

Data Refinement with Low-level Pointer Operations

Ivana Mijajlović¹ and Hongseok Yang^{2*}

¹ Queen Mary, University of London, UK

² ERC-ACI, Seoul National University, South Korea

Abstract. We present a method for proving data refinement in the presence of low-level pointer operations, such as memory allocation and deallocation, and pointer arithmetic. Surprisingly, none of the existing methods for data refinement, including those specifically designed for pointers, are sound in the presence of low-level pointer operations. The reason is that the low-level pointer operations allow an additional potential for obtaining the information about the implementation details of the module: using memory allocation and pointer comparison, a client of a module can find out which cells are internally used by the module, even without dereferencing any pointers. The unsoundness of the existing methods comes from the failure of handling this potential. In the paper, we propose a novel method for proving data refinement, called power simulation, and show that power simulation is sound even with low-level pointer operations.

1 Introduction

Data refinement [7] is a process in which the concrete representation of some abstract module is formally derived. Viewed from outside, the more concrete representation behaves the same as (or better than) the given abstract module. Thus, data refinement ensures that for every program, we can replace a given abstract module by the concrete one, while preserving (or even improving) the observable behavior of the program.

Our aim here is to develop a method of data refinement in the presence of *low-level* pointer operations, such as memory allocation and deallocation, and pointer arithmetic. Developing such methods is challenging, because low-level pointer operations allow subtle ways for accessing the internals of a module; without protecting the module internals from these accessing mechanisms (and thus ensuring that the module internals are only accessed by the module operations), we cannot have a sound method of data refinement.

The best-known accessing mechanism is the dereference of cross-boundary pointers. If a client program knows the location of some internal heap cell of a module, which we call a *cross-boundary pointer*, it can directly read or write that

* Yang was supported by R08-2003-000-10370-0 from the Basic Research Program of Korea Science & Engineering Foundation, and Brain Korea 21 project in 2005.

<pre> module counter1 { init() { *1=allocCell2(); **1=0; } inc() { **1=(**1)+1; } read() { *3=(**1); } final() { free(*1); *1=0; } } </pre>	<pre> module counter2 { init() { *1=0; } inc() { *1=(*1)+1; } read() { *3=(*1); } final() { *1=0; } } </pre>	<pre> module counter3 { init() { *1=alloc(); **1=0; } inc() { **1=(**1)+1; } read() { *3=(**1); } final() { free(*1); *1=0; } } </pre>
---	--	--

Fig. 1. Counter Modules

internal cell by dereferencing that location. Thus, such a program can detect the changes in the representation of a module, and invalidate the standard methods of data refinements. This problem of cross-boundary pointers is well known, and several methods for data refinement have been proposed specifically to solve this problem [12, 17, 1, 3, 2].

However, none of the existing data-refinement methods, including the ones designed for cross-boundary pointers, can handle another accessing mechanism, which we call *allocation-status testing*. This mechanism uses the memory allocator and pointer comparison (with specific integers) to find out which cells are used internally by a module. A representative example `check2` that implements this mechanism is `z=alloc(); if (z==2) then v=1 else v=2`. Assume that the memory allocator `alloc` nondeterministically chooses one inactive cell, and allocates the chosen cell. Under this assumption, `check2` can detect whether cell 2 is used internally by a module or not. If a module is currently using cell 2, the newly allocated cell in `check2` has to be different from 2, so that `check2` always assigns 2 to v . On the other hand, if a module is not using cell 2, so cell 2 is free, then the memory allocation in `check2` may or may not choose cell 2, and so, the variable v nondeterministically has value 1 or 2. Thus, by changing its nondeterministic behavior, `check2` “observes” the allocation status of cell 2.

Protecting the module internals from the allocation-status testing is crucial for sound data refinement; using the allocation-status testing, a client can detect space-optimizing data refinements. We explain the issue with the first two counter modules, `counter1` and `counter2`, in Fig. 1. Both modules implement a counter “object” with operations for incrementing the counter (`inc`) or reading the value of the counter (`read`). The main difference is that the second module uses less space than the first module. Let `allocCell2()` be a memory allocator that always selects cell 2: if cell 2 is inactive, `allocCell2()` allocates the cell; otherwise, i.e., if 2 is already allocated, then `allocCell2()` diverges. The first module is initialized by allocating cell 2 (`allocCell2()`) and storing the value of the counter in the allocated cell 2. The address of this newly allocated cell, namely 2, is kept in cell 1. On the other hand, the second module uses only cell 1, and stores the counter value directly to cell 1. The space-saving optimization in `counter2` can be detected by the command `check2` in the previous paragraph. When `check2` is run with `counter1`, it always assigns 1 to v , but when `check2` is run with the other module `counter2`, it can nondeterministically assign 1 or 2 to v . Thus, the optimization in `counter2` is not correct, because it generates a new behavior of the client program `check2`.

Here, we present a data-refinement method that handles both cross-boundary pointers and allocation-status testing. Our method is based on Mijajlović *et al.*'s

technique [12], which ensures correct data refinement in the presence of cross-boundary pointers, but as stressed there, not with allocation status testing. We provide a more general method which can cope well with both problems. The key idea of our method is to restrict the space optimization of a concrete module to nondeterministically allocated cells only, in order to hide the identities of the optimized cells from a client program, by making all the allocation-status testing fail to give any useful information. For instance, our method allows `counter3` in Fig. 1 to be optimized by `counter2`, because the internal cell in `counter3` is allocated nondeterministically. Note that even with `check2`, a client cannot detect this optimization, e.g. when cell 2 is free initially, `counter3.init(); detect2` nondeterministically assigns 1 or 2 to v , just as `counter2.init(); detect2` does. The precise formulation of our method uses a new notion of simulation – *power simulation*, to express this restriction on space optimization.

Related Work and Motivation It has long been known that pointers cause great difficulties in the treatment of data abstraction [8, 9], and this has led on to a non-trivial body of research [1, 3, 14, 11, 19, 17]. The focus of the present work (and [12]), on problems caused by low-level operations, sets it apart from all this other research.

Now, the reader might think that these problems arise only because of language bugs. Indeed, previous work has relied strongly on protection mechanisms of high-level, garbage collected languages. In such high-level languages, the nondeterministic memory allocation is harmless; it does not let one implement the allocation-status testing (because those languages forbid explicit deallocation and pointer arithmetic) and the nondeterministic allocation can even be treated deterministically using location renaming [19, 17]. Moreover, those high-level languages often have sophisticated type systems [3, 2] that limit cross-boundary pointers. However, we would counter that a comprehensive approach to abstraction cannot be based on linguistic restrictions. For, the fact of the existence of significant suites of infrastructure code – operating systems, database servers, network servers – argues against it. The architecture of this code is not enforced by linguistic mechanisms, and it is hard to see how it could be. Low-level code naturally uses cross-boundary pointers and address arithmetic. But it is a mistake to think that infrastructure code is unstructured; it often exhibits a large degree of pre-formal modularity. In this paper, we will demonstrate that there is no inherent reason why the *idea* of refinement of modules should not be applicable to it.

Outline We start the paper by defining the storage model and the programming language in Sec. 2 and 3. Then, in Sec. 4, we describe the problem of finding a sound data-refinement method, and show that the usual forward method of data refinement fails to be a solution for the problem. In Sec. 5, we introduce the notion of *power simulation*, and prove its soundness; so, power simulation is a solution for the problem. Finally, in Sec. 6, we conclude the paper. The missing proofs of lemmas and propositions appear in the full version of the paper [13].

2 Storage Model and Finite Local Action

Our storage model, St , is the RAM model in separation logic [18, 10]:

$$\text{Loc} = \{1, 2, \dots\} \quad \text{Int} = \{\dots, -2, -1, 0, 1, \dots\} \quad \text{St} = \text{Loc} \rightarrow_{\text{fin}} \text{Int}$$

A state $h \in \text{St}$ in the model is a finite mapping from locations to integer values; the domain of h denotes the set of currently allocated memory cells, and the “action” of h the contents of those allocated cells. Note that addresses are positive natural numbers, and so, they can be manipulated by arithmetic operations. We recall the disjointness predicate $h\#h'$ and the (partial) heap combining operator $h \cdot h'$ from separation logic. The predicate $h\#h'$ means that $\text{dom}(h) \cap \text{dom}(h') \neq \emptyset$; and, $h \cdot h'$ is defined only for such disjoint heaps h and h' , and in that case, it denotes the combined heap $h \cup h'$. We overload the disjointness predicate $\#$, and for states h and location sets L , we write $h\#L$ to mean that all locations in L are free in h (i.e., $\text{dom}(h) \cap L = \emptyset$).

