r_1) \models \rho_1$ then $(s, h, r \ominus r_1) \models \rho_2$.*

Proof. The proofs of all lemmas in this paper can be found in the technical report [28].

It follows that given \ominus_f , a lifting of resource capacity splitting to functions, then:

$$\frac{(\rho_l^2, \rho_u^2) = (\rho_l, \rho_u) \ominus_f (\rho_l^1, \rho_u^1)}{\text{RC}\langle \rho_l, \rho_u \rangle \vdash \text{RC}\langle \rho_l^1, \rho_u^1 \rangle \blacktriangleright \text{RC}\langle \rho_l^2, \rho_u^2 \rangle}$$

Entailments of the form $\rho \vdash \rho_1 \blacktriangleright \rho_2$ are of particular interest in the context of program verification as they naturally encode the restriction imposed at a method call and the remaining restriction after the execution of this method. For the proposed resource logic, we construct a general entailment system with frame inference by merging the entailment of resource constraints presented earlier with the entailment system corresponding to the underlying logic. Let the underlying entailment system be of the general form $\Psi \vdash \Phi \rightsquigarrow \Phi_r$ denoting that Ψ implies Φ with frame Φ_r . In sub-structural logics such as separation logic, the frame captures any residual state that is not required by the entailment. In pure logics where the program states are not changed, the frame is simply the antecedent of the entailment.

To support logics with disjunctions, the entailment system firstly deconstructs disjunctive antecedents (*e.g.*, using the rule $\boxed{\text{ENT-DISJ-LHS}}$) and consequents until formulas of the form $\mu \wedge \rho$ with a single resource constraint⁶ are encountered in both sides of the sub-entailments. The judgment system then applies the rule $\boxed{\text{ENT-CONJ}}$ that is slightly changed to handle resource constraints by splitting an entailment into two parts, namely *logical* part and *resource* part. The logical goal is solved by the entailment system $\mu_a \vdash \mu_c \rightsquigarrow \Phi_r$ of the underlying logic. The resource goal is solved by using

⁶ A conjunction of resource constraints can be simplified to either a single resource constraint or `false` as discussed in Sec. 3.2.

the resource entailment rules presented above. The solving process for the resource part leverages the entailment outcome Φ_r from the underlying logic, which is simply added to the antecedent of the resource entailment, to check the condition stated in Defn. 3 for the resource capacity splitting operation to be defined.

$$\frac{\boxed{\text{ENT-DISJ-LHS}} \quad \begin{array}{l} \Psi = \bigvee \exists v_i^* \cdot (\mu_i \wedge \rho_i) \\ \forall i \cdot (\mu_i \wedge \rho_i) \vdash \Phi \rightsquigarrow \Psi_r^i \\ \hline \Psi \vdash \Phi \rightsquigarrow \bigvee \exists v_i^* \cdot \Psi_r^i \end{array}}{\boxed{\text{ENT-CONJ}} \quad \begin{array}{l} \mu_a \vdash \mu_c \rightsquigarrow \Phi_r \\ \mu_a \wedge \Phi_r \wedge \rho_a \vdash \rho_c \blacktriangleright \rho_r \\ \hline \mu_a \wedge \rho_a \vdash \mu_c \wedge \rho_c \rightsquigarrow (\Phi_r \wedge \rho_r) \end{array}}$$

3.4 Hoare Logic for Resource Verification

Language. We provide a core strict imperative language with usual constructs, such as type declarations, method declarations, method calls, assignments, etc. to facilitate the verification for multiple front-end imperative languages. For simplicity, we choose a core language without while-loop constructs and assume a preprocessing step that applies an automatic translation into tail-recursive methods with reference-type parameters (declared by the keyword `ref`).

The pre and post conditions of a method are specified by the `requires` and `ensures` keywords, followed by logic formulas in the assertion language in Fig. 4. Resource-related assertions always appear in the method preconditions to denote resource requirements imposed on the caller for its execution. In contrast, resource assertions in the postconditions denote unspent/generated fuel returned to the caller, so that these assertions may not appear in the postconditions, depending on the analyzed resource. For example, as execution length (*i.e.*, a temporal resource) can only be consumed, it is safe and convenient to assume that the method consumes all the initially required resource; thus we can avoid the need for execution length related assertions in postconditions.