We specify a property of storage, using subsets of St directly, instead of syntactic formulas. We call such subsets of St *predicates*, and use semantic versions of separating conjunction $*$ and preciseness from separation logic:

$$p, q \in \text{Pred} \stackrel{\text{def}}{=} \wp(\text{St}) \quad p * q \stackrel{\text{def}}{=} \{h_p \cdot h_q \mid h_p \in p \wedge h_q \in q\} \quad \text{true} \stackrel{\text{def}}{=} \text{St}$$

p is precise $\stackrel{\text{def}}{\iff}$ for all h , there is at most one splitting $h_p \cdot h_0 = h$ of h s.t. $h_p \in p$.

An *action* r is a relation from St to $\text{St} \cup \{\text{av}, \text{flt}\}$. Intuitively, it denotes a nondeterministic client program that uses a module. Action r can output two types of errors, access violation av and memory fault flt . The first error av means that a client attempts to break the boundary between the client and the module, by accessing the internals of the module directly without using module operations. The second one, flt , means that a client tries to dereference a null or a dangling pointer. Note that if $\neg h[r]\text{flt}$, state h contains all the cells that r dereferences, except the newly allocated cells. As in separation logic, we write $\text{safe}(r, h)$ to indicate this (i.e., $\neg h[r]\text{flt}$).

A *finite local action* is an action that satisfies: *safety monotonicity*, *frame property*, *finite access property*, and *contents independence*. Intuitively, these four properties mean that each execution of the action accesses only finitely many heap cells. Some of the cells are accessed directly by pointer dereferencing, so that the contents of the cells affects the execution, while the other remaining cells are accessed only indirectly by the allocation-status testing, so that the execution only depends on the allocation status of the cells, not their contents. More precisely, we define the four properties as follows:¹

- **Safety Monotonicity:** if $h_0\#h_1$ and $\text{safe}(r, h_0)$, then $\text{safe}(r, h_0 \cdot h_1)$.
- **Frame Property:** if $\text{safe}(r, h_0)$ and $h_0 \cdot h_1[r]h'$, then $\exists h'_0. h' = h'_0 \cdot h_1 \wedge h_0[r]h'_0$.
- **Finite Access Property:** if $\text{safe}(r, h_0)$ and $h_0[r]h'_0$, then

$$\exists L \subseteq_{\text{fin}} \text{Loc}. \forall h_1. (h_1\#h_0 \wedge h_1\#h'_0 \wedge (\text{dom}(h_1) \cap L = \emptyset)) \Rightarrow h_0 \cdot h_1[r]h'_0 \cdot h_1.$$
- **Contents Independence:** if $\text{safe}(r, h_0)$ and $h_0 \cdot h_1[r]h'_0 \cdot h_1$, then $h_0 \cdot h_2[r]h'_0 \cdot h_2$ for all states h_2 with $\text{dom}(h_1) = \text{dom}(h_2)$.

¹ All the states free in the properties are universally quantified.

The first two properties are well-known locality properties from separation logic, and mean that if h_0 contains all the directly accessed cells by a “command” r , every computation from a bigger state $h_0 \cdot h_1$ is safe, and it can be tracked by some computation from the smaller state h_0 . The third condition expresses the converse; every computation from the smaller state h_0 can be extended to a computation from the bigger state $h_0 \cdot h_1$, as long as the extended part h_1 does not include directly accessed locations ($h_1 \# h_0 \wedge h_1 \# h'_0$) or indirectly accessed locations (i.e., $\text{dom}(h_1) \cap L = \emptyset$). Note that the finite set L contains all the indirectly accessed locations by the computation $h_0[r]h'_0$. The last one, contents independence, expresses that if $\text{safe}(r, h_0)$, the execution of r from a bigger state $h_0 \cdot h_1$ does not look at the contents of cells in h_1 ; it can only use the information that the locations in h_1 are allocated initially. At first glance, it may seem that contents independence follows from the frame property, but the following example suggests otherwise. Let \square be the empty state, and let r be an action defined by $h[r]v \Leftrightarrow h = v \wedge (h = \square \vee (1 \in \text{dom}(h) \wedge h(1)=2))$. This “command” r satisfies both the safety monotonicity and the frame property, but not the contents independence; even though $\text{safe}(r, \square)$ and $1 \notin \text{dom}(\square)$, “command” r behaves differently depending on the contents of cell 1. The finite access property and contents independence are new in this paper, and they play an important role in the soundness of our data-refinement method (Sect. 5.1).

Definition 1 (Finite Local Action). *A finite local action, in short FLA, is an action that satisfies safety monotonicity, frame property, finite access property, and contents independence. A finite local action is av-free iff it does not relate any state to av.*

The set of finite local actions has a structure rich enough to interpret programs with all the low-level pointer operations that have been considered in separation logic.² Let \mathcal{F} be the poset of FLAs ordered by the “graph-subset” relation \sqsubseteq^3 , and let $\mathcal{F}_{\text{noav}}$ be the sub-poset of \mathcal{F} consisting of av-free FLAs. Particularly interesting are the low-level pointer operations, such as the memory update, allocation and deallocation of a cell, and a test “ $*l \in I$ ” for location l and integer set I : if l is allocated and it contains a value in I , the test skips; if l is allocated but its value is not in I , the test blocks; otherwise (i.e., if l is not allocated), the test generates the memory fault flt . For instance, $\text{test}(1, \{3\})$ expresses the conditional statement $\text{if}(*1 \neq 3)\{\text{diverge}\}$. Note that $\text{test}(1, \{3\})$ generates flt precisely when the boolean condition $*1 \neq 3$ dereferences an inactive cell.

Lemma 1. *The poset $\mathcal{F}_{\text{noav}}$ of av-free FLAs contains the operations in Fig. 2.*

Lemma 2. *Both \mathcal{F} and $\mathcal{F}_{\text{noav}}$ are complete lattices that have the set union as their join operator: for every family $\{r_i\}_{i \in I}$ in each poset, $\bigsqcup_{i \in I} r_i$ is $\bigcup_{i \in I} r_i$.*

² Thus, the set of finite local actions, as a semantic domain, expresses the computational behavior of pointer programs more accurately than the set of local actions, just as the set of continuous functions is a more “accurate” semantic domain than that of monotone functions in the domain theory.

³ $r \sqsubseteq r'$ iff $\forall h \in \text{St}. \forall v \in \text{St} \cup \{\text{flt}, \text{av}\}. h[r]v \Rightarrow h[r']v$

Let l be a location, i an integer, n a positive natural number, and I a set of integers.

$$\begin{aligned}
h[\text{update}(l, i)]v &\stackrel{\text{def}}{\iff} \text{if } l \notin \text{dom}(h) \text{ then } v=\text{flt} \text{ else } v=h[l \mapsto i] \\
h[\text{cons}(l, n)]v &\stackrel{\text{def}}{\iff} \text{if } l \notin \text{dom}(h) \text{ then } v=\text{flt} \text{ else } (\exists l'. v=(h[l \mapsto l'] \cdot [l' \mapsto 0, \dots, l'+n-1 \mapsto 0])) \\
h[\text{dispose}(l)]v &\stackrel{\text{def}}{\iff} \text{if } l \notin \text{dom}(h) \text{ then } v=\text{flt} \text{ else } v.[l \mapsto h(l)]=h \\
h[\text{test}(l, I)]v &\stackrel{\text{def}}{\iff} \text{if } l \notin \text{dom}(h) \text{ then } v=\text{flt} \text{ else } (v=h \wedge h(l) \in I)
\end{aligned}$$

Fig. 2. Semantic Low-level Pointer Operations

3 Programming Language

The programming language is Dijkstra’s language of guarded commands [5] extended with low-level pointer operations and module operations. The syntax of the language is given by the grammar:

$$C ::= f \mid a \mid C; C \mid C \parallel C \mid P \mid \text{fix } P.C$$

where f, a, P are, respectively, chosen from three disjoint sets $\text{mop}, \text{aop}, \text{pid}$ of identifiers. The first construct f is a module operation declared in the “interface specification” mop . Before a command in our language gets executed, it is first “linked” to a specific module that implements the interface mop . This linked module provides the meaning of the command f . The second construct a is an atomic operation, which a client can execute without using the module operations. Usually, a denotes a low-level pointer operation. Note that the language does not provide a syntax for building specific pointer operations. Instead, we assume that the interpretation $\llbracket - \rrbracket_a$ of these atomic client operations as av -free FLAs is given along with aop , and that under this interpretation, aop includes at least all the pointer operations in Lemma 1, so that aop includes all the atomic pointer operations considered in separation logic. The remaining four constructs of the language are the usual compound commands from Dijkstra’s language: sequential composition $C; C$, nondeterministic choice $C \parallel C$, the call of a parameterless procedure P , and the recursive definition $\text{fix } P.C$ of a parameterless procedure. As in Dijkstra’s language, the construct $\text{fix } P.C$ not only defines a parameterless recursive procedure P , but also calls the defined procedure. We express that a command C does not have free procedure names, by calling C a *complete command*.