Hoare Logic. We observe that the resource consumption of each program statement is dependent on the tracked resource. As a result, the resource-aware Hoare logic needs to be adapted accordingly for each resource type. In terms of termination and non-termination reasoning, we are interested in the execution length as the tracked resource capacity. In the next section, we will construct a specific Hoare logic to reason about this resource.

4 (Non-)Termination Proofs via Resource Reasoning

For termination and non-termination reasoning, we have proposed three temporal constraints to capture: guaranteed termination `Term X`, guaranteed non-termination `Loop` and possible non-termination `MayLoop`, where X is a ranking function built from program variables. First, we define these constraints as resource capacity assertions, using the more general `RC` predicate. Next, we leverage the resource logic in Sec. 3, specialized in execution capacity, to construct a logic for termination and non-termination reasoning. A resource-based definition for the proposed temporal constraints is as follows:

Definition 4 (Temporal Constraints). *Temporal constraints are resource assertions over program execution lengths, such that $\text{Term } X \equiv \text{RC}\langle 0_f, \varpi \rangle$, $\text{Loop} \equiv \text{RC}\langle \infty_f, \infty_f \rangle$*

and $\text{MayLoop} \equiv \text{RC}\langle 0_f, \infty_f \rangle$ where 0_f and ∞_f denote the constant functions always returning 0 respectively ∞ . ϖ is a function of program variables to naturals, imposing a finite upper bound on the execution length of a terminating program.

Using the definition of resource entailments in Lemma 1, we formalize the set of valid entailments for temporal constraints below:

$$\begin{array}{l}
\text{MayLoop} \vdash \text{MayLoop} \blacktriangleright \text{MayLoop} \quad \text{Loop} \vdash \text{Term } X \blacktriangleright \text{Loop} \\
\text{MayLoop} \vdash \text{Term } X \blacktriangleright \text{MayLoop} \quad \text{Loop} \vdash \text{Loop} \blacktriangleright \text{MayLoop} \\
\text{MayLoop} \vdash \text{Loop} \blacktriangleright \text{MayLoop} \\
\text{Loop} \vdash \text{MayLoop} \blacktriangleright \text{Loop}
\end{array}
\quad
\frac{\mu \implies Y \leq_d X}{\mu \wedge \text{Term } X \vdash \text{Term } Y \blacktriangleright \text{Term } X -_d Y}$$

where \leq_d and $-_d$ are the ordering and the subtraction operation on the domain of the termination measures X and Y , respectively. All other decomposition attempts, such as $\text{Term } X \vdash \text{MayLoop} \blacktriangleright _$ and $\text{Term } X \vdash \text{Loop} \blacktriangleright _$, describe unfeasible splits. Thus in those cases, the entailment fails and an error is signaled.

4.1 From Termination Measures to Execution Capacity's Finite Upper Bounds

In Defn. 4, as X denotes a termination measure, a bounded function that decreases across recursive method calls, the resource upper bound ϖ must also follow. Thus, the mapping function from X to ϖ must be an *order-embedding* denoted by $\text{embed}(X)$. In our approach, the termination measure X is a list of arithmetic formulas over naturals since this formulation is simpler to write than a single but more complex termination measure and it can be used for a wider range of programs. In general, an order-embedding of lists of unbounded elements requires *ordinals*. However, transfinite ordinals are not suitable to model finite computational resources denoted by $\text{Term } X$.

By a *co-inductive* argument that every execution of a terminating method only computes finitely many different values, it follows that every non-negative element of a lexicographic termination measure applied to states of the corresponding call tree is upper-bounded. We then show that there always exists an order-embedding \mathcal{L} from the codomain of a termination measure (*i.e.*, tuples of bounded naturals) to naturals, such that $\text{embed}(X) = \mathcal{L} \circ X$.