We interpret commands using an instrumented denotational semantics; besides computing the usual state transformation, the semantics also checks whether each atomic client operation accesses the internals of a module, and for such illegal accesses, the semantics generates an access violation av .

To implement the instrumentation, we parameterize the semantics by what we call a *semantic module*. Let init and final be identifiers that are not in mop . A semantic module is a pair of a predicate p and a function η from $\text{mop} \cup \{\text{init}, \text{final}\}$ to $\mathcal{F}_{\text{noav}}$, such that (1) $\forall h, h'. (\text{safe}(\eta(\text{init}), h) \wedge h[\eta(\text{init})]h' \Rightarrow h' \in p * \text{true})$; (2) for all f in mop , $\forall h, h'. (\text{safe}(\eta(f), h) \wedge h \in p * \text{true} \wedge h[\eta(f)]h' \Rightarrow h' \in p * \text{true})$; (3) p is precise. Intuitively, the predicate p in the semantic module denotes the resource invariant for the module internals, and function η specifies the meaning of the module operations, initialization init and finalization final . The first condition of the semantic module requires initialization to establish the resource invariant,

$$\begin{array}{l}
\mu \in \mathcal{E} \stackrel{\text{def}}{=} \text{pid} \rightarrow \mathcal{F} \quad \llbracket C \rrbracket_{(p,\eta)} : \mathcal{E} \rightarrow \mathcal{F} \quad \llbracket C \rrbracket_{(p,\eta)}^c : \mathcal{F} \text{ (for complete } C) \\
\llbracket a \rrbracket_{(p,\eta)} \mu \stackrel{\text{def}}{=} \text{prot}(\llbracket a \rrbracket_a, p) \quad \llbracket C \rrbracket_{(p,\eta)} \mu \stackrel{\text{def}}{=} \llbracket C \rrbracket_{(p,\eta)} \mu \cup \llbracket C \rrbracket_{(p,\eta)} \mu \\
\llbracket f \rrbracket_{(p,\eta)} \mu \stackrel{\text{def}}{=} \eta(f) \quad \llbracket P \rrbracket_{(p,\eta)} \mu \stackrel{\text{def}}{=} \mu(P) \\
\llbracket C; C' \rrbracket_{(p,\eta)} \mu \stackrel{\text{def}}{=} \text{seq}(\llbracket C \rrbracket_{(p,\eta)} \mu, \llbracket C' \rrbracket_{(p,\eta)} \mu) \quad \llbracket \text{fix } P. C \rrbracket_{(p,\eta)} \mu \stackrel{\text{def}}{=} \text{fix } \lambda r. \llbracket C \rrbracket_{(p,\eta)} (\mu[P \rightarrow r]) \\
\llbracket C \rrbracket_{(p,\eta)}^c \stackrel{\text{def}}{=} \text{seq}(\text{seq}(\eta(\text{init}), \llbracket C \rrbracket_{(p,\eta)} \perp), \eta(\text{final})) \text{ (for complete } C)
\end{array}$$

where $\text{seq}: \mathcal{F} \times \mathcal{F} \rightarrow \mathcal{F}$ and $\text{prot}: \mathcal{F} \times \text{Pred} \rightarrow \mathcal{F}$ are defined as follows:

$$\begin{array}{l}
h[\text{prot}(r, p)]v \stackrel{\text{def}}{\iff} h[r]v \vee (v = \text{av} \wedge \neg h[r]\text{flt} \wedge \exists h_p, h_0. h = h_p \cdot h_0 \wedge h_p \in p \wedge h_0[r]\text{flt}) \\
h[\text{seq}(r, r')]v \stackrel{\text{def}}{\iff} (\exists h'. h[r]h' \wedge h'[r']v) \vee (h[r]\text{flt} \wedge v = \text{flt}) \vee (h[r]\text{av} \wedge v = \text{av})
\end{array}$$

Fig. 3. Semantics of Language

and the second condition, that the established resource invariant be preserved by module operations. The last condition is more subtle. It ensures that using the invariant p , we can determine which part of each state belongs to the module. Recall that a predicate q is precise iff every state h in $q * \text{true}$ has a unique splitting $h_q \cdot h_0 = h$ such that $h_q \in q$. Thus, if p is precise, then for every state h containing both the internals and externals of the module (i.e., $h \in p * \text{true}$), we can unambiguously split h into module-owned part h_p and client-owned part h_0 . This unambiguous splitting is used in the semantics to detect the access violation of the atomic client operations, and it also plays a crucial role in the soundness of our refinement method (Sect. 5.1). We remark that requiring the preciseness of the invariant p is not as restrictive as one might think, because most of the used resource invariants are precise; among the used resource invariants in separation logic, only one invariant is not precise, but even that invariant can safely be tightened to a precise one.⁴

Let \mathcal{E} be the poset of all functions from pid to \mathcal{F} ordered pointwise. Given semantic module (p, η) , we interpret a command as a continuous function $\llbracket - \rrbracket_{(p,\eta)}$ from \mathcal{E} to \mathcal{F} . For complete commands C , we consider an additional interpretation $\llbracket - \rrbracket_{(p,\eta)}^c$ that uses the least environment $\perp = \lambda P. \emptyset$, and runs the initialization and the finalization of the module (p, η) before and after $(\llbracket C \rrbracket_{(p,\eta)} \perp)$, respectively. The details of these two interpretations are shown in Fig. 3.

The most interesting part of the semantics lies in the interpretation of the atomic client operations. For each atomic operation a , its interpretation first looks up the original meaning $\llbracket a \rrbracket_a \in \mathcal{F}_{\text{noav}}$, which is given when the syntax of the language is defined. Then, the interpretation transforms the meaning into $\text{prot}(\llbracket a \rrbracket_a, p)$, “the p -protected execution of $\llbracket a \rrbracket_a$.” Intuitively, $\text{prot}(\llbracket a \rrbracket_a, p)$ behaves the same as $\llbracket a \rrbracket_a$, except that whenever $\llbracket a \rrbracket_a$ accesses the p -part of the input state, $\text{prot}(\llbracket a \rrbracket_a, p)$ generates av , thus indicating that there is an “access violation.” Since p is the resource invariant of the module, $\text{prot}(\llbracket a \rrbracket_a, p)$ notifies all illegal accesses to the module internals, by generating av .

Lemma 3. *The interpretation in Fig. 3 is well-defined.*

⁴ The only known unprecise invariant is $\text{listseg}(x, y)$ in [18], which means the existence of a (possibly cyclic) linked list segment from x to y . However, even that invariant can be made precise, if it is restricted to forbid a cycle in the list segment [15].

4 Data Refinement

The goal of this paper is to find a method for proving that a “concrete” module (q, ϵ) data-refines an “abstract” module (p, η) . In this section, we first formalize this goal by defining the notion of data refinement. Then, we demonstrate the difficulty of achieving the goal, by showing that the standard forward method is not sound in the presence of allocation-status testing.

We use the notion of data refinement that Mijajlović *et al.* devised in order to handle cross-boundary pointers. Usually, data refinement is a relation between modules defined by substitutability: a module (q, ϵ) data-refines another module (p, η) iff for all complete commands C using (p, η) , substituting the concrete module (q, ϵ) for the abstract module (p, η) improves the behavior of C , i.e., C becomes more deterministic with the concrete module. Mijajlović *et al.* weakened this usual notion of data refinement, by dropping the requirement about improvement for error-generating input states: if C with the abstract module (p, η) generates an access violation av or a memory fault flt from an input state h , then for this input h , the data refinement does not constrain the execution of C with the concrete module (q, ϵ) , and allows it to generate any outputs. In this paper, we use the following formalization of this weaker notion of data refinement:

Definition 2 (Data Refinement). *A module (q, ϵ) data-refines another module (p, η) iff for all complete commands C and all states h , if $\llbracket C \rrbracket_{(p, \eta)}^c$ does not generate an error from h (i.e., $\neg h[\llbracket C \rrbracket_{(p, \eta)}^c] \text{av} \wedge \neg h[\llbracket C \rrbracket_{(p, \eta)}^c] \text{flt}$), then*

$$(\neg h[\llbracket C \rrbracket_{(q, \epsilon)}^c] \text{av} \wedge \neg h[\llbracket C \rrbracket_{(q, \epsilon)}^c] \text{flt}) \wedge (\forall h'. h[\llbracket C \rrbracket_{(q, \epsilon)}^c] h' \Rightarrow h[\llbracket C \rrbracket_{(p, \eta)}^c] h').$$

The main benefit of considering this notion of data refinement is that a proof method for data refinement does not have to do anything special in order to handle the cross-boundary pointers. Recall that flt means that a command tries to dereference dangling pointers or nil , and av means that a command attempts to dereference the internal cells of a module without using module operations. Thus, if a command C does not generate an error from an input state h , then all the cells that C directly dereferences during execution must be allocated and belong to the “client” portion of the state; in particular, C does not dereference any cross-boundary pointers directly. Since the data refinement now asks for the improvement of only the error-free computations of C , a proof method for data refinement can ignore the “bad” computations where C dereferences cross-boundary pointers.

Unfortunately, even with this weaker notion of data refinement, standard proof methods for data refinement are not sound; they fail to deal with the allocation-status testing. We explain this soundness problem making use of the notion of the forward simulation in [12]. As pointed out in their work, while successfully dealing with the cross-boundary pointer dereferencing problem, the forward method is not sound for allocation-status testing.

The key concept of the forward simulation in [12] is an operator fsim that maps a pair (R_0, R_1) of state relations to a relation $\text{fsim}(R_0, R_1)$ between FLAs.

Intuitively, $r'[\text{fsim}(R_0, R_1)]r$ means that given R_0 -related input states h' and h , if r does not generate an error from h , then (1) r' does not generate an error from h' and (2) every output of r' from h' is R_1 -related to some outcome of r . More precisely, $r'[\text{fsim}(R_0, R_1)]r$ iff for all states h' and h , if $(h'[R_0]h \wedge \neg h[r]\text{flt} \wedge \neg h[r]\text{av})$, then

$$(\neg h'[r']\text{flt} \wedge \neg h'[r']\text{av}) \wedge (\forall h'_1. h'[r']h'_1 \Rightarrow \exists h_1. h[r]h_1 \wedge h'_1[R_1]h_1).$$

The condition about the absence of errors comes from the fact that the data refinement considers only error-free computations. Except this condition, the way of relating two actions (or commands) in $\text{fsim}(R_0, R_1)$ is fairly standard in the work on data refinement [6, 4].

Let Δ be the diagonal relation on states⁵, and for state relations R_0 and R_1 , let $R_0 * R_1$ be their relational separating conjunction [17]: $h'[R_0 * R_1]h$ iff h' and h are, respectively, split into $h'_0 \cdot h'_1 = h'$ and $h_0 \cdot h_1 = h$ such that the first parts h'_0, h_0 are related by R_0 and the second parts h'_1, h_1 by R_1 .⁶ The formal definition of forward simulation is given below:

Definition 3 (Forward Simulation). *Let $(q, \epsilon), (p, \eta)$ be semantic modules, and R a relation s.t. $R \subseteq q \times p$. Module (q, ϵ) forward-simulates (p, η) by R iff*

1. $\epsilon(\text{init})[\text{fsim}(\Delta, R * \Delta)]\eta(\text{init})$ and $\epsilon(\text{final})[\text{fsim}(R * \Delta, \Delta)]\eta(\text{final})$;
2. $\forall f \in \text{mop}. \epsilon(f)[\text{fsim}(R * \Delta, R * \Delta)]\eta(f)$.

The relation $R * \Delta$ here expresses that the corresponding states of r' and r can, respectively, be partitioned into the module and client parts; the module parts of r' and r are related by R , but the client parts of r' and r are the same.

The forward simulation is not sound: there are modules $(q, \epsilon), (p, \eta)$ such that the concrete module (q, ϵ) forward-simulates the abstract module (p, η) by some $R \subseteq q \times p$, but it does not data-refine it. The main reason of this unsoundness is that the low-level pointer operations in our language, especially those implementing allocation-status testing, break the underlying assumption of the forward simulation. The forward simulation assumes a language where if a command C does not call module operations, then for all relations $R \subseteq q \times p$, the command “forward-simulates” itself by R : $\llbracket C \rrbracket_{(q, \epsilon)} \mu' [\text{fsim}(R * \Delta, R * \Delta)] \llbracket C \rrbracket_{(p, \eta)} \mu$ for all μ', μ that define $\text{fsim}(R * \Delta, R * \Delta)$ -related “procedures”. Our language, however, does not satisfy this assumption; if an atomic client command a implements the allocation-status testing, it is not related to itself by $\text{fsim}(R * \Delta, R * \Delta)$ in general. For instance, having a concrete module (q, ϵ) and the abstract one (p, η) and a relation R between them, consider an atomic command $\text{cons}(2, 1)$ that allocates one new cell initialized to 0 and assigns its address to cell 2; in case that cell 2 is not allocated initially, $\text{cons}(2, 1)$ generates flt .⁷ Let R be defined by $h'_0[R]h_0 \Leftrightarrow h'_0 = [] \wedge h_0 = [1 \rightarrow 2]$. Then, $h'[R * \Delta]h$ iff there is some state h_1 such that $1 \notin \text{dom}(h_1) \wedge h' = [] \cdot h_1 \wedge h = [1 \rightarrow 2] \cdot h_1$. Thus, states $h' = [2 \rightarrow 0]$ and $h = [1 \rightarrow 2, 2 \rightarrow 0]$ are $R * \Delta$ -related. We will now consider the execution of

⁵ $h'[\Delta]h \stackrel{\text{def}}{\Leftrightarrow} h' = h$

⁶ $h'[R_0 * R_1]h \stackrel{\text{def}}{\Leftrightarrow} \exists h'_0, h'_1, h_0, h_1. h'_0 \cdot h'_1 = h' \wedge h_0 \cdot h_1 = h \wedge h'_0[R_0]h_0 \wedge h'_1[R_1]h_1$.

⁷ $h[\llbracket \text{cons}(2, 1) \rrbracket_a]v \stackrel{\text{def}}{\Leftrightarrow}$ if $2 \notin \text{dom}(h)$ then $v = \text{flt}$ else $\exists n. v = h[2 \rightarrow n] \cdot [n \rightarrow 0]$.

$\text{cons}(2, 1)$ from these $R*\Delta$ -related states h' and h . When $\text{cons}(2, 1)$ is run from h' (with the concrete module (q, ϵ)), it can allocate cell 1 and give the output state $h'_1 = [1 \rightarrow 0, 2 \rightarrow 1]$ (i.e., $h_1[\llbracket \text{cons}(2, 1) \rrbracket_{(q, \epsilon)} \mu'] h'_1$), because cell 1 is free initially (i.e., $1 \notin \text{dom}(h')$). However, when the same command is run from h (with the abstract module (p, η)), it cannot allocate cell 1, because 1 is already active in h (i.e., $1 \in \text{dom}(h)$). In this case, all the output states of $\text{cons}(2, 1)$ have the form $[1 \rightarrow 0, 2 \rightarrow n, n \rightarrow 0]$ for some $n \in \text{Nats} - \{1, 2\}$. Note that the state $h'_1 = [1 \rightarrow 0, 2 \rightarrow 1]$ is not $R*\Delta$ -related to any such outputs $[1 \rightarrow 0, 2 \rightarrow n, n \rightarrow 0]$. Thus, we cannot have that $(\llbracket \text{cons}(2, 1) \rrbracket_{(q, \epsilon)} \mu') [\text{fsim}(R*\Delta, R*\Delta)] (\llbracket \text{cons}(2, 1) \rrbracket_{(p, \eta)} \mu)$. In the full version of the paper [13], we have used these R and cons to construct a counter example for the soundness of the forward simulation.

5 Power Simulation

We now present the main result of this paper: a new method for data refinement, called power simulation, and its soundness proof.

The key idea of power simulation is to use the state-set lifting $\text{lft}(r)$ of a FLA:

$$\text{lft}(r) : \wp(\text{St}) \leftrightarrow (\wp(\text{St}) \cup \{\text{flt}, \text{av}\})$$

$$H[\text{lft}(r)]V \stackrel{\text{def}}{\iff} (V \subseteq \text{St} \wedge \forall h' \in V. \exists h \in H. h[r]h') \vee ((V = \text{av} \vee V = \text{flt}) \wedge \exists h \in H. h[r]V).$$

Given an input state set H , the “lifted command” $\text{lft}(r)$ runs r for all the states in H , chooses some states among the results, and returns the set V of the chosen states. Note that V might not contain some possible outputs from H ; so, $\text{lft}(r)$ is different from the usual direct image map of r , and in general, it is a relation rather than a function. For each module (p, η) , we write $\text{lft}(\eta)$ for the lifting of all module operations (i.e., $\forall f \in \text{mop}. \text{lft}(\eta)(f) = \text{lft}(\eta(f))$), and call $(p, \text{lft}(\eta))$ the *lifting* of (p, η) .

The power simulation is the usual forward simulation of a *lifted* “abstract” module by a *normal* “concrete” module. Suppose that we want to show that a concrete module (q, ϵ) data-refines an abstract module (p, η) . Define a power relation to be a relation between states and state sets. Intuitively, the power simulation says that to prove this data refinement, we only need to find a “good” power relation $\mathcal{R} \subseteq \text{St} \times \wp(\text{St})$ such that every concrete-module operation $\epsilon(k)$ “forward-simulates” the corresponding lifted abstract-module operation $\text{lft}(\eta(k))$ by \mathcal{R} . The official definition of power simulation formalizes this intuition by specifying (1) which power relation should be considered good for given modules (q, ϵ) and (p, η) , and (2) what it means that a normal command “forward-simulates” a lifted command. For the first, we use the *expansion* operator and *admissibility* condition for power relations. For the second, we use the operator psim that maps a power-relation pair to a relation on FLAs. We will now define these sub-components of power simulation, and use them to give the formal definition of power simulation.