Lemma 2. *If the termination of a program can be proven by a given lexicographic termination measure, then for each call tree τ of the program, every element of the termination measure applied to the program states corresponding to the nodes in the call tree τ is bounded.*

If every element x_i , where $0 \leq i \leq n - 1$, of a lexicographic termination measure $[x_n, x_{n-1}, \dots, x_0]$ corresponding to a given call tree τ is bounded by a constant k , we can use the base $b = k+1$ to construct a possible order-embedding function $\mathcal{D}([x_n, x_{n-1}, \dots, x_0]) = x_n * b^n + x_{n-1} * b^{n-1} + \dots + x_0$. The function \mathcal{D} preserves the order of the given measure along every trace of τ , as stated by Lemma 3.

Lemma 3. *For all $x_n, \dots, x_0, y_n, \dots, y_0 \in \mathbb{N}$ such that $\forall i \in \{0..n-1\} \cdot x_i, y_i < b$, $[x_n, \dots, x_0] >_l [y_n, \dots, y_0]$ iff $\mathcal{D}([x_n, \dots, x_0]) > \mathcal{D}([y_n, \dots, y_0])$, where $>_l$ is the lexicographic ordering.*

$\frac{\text{CheckMin}(\Psi_1) \quad \text{CheckMin}(\Psi_2)}{\text{CheckMin}(\Psi_1 \vee \Psi_2)}$	$\frac{\mu \vdash \rho_l = 0}{\text{CheckMin}(\mu \wedge \text{RC}\langle \rho_l, \rho_u \rangle)}$
$\boxed{\text{FV-CALL}}$	
$t_0 \text{ mn}((t v)^*) (\Psi_{\text{Pre}}, \Phi_{\text{Post}}) \{code\} \in \text{Prog}$	$\boxed{\text{FV-RET}}$
$\frac{\Psi \vdash \text{RC}\langle 1, 1 \rangle \rightsquigarrow \Theta \quad \Theta \vdash \Psi_{\text{Pre}} \rightsquigarrow \Phi \quad \Psi_r = \Phi \wedge \Phi_{\text{Post}}}{\vdash \{\Psi\} \text{ mn}(v^*) \{\Psi_r\}}$	$\frac{\text{CheckMin}(\Psi)}{\vdash \{\Psi\} \text{ return } v \{\Psi \wedge \text{res} = v'\}}$

Fig. 6. Hoare Verification Rules: Method Call and Return

In general, such a bounded constant k for a call tree τ can be determined by a function \mathcal{K} of initial values of the call tree's variables. Since the execution of a loop has only a single trace, the order-embedding \mathcal{D} , constructed from the constant k , would be enough to ensure the sufficiency of execution capacity for the loop. However, in order to give a proper estimate of the execution capacity for more complex recursion patterns, especially when the termination measures are based on the depth of recursion, we propose using a more refined embedding for a call tree, that is $\mathcal{L} = \begin{cases} \mathcal{D} & \mathcal{N} \leq 1 \\ \mathcal{N}^{\mathcal{D}} & \mathcal{N} > 1 \end{cases}$, where \mathcal{N} is the maximum number of children for each node of the call tree.

Therefore, given the termination measure X of a terminating program, there always exists an order-embedding \mathcal{L} from the codomain of X to naturals. The function \mathcal{L} can be constructed from initial values of program variables and the call trees corresponding to these initial values. As a result, $\text{embed}(X) = \mathcal{L} \circ X$ is a function from program variables to naturals, which describes an upper bound on the number of method calls taken by any execution of the program.

4.2 Termination and Non-termination Verification

Here we elaborate on the construction of both termination and non-termination proofs based on Defn. 4 and the verification framework in Fig. 6 for tracking execution length as resource. Although execution length can be tracked at various levels of granularities, we choose to track it only at method calls (*i.e.*, as the total number of method calls) in order to simplify the verification rules and the operational semantics. In Fig. 6, we only outline the Hoare logic rules for the method call and the return statements, which are especially relevant to the verification of execution lengths as they encode the resource consumption. The Hoare rules for other constructs are standard because they do not interact with the resource of interest.

As a standard preprocessing step, we check that all predicate invariants are satisfied, including the invariants of resource constraints: the resource assertion $\text{RC}\langle \rho_l, \rho_u \rangle$ in precondition Ψ_{Pre} is consistent if $0 \leq \rho_l \leq \rho_u$, that is, for each disjunct $\mu \wedge \text{RC}\langle \rho_l, \rho_u \rangle$ of Ψ_{Pre} it follows that $\mu \vdash \rho_u \geq \rho_l \wedge \rho_l \geq 0$. We observe that the invariant check on $\text{Term } X$ requires that every element of X be non-negative to ensure a non-negative upper-bound $\mathcal{L} \circ X$, so that the execution capacity satisfies the invariant $0 \leq 0_f \leq \mathcal{L} \circ X$.