We explain operator psim first. For power relations \mathcal{R}_0 and \mathcal{R}_1 , $\text{psim}(\mathcal{R}_0, \mathcal{R}_1)$ relates a “concrete” FLA r' with an “abstract” r iff for every \mathcal{R}_0 -related input state h' and state set H , if $\text{lft}(r)$ does not generate an error from H , then all

the outputs of r' from h' are \mathcal{R}_1 -related to some output state sets of $\text{lft}(r)$ from H . More precisely, $r'[\text{psim}(\mathcal{R}_0, \mathcal{R}_1)]r$ iff for all h' and H , if $h'[\mathcal{R}_0]H$ and neither $H[\text{lft}(r)]\text{ft}$ nor $H[\text{lft}(r)]\text{av}$, then

$$(\neg h'[r']\text{ft} \wedge \neg h'[r']\text{av}) \wedge (\forall h'_1. h'[r']h'_1 \Rightarrow \exists H_1. H[\text{lft}(r)]H_1 \wedge h'_1[\mathcal{R}_1]H_1).$$

Note that this definition is the lifted version of fsim in Sec. 4; except that it considers the lifted computation $\text{lft}(r)$, instead of the usual computation r , it coincides with the definition of fsim . In the definition of power simulation, we will use this psim to express the “forward-simulation” of a lifted command by a normal command.

Next, we define the expansion operator $-\otimes\Delta$ for power relations. The expansion $\mathcal{R}\otimes\Delta$ of a power relation \mathcal{R} is a power relation defined as follows:

$$h[\mathcal{R}\otimes\Delta]H \stackrel{\text{def}}{\Leftrightarrow} \exists h_r, h_0, H_r. (h = h_r \cdot h_0 \wedge h_r[\mathcal{R}]H_r \wedge H = H_r * \{h_0\}).$$

Intuitively, the definition means that h and H are obtained by extending \mathcal{R} -related state h_r and state sets H_r by the same state h_0 . Usually, \mathcal{R} is a “coupling” power relation that connects the internals of two modules, and $\mathcal{R}\otimes\Delta$ expands this coupling relation to the relation for the entire memory, by asking that the added client parts must be identical.

The final subcomponent of power simulation is the admissibility condition for power relations. A power relation \mathcal{R} is *admissible* iff for every \mathcal{R} -related state h and state set H (i.e., $h[\mathcal{R}]H$), we have that⁸

$$H \neq \emptyset \wedge (\forall L \subseteq_{\text{fin}} \text{Loc} - \text{dom}(h). \exists H_1 \subseteq H. (H_1 \neq \emptyset \wedge h[\mathcal{R}]H_1 \wedge \forall h_1 \in H_1. h_1 \# L)).$$

The first conjunct in the admissibility condition means that all related state sets must contain at least one state. The second conjunct is about the “free cells” in these related state sets. It means that if $h[\mathcal{R}]H$, state set H collectively has at least as many free cells as h : for every finite collection L of free cells in h , set H contains states that do not have any of the cells in L , and, moreover, the set H_1 of such states itself collectively has as many free cells as h . To understand the second conjunct more clearly, consider power relations $\mathcal{R}_0, \mathcal{R}_1, \mathcal{R}_2$ defined as follows:

$$\begin{aligned} h[\mathcal{R}_0]H &\stackrel{\text{def}}{\Leftrightarrow} h=[3\rightarrow 1] \wedge H=\{[3\rightarrow 5]\} & h[\mathcal{R}_1]H &\stackrel{\text{def}}{\Leftrightarrow} h=[3\rightarrow 1] \wedge H=\{[3\rightarrow 5, 4\rightarrow 5]\} \\ h[\mathcal{R}_2]H &\stackrel{\text{def}}{\Leftrightarrow} h=[3\rightarrow 1] \wedge \exists L \subseteq_{\text{fin}} \text{Loc}. H=\{[3\rightarrow 5, n\rightarrow 5] \mid n \notin L \cup \{3\}\} \end{aligned}$$

The first power relation \mathcal{R}_0 is admissible, because set $\{[3\rightarrow 5]\}$ has only one state $[3\rightarrow 5]$ that has the exactly same free cells, namely all cells other than 3, as state $[3\rightarrow 1]$. On the other hand, \mathcal{R}_1 is not admissible, because the (unique) state in $\{[3\rightarrow 5, 4\rightarrow 5]\}$ has an active cell 4 that is not free in $[3\rightarrow 1]$. The last relation \mathcal{R}_2 is tricky; relation \mathcal{R}_2 is admissible, even though for all \mathcal{R}_2 -related h and H , every state in H has more active cells than h . The intuitive reason for this is that for every free cell in $[3\rightarrow 1]$, set H contains a state that does not contain the cell, and so, it collectively has as many free cells as $[3\rightarrow 1]$; in a sense, by having sufficiently many states, H hides the identity of the additional cell n . The formal proof that \mathcal{R}_2 satisfies the second conjunct of the admissibility condition proceeds as

⁸ Recall that $h_1 \# L$ iff $\text{dom}(h_1) \cap L = \emptyset$.

follows. Consider H, h', L_1 such that $h'[\mathcal{R}_2]H$ and $L_1 \subseteq_{fin} (\text{Loc} - \text{dom}(h))$. By the definition of \mathcal{R}_2 , there exists a finite location set L such that $H = \{[3 \rightarrow 5, n \rightarrow 5] \mid n \notin L \cup \{3\}\}$. Let $H_1 = \{[3 \rightarrow 5, n \rightarrow 5] \mid n \notin L \cup L_1 \cup \{3\}\}$. The defined set H_1 is a nonempty subset of H . We now prove that H_1 is in fact the required subset of H in the admissibility condition. Since $h'[\mathcal{R}_2]H_1$, $h'[\mathcal{R}_2]H_1$ follows from the definition of \mathcal{R}_2 and H_1 . We also have that $\forall h_1 \in H_1. \text{dom}(h_1) \cap L_1 = \emptyset$, because $\text{dom}(h_1) \cap L_1 \subseteq \{3\}$ but L_1 does not contain 3 ($h' = [3 \rightarrow 1] \# L_1$).

Using the expansion operator and admissibility condition, we can define the criteria for deciding which power relation should be considered “good” for given modules (q, ϵ) and (p, η) . The criteria is: a power relation should be the expansion $\mathcal{R} \otimes \Delta$ of an admissible \mathcal{R} for the module internals (i.e., $\mathcal{R} \subseteq q \times \wp(p)$). The following lemma, which we will prove later in Sec. 5.1, provides the justification of this criteria:

LEMMA 4: For all q, p , and all power relations $\mathcal{R} \subseteq q \times \wp(p)$, if \mathcal{R} is admissible and q is precise, then $\forall r \in \mathcal{F}_{\text{nov}}. \text{prot}(r, q)[\text{psim}(\mathcal{R} \otimes \Delta, \mathcal{R} \otimes \Delta)]\text{prot}(r, p)$.

To see the significance of this lemma, recall that the forward simulation in Sec. 4 failed to be sound mainly because some atomic client operations are not related to themselves by fsim . The lemma indicates that as long as we are using admissible power relation \mathcal{R} , we do not have such a problem for psim : if \mathcal{R} is admissible, then for all atomic client operations a and all environment pairs (μ', μ) with $\text{psim}(\mathcal{R} \otimes \Delta, \mathcal{R} \otimes \Delta)$ -related procedures, we have that $\llbracket a \rrbracket_{(q, \epsilon)} \mu' [\text{psim}(\mathcal{R} \otimes \Delta, \mathcal{R} \otimes \Delta)] \llbracket a \rrbracket_{(p, \eta)} \mu$.

We now define the power simulation of an abstract module (p, η) by a concrete module (q, ϵ) . Let \mathcal{R} be an admissible power relation such that $\mathcal{R} \subseteq q \times \wp(p)$, and let ID be the “identity” power relation defined by: $h[\text{ID}]H \stackrel{\text{def}}{=} \{h\} = H$.

Definition 4 (Power Simulation). Module (q, ϵ) power-simulates (p, η) by \mathcal{R} iff

1. $\epsilon(\text{init})[\text{psim}(\text{ID}, \mathcal{R} \otimes \Delta)]\eta(\text{init})$ and $\epsilon(\text{final})[\text{psim}(\mathcal{R} \otimes \Delta, \text{ID})]\eta(\text{final})$;
2. $\forall f \in \text{mop}. \epsilon(f)[\text{psim}(\mathcal{R} \otimes \Delta, \mathcal{R} \otimes \Delta)]\eta(f)$.

Example 1. We demonstrate power simulation using the semantic modules (q, ϵ) and (p, η) that, respectively, correspond to `counter2` and `counter3` in Fig. 1. Recall that both `counter2` and `counter3` implement a counter “object” with two operations, `inc` for incrementing the counter and `read` for reading the value of the counter; the main difference is that `counter3` uses two cells, namely cell 1 and a newly allocated one, to track the value of the counter, while `counter2` uses only cell 1 for the same purpose. The corresponding semantic modules, (q, ϵ) for `counter2` and (p, η) for `counter3`, are defined in Fig. 4. Note that the resource invariant p indicates that `counter3` uses two cells 1 and n internally, and the invariant q that `counter2` uses only one cell 1 internally. We will now show that the space saving in `counter2` is correct, by proving that (q, ϵ) power-simulates (p, η) .

The first step of power simulation is to find an admissible power relation that couples the internals of (q, ϵ) and (p, η) . For this, we use the following \mathcal{R} :

$$\begin{aligned}
h \in p &\stackrel{\text{def}}{\iff} \exists n, n'. n' \neq 1 \wedge n \geq 0 \wedge h = [1 \rightarrow n', n' \rightarrow n] \\
h[\eta(\text{init})]v &\stackrel{\text{def}}{\iff} \text{if } (1 \notin \text{dom}(h)) \text{ then } v = \text{flt} \text{ else } \exists n. n \notin \text{dom}(h) \wedge v = h[1 \rightarrow n] \cdot [n \rightarrow 0] \\
h[\eta(\text{inc})]v &\stackrel{\text{def}}{\iff} \text{if } (1 \notin \text{dom}(h) \vee h(1) \notin \text{dom}(h)) \text{ then } v = \text{flt} \text{ else } v = h[h(1) \rightarrow (h(h(1)) + 1)] \\
h[\eta(\text{read})]v &\stackrel{\text{def}}{\iff} \text{if } (1 \notin \text{dom}(h) \vee h(1) \notin \text{dom}(h) \vee 3 \notin \text{dom}(h)) \text{ then } v = \text{flt} \text{ else } v = h[3 \rightarrow h(h(1))] \\
h[\eta(\text{final})]v &\stackrel{\text{def}}{\iff} \text{if } (1 \notin \text{dom}(h) \vee h(1) \notin \text{dom}(h)) \text{ then } v = \text{flt} \\
&\quad \text{else } \exists h_0. v = h_0[1 \rightarrow 0] \wedge h = h_0 \cdot [h(1) \rightarrow h(h(1))] \\
h \in q &\stackrel{\text{def}}{\iff} \exists n. n \geq 0 \wedge h = [1 \rightarrow n] \\
h[\epsilon(\text{init})]v &\stackrel{\text{def}}{\iff} \text{if } (1 \notin \text{dom}(h)) \text{ then } v = \text{flt} \text{ else } v = h[1 \rightarrow 0] \\
h[\epsilon(\text{inc})]v &\stackrel{\text{def}}{\iff} \text{if } (1 \notin \text{dom}(h)) \text{ then } v = \text{flt} \text{ else } v = h[1 \rightarrow (h(1) + 1)] \\
h[\epsilon(\text{read})]v &\stackrel{\text{def}}{\iff} \text{if } (1 \notin \text{dom}(h) \vee 3 \notin \text{dom}(h)) \text{ then } v = \text{flt} \text{ else } v = h[3 \rightarrow h(1)] \\
h[\epsilon(\text{final})]v &\stackrel{\text{def}}{\iff} \text{if } (1 \notin \text{dom}(h)) \text{ then } v = \text{flt} \text{ else } v = h[1 \rightarrow 0]
\end{aligned}$$

Fig. 4. Definition of Module (p, η) and (q, ϵ)

$$h[\mathcal{R}]H \stackrel{\text{def}}{\iff} \exists L, n. L \subseteq_{\text{fin}} \text{Loc} \wedge n \geq 0 \wedge h = [1 \rightarrow n] \wedge H = \{[1 \rightarrow n', n' \rightarrow n] \mid n' \notin L \cup \{1\}\}.$$

Intuitively, $h[\mathcal{R}]H$ means that all the states in H and state h represent the same counter having the value $h(1)$, and moreover, H collectively has as many free cells as h .

The next step is to show that all the corresponding module operations of (q, ϵ) and (p, η) are related by psim . Here we only show that $\epsilon(\text{init})$ and $\eta(\text{init})$ are $\text{psim}(\text{ID}, \mathcal{R} \otimes \Delta)$ -related. Consider h and H related by the “identity relation” ID . Then, by the definition of ID , set H must be the singleton set containing the heap h . Thus, it suffices to show that if $\text{lft}(\eta(\text{init}))$ does not generate an error from $\{h\}$, all the outputs of $\epsilon(\text{init})$ from h are $\mathcal{R} \otimes \Delta$ -related to some output state sets of $\text{lft}(\eta(\text{init}))$ from $\{h\}$. Suppose that $\text{lft}(\eta(\text{init}))$ does not generate an error from $\{h\}$. Then, $\eta(\text{init})$ cannot output flt from h , and so, cell 1 should be in $\text{dom}(h)$. From this, it follows that the concrete initialization $\epsilon(\text{init})$ does not generate an error from h . We now check the non-error outputs of $\epsilon(\text{init})$. When started from h , the concrete initialization $\epsilon(\text{init})$ has only one non-error output, namely state $h[1 \rightarrow 0]$. We split this output state $h[1 \rightarrow 0]$ into $[1 \rightarrow 0]$ and the remainder h_0 . By the definition of \mathcal{R} , the first part $[1 \rightarrow 0]$ of the splitting is \mathcal{R} -related to $H_r = \{[1 \rightarrow n', n' \rightarrow 0] \mid n' \notin \text{dom}(h)\}$. Thus, extending $[1 \rightarrow 0]$ and H_r by the remainder h_0 gives $\mathcal{R} \otimes \Delta$ -related state $[1 \rightarrow 0] \cdot h_0 = h[1 \rightarrow 0]$ and state set $H_r * \{h_0\}$. The state set $H_r * \{h_0\}$ is equal to $\{h[1 \rightarrow n'] \cdot [n' \rightarrow 0] \mid n' \notin \text{dom}(h)\}$, and so, it is a possible output of $\text{lft}(\eta(\text{init}))$ from $\{h\}$ by the definition of $\text{lft}(\eta(\text{init}))$. We have just shown that the output $h[1 \rightarrow 0]$ is $\mathcal{R} \otimes \Delta$ -related to some output of $\text{lft}(\eta(\text{init}))$, as required.

The nondeterministic allocation in the abstract initialization $\eta(\text{init})$ is crucial for the correctness of data refinement. Suppose that we change the initialization of the abstract module such that it allocates a specific cell 2:

$$h[\eta(\text{init})]v \stackrel{\text{def}}{\iff} \text{if } (1 \notin \text{dom}(h)) \text{ then } (v = \text{flt}) \text{ else } (2 \notin \text{dom}(h) \wedge v = h[1 \rightarrow 2] \cdot [2 \rightarrow 0])$$

Then, (q, ϵ) no longer data-refines (p, η) ;⁹ by testing the allocation status of cell 2 using memory allocation and pointer comparison, a client command can detect

⁹ Even when we replace p by a more precise invariant $\{[1 \rightarrow 2, 2 \rightarrow n] \mid n \geq 0\}$, module (q, ϵ) does not data-refine (p, η) .

the replacement of (p, η) by (q, ϵ) , and exhibit a behavior that is only possible with (q, ϵ) , but not with (p, η) . Power simulation correctly captures this failure of data refinement. More specifically, for all power relations $\mathcal{R} \subseteq q \times \wp(p)$ if $\epsilon(\text{init})[\text{psim}(\text{ID}, \mathcal{R} \otimes \Delta)]\eta(\text{init})$, then \mathcal{R} cannot be admissible. To see the reason, suppose that $\epsilon(\text{init})[\text{psim}(\text{ID}, \mathcal{R} \otimes \Delta)]\eta(\text{init})$. When $\epsilon(\text{init})$ and $\text{lft}(\eta(\text{init}))$ are run from ID-related $[1 \rightarrow 0]$ and $\{[1 \rightarrow 0]\}$, $\epsilon(\text{init})$ outputs $[1 \rightarrow 0]$ and $\text{lft}(\eta(\text{init}))$ outputs $\{[1 \rightarrow 2, 2 \rightarrow 0]\}$ or \emptyset . Thus, by the definition of $\text{psim}(\text{ID}, \mathcal{R} \otimes \Delta)$, $[1 \rightarrow 0]$ should be $\mathcal{R} \otimes \Delta$ -related to $\{[1 \rightarrow 2, 2 \rightarrow 0]\}$ or \emptyset . Then, by the definition of $\mathcal{R} \otimes \Delta$, state $[1 \rightarrow 0]$ is \mathcal{R} -related to $\{[1 \rightarrow 2, 2 \rightarrow 0]\}$ or \emptyset . In either case, \mathcal{R} is not admissible; the first case violates the second conjunct about the free cells in the admissibility condition, and the second case violates the first conjunct about the nonemptiness. \square

5.1 Soundness of Power Simulation

The soundness of power simulation follows from the fact that every atomic client operation is related to itself by psim (Lemma 4), all language constructs preserve psim (Lemma 5,6) and $\text{psim}(\text{ID}, \text{ID})$ is precisely the improvement requirement in the definition of data refinement (Lemma 7). Among these lemmas, we give the proof of only the most important one, Lemma 4. Then, we show how all the lemmas are used to give the soundness of power simulation. The missing proofs appear in the full version of this paper [13].