In the method call rule $\boxed{\text{FV-CALL}}$, the available execution capacity is first decreased by one step, denoted by $\text{RC}\langle 1, 1 \rangle$, to account the cost of method call, followed by a check that the callee's requirements are met. This check is translated into an entailment for proving the method precondition. Finally, the poststate after this method call is computed. With the help of the resource-enhanced entailment system introduced in

Sec. 3.3, both logical and resource proving are combined into one entailment, resulting in a standard-looking Hoare rule for method call.

In addition, specifically for temporal constraints, two entailments $\Psi \vdash \text{RC}\langle 1, 1 \rangle \rightsquigarrow \Theta$ and $\Theta \vdash \Psi_{\text{Pre}} \rightsquigarrow \Phi$ can be combined into $\Psi \vdash \Psi_{\text{Pre}} \rightsquigarrow \Phi$ by using a new entailment \vdash_t for temporal constraints.

Definition 5 (Unit Reduction Temporal Entailments). *Given temporal constraints θ , θ_1 and θ_2 , $\theta \vdash_t \theta_1 \blacktriangleright \theta_2$ iff $\forall s, h, r \cdot$ if $(s, h, r) \models \theta$ then $(s, h, r \ominus (1, 1)) \models \theta_1 \blacktriangleright \theta_2$.*

Therefore, if θ is Loop or MayLoop then $\theta \vdash_t \theta_1 \blacktriangleright \theta_2$ iff $\theta \vdash \theta_1 \blacktriangleright \theta_2$. If θ is Term X then $\mu \wedge \text{Term } X \vdash_t \text{Term } Y \blacktriangleright \text{Term } ((X -_d 1_d) -_d Y)$ if $\mu \implies Y <_d X$, where 1_d is the unit of termination measures' domain. Basically, the check $Y <_d X$ is equivalent to the check that termination measures are decreasing across recursive method calls in the traditional termination proof. By introducing the temporal entailment \vdash_t , we obtain a resource-based temporal logic which is related to only the temporal constraints and thus the underlying resource reasoning becomes implicit.

In the method return rule $\boxed{\text{FV-RET}}$, the CheckMin predicate, which is also defined in Fig. 6, ensures that the specified minimum computation resource has been completely consumed when the method returns. Note that if the method does not terminate, the minimum guaranteed execution length is always satisfied since the actual return point is never reached. For temporal constraints, CheckMin holds for any Term X and MayLoop as the lower bounds in their execution capacities are always 0. In non-termination cases, CheckMin($\mu \wedge \text{Loop}$) only holds when μ is unsatisfiable. This check ensures that a return statement cannot be executed/reachable from a state satisfying Loop.

We now state the soundness of this resource-aware Hoare logic as follows:

Theorem 1. *The standard Hoare rules (e.g., assignment, conditional, sequential composition) and the Hoare rules for method call and return are sound.*

Proof. The proof can be found in the technical report [28].

4.3 Flow-Insensitive Temporal Logic

Observe that the current formulation of the temporal logic with temporal constraints is *flow-sensitive* since the entailment $\theta \vdash_t \theta_1 \blacktriangleright \theta_2$ might return a residue θ_2 distinct from θ . However, with the following observations, we can formalize a *flow-insensitive* version of the temporal logic and provide a further abstraction on the resource-based framework presented so far.

First, it is possible to refine the granularity of the termination and non-termination verification by tracking only execution lengths of (mutually) recursive method calls. Second, using König's lemma [27], it is sufficient to inspect individual execution traces in the call tree for deciding just termination or non-termination, instead of tracking the total execution length of all traces in the call tree. That is, a program terminates iff every execution trace is finite; otherwise, the program is non-terminating.

Based on these observations, the tracked resource will be abstracted to capture the execution capacity required for the longest trace in the call tree, instead of the execution capacity required for the remaining program. With this, the resource (for the longest trace allowed) remains unchanged after each splitting operation, which determines the residue resource needed for subsequent method calls. Thus, for every method, we endeavor to provide a single abstract resource that is sufficient for executing a given method call and also its remaining code sequences.

Benchmarks	Programs	Term	Loop	MayLoop	PC(s)	TC(s)	Overhead (%)
Invel	59	137	81	12	14.88	15.96	6.77
AProVE	124	534	120	8	15.73	17.21	8.60
Pasta	44	219	10	3	4.95	5.79	14.51
Others	48	194	32	22	7.35	8.78	16.29
Totals/(%)	275	1084 (79.0%)	243 (17.7%)	45 (3.3%)	42.91	47.74	10.12%

Fig. 7. Termination Verification for Numerical Programs

By using this abstraction, we can obtain a formulation on temporal entailment that ensures $\theta \vdash_t \theta_1 \blacktriangleright \theta$ whereby the temporal constraint in residue is always identical to the one in the antecedent. Hence, the operator $-_d$ can be fully circumvented. Moreover, the finite upper bound ϖ used for the definition of Term X in Defn. 4 can be determined as $\varpi = \mathcal{D} \circ X$, instead of the larger $\mathcal{L} \circ X$. As a result, without any change to the Hoare rules, during a method’s verification, the same initial resource capacity is used for the verification of call traces and thus facilitating a simpler verification procedure for temporal constraint. As a direct outcome of this abstraction, the temporal assertions Loop, MayLoop and Term X are now *flow-insensitive*, and therefore closer to the pure logic form, as opposed to the sub-structural form of resource logics. Note that flow-insensitive label applies to only the temporal constraints. In general, program states (*e.g.*, denoted by separation logic as the underlying logic) remain flow-sensitive since they might be changed due to changes on heap state and program variables.

5 Experiments

We have implemented the proposed termination logic into an automated verification system, called *HipTNT*. The integration of the termination logic into an existing system allows us to utilize the infrastructure that has been developed for some richer specification logics, such as separation logic, beyond a simple first-order logic. Consequently, we are able to *specify* and *verify* both termination and non-termination properties, in addition to correctness properties for a much wider class of programs, including heap-manipulating programs. In this system, the final proof obligations are automatically discharged by off-the-shelf provers, such as Z3 [16]. The expressivity of our new integrated logic is shown in the following experimental results, in which the lexicographic order is needed for about 25% of the considered programs in our experiments.

5.1 Numerical Programs

HipTNT was evaluated using a benchmark of over 200 small numerical programs selected from a variety of sources: (i) from the literature, such as [13,11], (ii) from benchmarks used by other systems (that are AProVE [19], Invel [37] and Pasta [18]) and (iii) some realistic programs, such as the Microsoft Zune’s clock driver that has a leap-year non-termination bug. Most of the methods in these benchmark programs contain either terminating or non-terminating code fragments, expressed in (mutual) recursive calls or (nested) loops. To construct these benchmarks we added the novel termination specifications to the original examples from the analysis tools for termination and non-termination. We have chosen these benchmarks in order to show the usability and

practicality of our approach. A comparison with these tools would be of less relevance as our proposal focuses on verifying the given specifications rather than infer them.

Fig. 7 summarizes the characteristics and the verification times for a benchmark of numerical programs. Columns 3-5 describe the number of preconditions that have been specified and successfully verified as *terminating*, *non-terminating* or *unknown*, respectively. As hoped for, the number of preconditions annotated by MayLoop occupies the smallest fragment (about 3%) of the total number of preconditions. Such MayLoop constraints were only used in some unavoidable scenarios, such as (non-)termination depends on unpredictable user input or non-deterministic assignments to variables. In contrast, the `Term` constraints (with the given measures) are in the majority because most of the methods are expected to be terminating, except for the Invel benchmark which focuses on mostly non-terminating programs.

Our verification system can perform both correctness and termination proofs. Column 7 (**TC**) gives the total timings (in seconds) needed to perform both termination and correctness proofs for all the programs in each row, while column 6 (**PC**) gives the timings needed for just correctness proofs. The difference in the two timings represents the small overhead needed for termination and non-termination reasoning.