Lemma 4. *For all predicates q, p , and all power relations $\mathcal{R} \subseteq q \times \wp(p)$, if q is precise and \mathcal{R} is admissible, then $\forall r \in \mathcal{F}_{\text{noav}}. \text{prot}(r, q)[\text{psim}(\mathcal{R} \otimes \Delta, \mathcal{R} \otimes \Delta)]\text{prot}(r, p)$.*

Note that the lemma requires that the “resource invariant” q for the “concrete module” be precise, and that r be a av -free finite local action. None of these requirements can be omitted, because the requirements are used crucially in the proof of the lemma.

Proof. Let r_q be $\text{prot}(r, q)$ and let r_p be $\text{prot}(r, p)$. Pick arbitrary $[\mathcal{R} \otimes \Delta]$ -related h and H such that $\text{lft}(r_p)$ does not generate an error from H . Since $h[\mathcal{R} \otimes \Delta]H$, state h and state set H can, respectively, be split into $h_q \cdot h_0 = h$ and $H = H_p * \{h_0\}$ for some h_q, h_0, H_p such that $h_q[\mathcal{R}]H_p$. We note two facts about these splittings. First, set H_p contains a state that is disjoint from h_0 . Since h_q and H_p are related by the admissible relation \mathcal{R} and $\text{dom}(h_q)$ is disjoint from $\text{dom}(h_0)$, there is a nonempty subset of H_p such that every h_1 in the subset satisfies $h_1 \# \text{dom}(h_0)$. We pick one state from this subset, and call it h_p . Second, the state h_p in H_p and the part h_q of the splitting of h , respectively, belong to p and q . This second fact follows since $h_q[\mathcal{R}]H_p$ and $\mathcal{R} \subseteq q \times \wp(p)$. We sum up the obtained properties about H_p, h_0, h_q, h_p below:

$$H = H_p * \{h_0\} \wedge h = h_q \cdot h_0 \wedge h_q[\mathcal{R}]H_p \wedge h_p \# h_0 \wedge h_p \in p \wedge h_q \in q.$$

We now prove that r_q does not generate an error from h . Since the lifted command $\text{lft}(r_p)$ does not generate an error from H and state $h_p \cdot h_0$ is in this input state set H , we have that $\neg h_p \cdot h_0[r_p]\text{flt} \wedge \neg h_p \cdot h_0[r_p]\text{av}$. This absence of errors of r_p ensures one important property of r : r cannot generate flt from h_0 .

To see the reason, note that h_p is in p , and that $\neg h_p \cdot h_0[r]\text{flt}$ since $\neg h_p \cdot h_0[r_p]\text{flt}$. So, if $h_0[r]\text{flt}$, then by the definition of **prot**, we have that $h_p \cdot h_0[r_p]\text{av}$, which contradicts $\neg h_p \cdot h_0[r_p]\text{av}$. We will use this property of r to show $\neg h[r_q]\text{flt}$ and $\neg h[r_q]\text{av}$. Since $h = h_0 \cdot h_q$ and $\neg h_0[r]\text{flt}$, by the safety monotonicity of r , we have that $\neg h[r]\text{flt}$. Thus, $\neg h[r_q]\text{flt}$ by the definition of **prot**. For $\neg h[r_q]\text{av}$, we have to show that

$$\neg h[r]\text{av} \wedge (h[r]\text{flt} \vee (\forall m_q, m_0 \in \text{St}. (m_q \cdot m_0 = h \wedge m_q \in q) \Rightarrow \neg m_0[r]\text{flt})).$$

Since r is **av**-free, it does not output **av** for any input states. For the second conjunct, consider a splitting $m_q \cdot m_0$ of h such that $m_q \in q$. Then, since $h = h_q \cdot h_0$, $h_q \in q$ and q is precise, we should have that $m_q = h_q$ and $m_0 = h_0$. Since $\neg h_0[r]\text{flt}$, it follows that $\neg m_0[r]\text{flt}$.

Finally, we prove that every output state of r_q from h is $\mathcal{R} \otimes \Delta$ -related to some output state set of $\text{lft}(r_p)$ from H . In the proof, we will use $\neg h_0[r]\text{flt}$, which we have shown in the previous paragraph. Consider a state h' such that $h[r_q]h'$. Since $h = h_0 \cdot h_q$, by the definition of **prot**(r, q), we have that $h_0 \cdot h_q[r]h'$. Since $\neg h_0[r]\text{flt}$, we can apply the frame property of r to this computation, and obtain a substate h'_0 of h' such that $h' = h'_0 \cdot h_q$. Let L_0 be the finite set that includes all the indirectly accessed locations by the “computation” $h_0 \cdot h_q[r]h'_0 \cdot h_q$; L_0 is guaranteed to exist by the finite access property of r . Let L be the location set $(L_0 \cup \text{dom}(h_0) \cup \text{dom}(h'_0)) - \text{dom}(h_q)$. Since $h_q[\mathcal{R}]H_p$ and \mathcal{R} is admissible, there is a subset H_1 of H_p such that

$$H_1 \subseteq H_p \wedge H_1[\mathcal{R}]h_q \wedge \forall h_1 \in H_1. h_1 \# L.$$

We will show that $H_1 * \{h'_0\}$ is the required output state set. Since h_q and H_1 are \mathcal{R} -related, their h'_0 -extensions, $h_q \cdot h'_0$ and $H_1 * \{h'_0\}$, have to be $\mathcal{R} \otimes \Delta$ -related. Thus, it remains to show that $H = H_p * \{h_0\}[\text{lft}(r_p)]H_1 * \{h'_0\}$. Instead of proving this relationship directly, we will prove that

$$H_1 * \{h_0\}[\text{lft}(r_p)]H_1 * \{h'_0\}.$$

Because, then, the definition of $\text{lft}(r_p)$ will ensure that we also have the required computation. For every m in $H_1 * \{h'_0\}$, there is a state $m_1 \in H_1$ such that $m = m_1 \cdot h'_0$. By the choice of H_1 , we have $m_1 \in H_p \wedge m_1 \# L$. Then, there exist splitting $n_1 \cdot n_2 = m_1$ of m_1 and splitting $o_2 \cdot o_3 = h_q$ of h_q with the property that $n_1 \# h_q$ and $\text{dom}(n_2) = \text{dom}(o_2)$. Note that $n_1 \# h_q$ implies $n_1 \# L_0$, because $L_0 \subseteq L \cup \text{dom}(h_q)$ and $n_1 \cdot n_2 \# L$. We obtain a new computation of r_p as follows:

$$\begin{aligned} h_q \cdot h_0[r]h_q \cdot h'_0 &\Longrightarrow n_1 \cdot h_q \cdot h_0[r]n_1 \cdot h_q \cdot h'_0 && (\because \text{the finite access property of } r) \\ &\Longrightarrow n_1 \cdot o_2 \cdot o_3 \cdot h_0[r]n_1 \cdot o_2 \cdot o_3 \cdot h'_0 && (\because h_q = o_2 \cdot o_3) \\ &\Longrightarrow n_1 \cdot n_2 \cdot o_3 \cdot h_0[r]n_1 \cdot n_2 \cdot o_3 \cdot h'_0 && (\because \text{the contents independence of } r) \\ &\Longrightarrow n_1 \cdot n_2 \cdot h_0[r]n_1 \cdot n_2 \cdot h'_0 && (\because \text{the frame property of } r) \\ &\Longrightarrow m_1 \cdot h_0[r]m && (\because m_1 = n_1 \cdot n_2 \wedge m_1 \cdot h'_0 = m) \\ &\Longrightarrow m_1 \cdot h_0[r_p]m && (\because \text{the definition of } \text{prot}(r, p)) \end{aligned}$$

Note that the input $m_1 \cdot h_0$ of the obtained computation belongs to the state set $H_1 * \{h_0\}$. We just have shown $H_1 * \{h_0\}[\text{lft}(r_p)]H_1 * \{h'_0\}$. \square

Lemma 5. For all power relations $\mathcal{R}_0, \mathcal{R}_1, \mathcal{R}_2$ and all FLAs r_0, r'_0, r_1, r'_1 , if $r'_0[\text{psim}(\mathcal{R}_0, \mathcal{R}_1)]r_0 \wedge r'_1[\text{psim}(\mathcal{R}_1, \mathcal{R}_2)]r_1$, then $\text{seq}(r'_0, r'_1)[\text{psim}(\mathcal{R}_0, \mathcal{R}_2)]\text{seq}(r_0, r_1)$.