5.2 Heap-Manipulating Programs

We have also conducted termination reasoning on our own benchmark of heap-based programs using various data structures, on problems under Java Bytecode categories of the Termination Competition [33] and on some medium programs taken from the SIR/Siemens test suite [17]. Due to the tight integration with the underlying logics, the task of specifying and verifying the termination properties was easy even though some of the programs use non-trivial data structures (*e.g.*, Red-Black and AVL-trees), or non-linear constraints (*e.g.*, the BigNat program, which implements infinite precision natural numbers (by linked lists) with procedures for some arithmetic operations, in addition to a fast multiplication method based on the Karatsuba algorithm). We report that the termination and expected non-termination of all programs in these benchmarks are verified successfully with a small overhead ($< 5\%$). Due to space limitation, the detailed experimental results and discussion on these benchmarks are put in [28].

6 Related Work and Conclusion

There exists a rich body of related works on automatic analysis for termination [30,8,14], non-termination [21,37,10], and both [19]. However, they consider termination and non-termination reasoning as distinct from functional correctness reasoning. Therefore, these works cannot leverage the result of functional correctness analysis to conduct more intricate (non-)termination reasoning. Recently, Brockschmidt *et al.* [9] propose a cooperation between safety and termination analysis to find sufficient supporting invariants for the construction of termination arguments but not considering non-termination. Chen *et al.* [12] introduce a similar approach for proving only non-termination. Our proposal complements these works since our aim is to construct a logic where termination and non-termination properties are directly integrated into specification logics, and thus utilize the available infrastructure on functional correctness proofs. We have

achieved this, and have also successfully evaluated its applicability on a wide range of programs, covering both numerical and heap-based programs.

Related to resource verification, [4] introduces a resource logic for a low-level language. While this logic avoids the need of auxiliary counters, it redefines the semantic model of the underlying logic to track the resource consumption via logical assertions, making the proposal harder to retrofit to other logics. Moreover, this logic only targets partial correctness, so that it does not take into account infinite resource consumption.

There are some works that are based on the well-foundedness of inductive definitions of heap predicates [7,11] or user-defined quantitative functions over data structures [20] to prove termination of heap-manipulating programs. On one hand, they do not require any explicit ranking function. On the other hand, these approaches might have problems with programs like the Karatsuba multiplication method, in which the arguments of the recursive calls are not substructures of the input lists. In addition, the automated tools, such as AProVE and COSTA, cannot prove the termination of this method. In contrast, our approach is more flexible as it allows explicit termination measures, that are possibly non-linear, for proving programs' termination. These termination measures can be constructed from not only the heap structures but also the values of the data structures' elements. For example, we use the actual value of the natural presented by a linked list to bound the execution of the Karatsuba method. Moreover, we also allow non-termination to be specified and verified for these programs. We believe that relatively complex examples, such as the Karatsuba method, highlight the benefits of our approach, which trades a lower level of automation but gains additional power.

The comparison of our approach with the other specification languages, *i.e.* Dafny [31], JML [29], etc., has been discussed in Sec. 1. Another closely related work to ours is that of Nakata and Uustalu [34]. In this work, a Hoare logic for reasoning about non-termination of simple While programs (without method calls) was introduced. The logic is based on a trace-based semantics, in which the infiniteness of non-terminating traces is defined by coinduction. However, induction is still needed to define the finiteness of traces. In contrast, with resources, we can unify the semantics of the proposed termination and non-termination temporal constraints and allow the Hoare logic for functional correctness to be enhanced for termination and non-termination reasoning with minor changes. Moreover, our logic allows interprocedural verification in a modular fashion.

Conclusion. Termination reasoning has been intensively studied in the past, but it remains a challenge for the technology developed there to keep up with improvements to specification logic infrastructure, and vice versa. We propose an approach that would combine the two areas more closely together, through a tightly coupled union. Our unique contribution is to embed both termination and non-termination reasoning directly into specification logics, and to do so with the help of temporal entailment. We also show how its properties can be captured by a resource logic based on execution capacity, and how it could be abstracted into a flow-insensitive temporal logic. We believe this approach would have benefits. Its expressiveness is immediately enhanced by any improvement to the underlying logics. It can also benefit from infrastructures that have been developed for the underlying logics, including those that are related to program analysis. In particular we believe that a possible future avenue for investigation is to use the safety specifications as a basis for termination specification inference. Last, but not

least, it has placed termination and non-termination reasoning as a first-class concept, much like what was originally envisioned by Hoare's logic for total correctness.

Acknowledgement. This work was supported by MoE/NUS research project R-252-000-469-112.

References

1. Abel, A.: Type-based termination of generic programs. *SCP* 74(8) (2009)
2. Ahrendt, W., Baar, T., Beckert, B., Bubel, R., Giese, M., Hähnle, R., Menzel, W., Mostowski, W., Roth, A., Schlager, S., Schmitt, P.H.: The KeY tool. *SSM* 4 (2005)
3. Albert, E., Arenas, P., Genaim, S., Puebla, G., Zanardini, D.: COSTA: Design and Implementation of a Cost and Termination Analyzer for Java Bytecode. In: de Boer, F.S., Bonsangue, M.M., Graf, S., de Roever, W.-P. (eds.) *FMCO 2007*. LNCS, vol. 5382, pp. 113–132. Springer, Heidelberg (2008)
4. Aspinall, D., Beringer, L., Hofmann, M., Loidl, H.-W., Momigliano, A.: A program logic for resources. *TCS* 389(3) (2007)
5. Atkey, R.: Amortised resource analysis with separation logic. *LMCS* 7(2) (2011)
6. Baudin, P., Cuoq, P., Filliâtre, J.-C., Marché, C., Monate, B., Moy, Y., Prevosto, V.: *ACSL: ANSI/ISO C Specification Language Version 1.8*. (2013), <http://frama-c.com/acsl.html>
7. Berdine, J., Cook, B., Distefano, D., O'Hearn, P.W.: Automatic termination proofs for programs with shape-shifting heaps. In: Ball, T., Jones, R.B. (eds.) *CAV 2006*. LNCS, vol. 4144, pp. 386–400. Springer, Heidelberg (2006)
8. Bradley, A.R., Manna, Z., Sipma, H.B.: The polyranking principle. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) *ICALP 2005*. LNCS, vol. 3580, pp. 1349–1361. Springer, Heidelberg (2005)
9. Brockschmidt, M., Cook, B., Fuhs, C.: Better termination proving through cooperation. In: Sharygina, N., Veith, H. (eds.) *CAV 2013*. LNCS, vol. 8044, pp. 413–429. Springer, Heidelberg (2013)
10. Brockschmidt, M., Ströder, T., Otto, C., Giesl, J.: Automated detection of non-termination and `NullPointerException`s for java bytecode. In: Beckert, B., Damiani, F., Gurov, D. (eds.) *FoVeOOS 2011*. LNCS, vol. 7421, pp. 123–141. Springer, Heidelberg (2012)
11. Brotherston, J., Bornat, R., Calcagno, C.: Cyclic proofs of program termination in separation logic. In: *POPL* (2008)
12. Chen, H.-Y., Cook, B., Fuhs, C., Nimkar, K., O'Hearn, P.: Proving nontermination via safety. In: Ábrahám, E., Havelund, K. (eds.) *TACAS 2014 (ETAPS)*. LNCS, vol. 8413, pp. 156–171. Springer, Heidelberg (2014)
13. Cook, B., Gulwani, S., Lev-Ami, T., Rybalchenko, A., Sagiv, M.: Proving conditional termination. In: Gupta, A., Malik, S. (eds.) *CAV 2008*. LNCS, vol. 5123, pp. 328–340. Springer, Heidelberg (2008)
14. Cook, B., Podelski, A., Rybalchenko, A.: Termination proofs for systems code. In: *PLDI* (2006)
15. Cuoq, P., Kirchner, F., Kosmatov, N., Prevosto, V., Signoles, J., Yakobowski, B.: *Frama-C*. In: Eleftherakis, G., Hinchey, M., Holcombe, M. (eds.) *SEFM 2012*. LNCS, vol. 7504, pp. 233–247. Springer, Heidelberg (2012)
16. de Moura, L., Bjørner, N.S.: *Z3: An Efficient SMT Solver*. In: Ramakrishnan, C.R., Rehof, J. (eds.) *TACAS 2008*. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008)
17. Do, H., Elbaum, S.G., Rothermel, G.: Supporting Controlled Experimentation with Testing Techniques: An Infrastructure and its Potential Impact. In: *ESE*, vol. 10 (2005)