Lemma 6. For all power relations $\mathcal{R}_0, \mathcal{R}_1$, sets I and I -indexed families $\{r'_i\}_{i \in I}$, $\{r_i\}_{i \in I}$ of FLAs, if $\forall i \in I. r'_i[\text{psim}(\mathcal{R}_0, \mathcal{R}_1)]r_i$, then $\bigcup_{i \in I} r'_i[\text{psim}(\mathcal{R}_0, \mathcal{R}_1)] \bigcup_{i \in I} r_i$.

Theorem 1 (Abstraction). Let $(q, \epsilon), (p, \eta)$ be semantic modules, and \mathcal{R} be an admissible power relation s.t. $\mathcal{R} \subseteq q \times \wp(p)$. If (q, ϵ) power-simulates (p, η) by \mathcal{R} , then for all commands C and all environments μ, μ' , we have that

$$(\forall P. \mu'(P)[\text{psim}(\mathcal{R} \otimes \Delta, \mathcal{R} \otimes \Delta)]\mu(P)) \Rightarrow \llbracket C \rrbracket_{(q, \epsilon)} \mu'[\text{psim}(\mathcal{R} \otimes \Delta, \mathcal{R} \otimes \Delta)] \llbracket C \rrbracket_{(p, \eta)} \mu.$$

Proof. We use induction on the structure of C . When C is a module operation f or a procedure name P , the theorem follows from the assumption: (q, ϵ) power-simulates (p, η) by \mathcal{R} , and for all $P, \mu'(P)[\text{psim}(\mathcal{R} \otimes \Delta, \mathcal{R} \otimes \Delta)]\mu(P)$. When C is an atomic client operation a , the theorem holds because of Lemma 4. The remaining three cases follow from the closedness of $\text{psim}(\mathcal{R} \otimes \Delta, \mathcal{R} \otimes \Delta)$ in Lemma 5 and 6: $\text{psim}(\mathcal{R} \otimes \Delta, \mathcal{R} \otimes \Delta)$ is closed under arbitrary union and seq . This closedness property directly implies that the induction step goes through for the cases of $C_1 \llbracket C_2$ and $C_1; C_2$. For $\text{fix } P.C'$, we note that the closedness under arbitrary union implies that $\text{psim}(\mathcal{R} \otimes \Delta, \mathcal{R} \otimes \Delta)$ is complete,¹⁰ and that this completeness is what we need to prove the induction step for $\text{fix } P.C'$. \square

Lemma 7 (Identity Extension). A module (q, ϵ) data-refines another module (p, η) iff for all complete commands C , we have that $\llbracket C \rrbracket_{(q, \epsilon)}^c [\text{psim}(\text{ID}, \text{ID})] \llbracket C \rrbracket_{(p, \eta)}^c$.

Theorem 2 (Soundness). If a module (q, ϵ) power-simulates another module (p, η) by an admissible power relation $\mathcal{R} \subseteq q \times \wp(p)$, then (q, ϵ) data-refines (p, η) .

Proof. Suppose that a module (q, ϵ) power-simulates another module (p, η) by an admissible power relation $\mathcal{R} \subseteq q \times \wp(p)$. We will show that for all complete commands C , $\llbracket C \rrbracket_{(q, \epsilon)} [\text{psim}(\text{ID}, \text{ID})] \llbracket C \rrbracket_{(p, \eta)}$, because, then, module (q, ϵ) should data-refine (p, η) (Lemma 7). Pick an arbitrary complete program C . Let μ be an environment that maps all program identifiers to the empty relation. By Lemma 6, $\mu(P)[\text{psim}(\mathcal{R} \otimes \Delta, \mathcal{R} \otimes \Delta)]\mu(P)$ for all P in pid . From this, we derive the required relationship as follows:

$$\begin{aligned} & (\forall P \in \text{pid}. \mu(P)[\text{psim}(\mathcal{R} \otimes \Delta, \mathcal{R} \otimes \Delta)]\mu(P)) \\ \Rightarrow & \llbracket C \rrbracket_{(q, \epsilon)} \mu [\text{psim}(\mathcal{R} \otimes \Delta, \mathcal{R} \otimes \Delta)] \llbracket C \rrbracket_{(p, \eta)} \mu && (\because \text{Theorem 1}) \\ \Rightarrow & \llbracket C \rrbracket_{(q, \epsilon)}^c [\text{psim}(\text{ID}, \text{ID})] \llbracket C \rrbracket_{(p, \eta)}^c && (\because \text{Lemma 5 and Def. of } \llbracket - \rrbracket^c) \end{aligned} \quad \square$$

6 Conclusion

In this paper, we have proposed a new data-refinement method, called power simulation, for programs with low-level pointer operations, and provided a non-trivial soundness proof of the method.

¹⁰ $\text{psim}(\mathcal{R}_1, \mathcal{R}_2)$ relates the least FLA to itself, and is chain-complete.

The very idea of relating a state to a state set in power simulation comes from Reddy's method for data refinement [16]. In order to have a single complete data-refinement method for a language *without pointers*, he lifted forward simulation such that all the components of the simulation become about state sets, instead of states. However, the details of the two methods, such as the admissibility condition for coupling relations and the lifting operator for commands, are completely different.

References

1. A. Banerjee and D. A. Naumann. Representation independence, confinement and access control (extended abstract). In *POPL'02*, pages 166–177. ACM, 2002.
2. C. Boyapati, B. Liskov, and L. Shrira. Ownership types for object encapsulation. In *POPL'03*, pages 213–223. ACM, 2003.
3. D. G. Clarke, J. Noble, and J. M. Potter. Simple ownership types for object containment. In *ECOOP'01*, volume 2072 of *Lecture Notes in Computer Science*, pages 53–76. Springer-Verlag, 2001.
4. W.-P. de Roever and K. Engelhardt. *Data Refinement: Model-Oriented Proof Methods and their Comparison*. Number 47 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge, UK, 1998.
5. E. Dijkstra. *A Discipline of Programming*. Prentice-Hall, Englewood Cliffs, 1976.
6. J. He, C. A. R. Hoare, and J. W. Sanders. Data refinement refined. In *ESOP'86*, volume 213 of *Lecture Notes in Computer Science*, pages 187–196. Springer-Verlag, 1986.
7. C. A. R. Hoare. Proof of correctness of data representations. *Acta Informatica*, 1:271–281, 1972.
8. J. Hogg. Islands: Aliasing protection in object-oriented languages. In *OOPLA'91*, pages 271–285. ACM, 1991.
9. J. Hogg, D. Lea, A. Wills, D. deChampeaux, and R. Holt. The geneva convention on the treatment of object aliasing. *OOPS Messenger*, 3(2):11–16, 1992.
10. S. Istiaq and P. W. O'Hearn. BI as an assertion language for mutable data structures. In *POPL'01*, pages 14–26, London, 2001. ACM.
11. K. R. M. Leino and G. Nelson. Data abstraction and information hiding. *ACM Trans. on Program. Lang. and Syst.*, 24(5):491–553, 2002.
12. I. Mijajlovic, N. Torp-Smith, and P. O'Hearn. Refinement and separation contexts. In *FSTTCS'04*, volume 3328 of *Lecture Notes in Computer Science*, pages 421–433. Springer-Verlag, 2004.
13. I. Mijajlovic and H. Yang. Data refinement with low-level pointer operations. Manuscript, 2005. Available at <http://ropas.snu.ac.kr/~hyang/paper/full-ps.ps>.
14. D. Naumann and M. Barnett. Towards imperative modules: Reasoning about invariants and sharing of mutable state. In *LICS'04*, pages 313–323. IEEE, 2004.
15. P. W. O'Hearn, H. Yang, and J. C. Reynolds. Separation and information hiding. In *POPL'04*, pages 268–280, Venice, 2004. ACM.
16. U. S. Reddy. Talk at MFPS'00, Hokoken, New Jersey, USA, 2000.
17. U. S. Reddy and H. Yang. Correctness of data representations involving heap data structures. *Science of Computer Programming*, 50(1):257–305, 2004.
18. J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS'02*, volume 17, pages 55 – 74, Copenhagen, 2002. IEEE.
19. I. Stark. Categorical models for local names. *Lisp and Symbolic Comput.*, 9(1):77–107, 1996.