18. Falke, S., Kapur, D.: A term rewriting approach to the automated termination analysis of imperative programs. In: Schmidt, R.A. (ed.) CADE-22. LNCS, vol. 5663, pp. 277–293. Springer, Heidelberg (2009)
19. Giesl, J., Schneider-Kamp, P., Thiemann, R.: aProVE 1.2: automatic termination proofs in the dependency pair framework. In: Furbach, U., Shankar, N. (eds.) IJCAR 2006. LNCS (LNAI), vol. 4130, pp. 281–286. Springer, Heidelberg (2006)
20. Gulwani, S., Mehra, K.K., Chilimbi, T.: Speed: Precise and efficient static estimation of program computational complexity. In: POPL (2009)
21. Gupta, A., Henzinger, T.A., Majumdar, R., Rybalchenko, A., Xu, R.-G.: Proving non-termination. In: POPL (2008)
22. Hehner, E.C.R.: Termination is timing. In: van de Snepscheut, J.L.A. (ed.) MPC 1989. LNCS, vol. 375, pp. 36–47. Springer, Heidelberg (1989)
23. Hehner, E.C.R.: Specifications, programs, and total correctness. SCP 34(3) (1999)
24. Hoare, C.A.R., He, J.: Unifying Theories of Programming. Prentice-Hall (1998)
25. Hoffmann, J., Hofmann, M.: Amortized resource analysis with polynomial potential. In: Gordon, A.D. (ed.) ESOP 2010. LNCS, vol. 6012, pp. 287–306. Springer, Heidelberg (2010)
26. Jones, C.B.: Balancing expressiveness in formal approaches to concurrency (2013)
27. Kleene, S.: Mathematical logic. Wiley (1967)
28. Le, T.C., Gherghina, C., Hobor, A., Chin, W.-N.: A Resource-Based Logic for Termination and Non-Termination Proofs, Technical Report (2014), <http://loris-7.ddns.comp.nus.edu.sg/~project/hiptnt/HipTNT.pdf>
29. Leavens, G.T., Baker, A.L., Ruby, C.: Preliminary design of JML: A behavioral interface specification language for Java. SIGSOFT Softw. Eng. Notes 31(3), 1–38 (2006)
30. Lee, C.S., Jones, N.D., Ben-Amram, A.M.: The size-change principle for program termination. In: POPL (2001)
31. Leino, K.R.M.: Dafny: An automatic program verifier for functional correctness. In: Clarke, E.M., Voronkov, A. (eds.) LPAR-16 2010. LNCS, vol. 6355, pp. 348–370. Springer, Heidelberg (2010)
32. Magill, S., Tsai, M.-H., Lee, P., Tsay, Y.-K.: Automatic numeric abstractions for heap-manipulating programs. In: POPL (2010)
33. Marché, C., Zantema, H.: The termination competition. In: Baader, F. (ed.) RTA 2007. LNCS, vol. 4533, pp. 303–313. Springer, Heidelberg (2007)
34. Nakata, K., Uustalu, T.: A hoare logic for the coinductive trace-based big-step semantics of while. In: Gordon, A.D. (ed.) ESOP 2010. LNCS, vol. 6012, pp. 488–506. Springer, Heidelberg (2010)
35. Nguyen, H.H., David, C., Qin, S.C., Chin, W.-N.: Automated Verification of Shape and Size Properties Via Separation Logic. In: Cook, B., Podelski, A. (eds.) VMCAI 2007. LNCS, vol. 4349, pp. 251–266. Springer, Heidelberg (2007)
36. Reynolds, J.: Separation Logic: A Logic for Shared Mutable Data Structures. In: LICS (2002)
37. Velroyen, H., Rümmer, P.: Non-termination checking for imperative programs. In: Beckert, B., Hähnle, R. (eds.) TAP 2008. LNCS, vol. 4966, pp. 154–170. Springer, Heidelberg (2008)
38. Xi, H.: Dependent Types for Program Termination Verification. In: LICS (2001)