

Abstraction and Refinement for Local Reasoning

Thomas Dinsdale-Young, Philippa Gardner and Mark Wheelhouse

Department of Computing, Imperial College London, 180 Queen's Gate, London, SW7 2AZ

Received February 2011

Local reasoning has become a well-established technique in program verification, which has been shown to be useful at many different levels of abstraction. In separation logic, we use a low-level abstraction that is close to how the machine sees the program state. In context logic, we work with high-level abstractions that are close to how the clients of modules see the program state. We apply program refinement to local reasoning, demonstrating that high-level, abstract local reasoning is sound for module implementations. We consider two approaches: one that preserves the high-level locality at the low level; and one that breaks the high-level ‘fiction’ of locality.

Contents

| | | |
|-----|---|----|
| 1 | Introduction | 2 |
| 2 | Preliminaries | 5 |
| 2.1 | Programming Language | 5 |
| 2.2 | State Model | 5 |
| 2.3 | Axiomatic Semantics | 7 |
| 3 | Abstract Modules | 12 |
| 3.1 | Heap Module | 12 |
| 3.2 | Tree Module | 14 |
| 3.3 | List Module | 16 |
| 3.4 | Combining Abstract Modules | 20 |
| 4 | Module Translations | 21 |
| 4.1 | Modularity | 22 |
| 5 | Locality-Preserving Translations | 22 |
| 5.1 | Proof of Soundness of Locality-Preserving Translations | 27 |
| 5.2 | Eliminating Crust | 31 |
| 5.3 | Including the Conjunction Rule | 31 |
| 5.4 | Module Translation: $\tau_1 : \mathbb{T} \rightarrow \mathbb{H}$ | 32 |
| 5.5 | Module Translation: $\tau_2 : \mathbb{T} \rightarrow \mathbb{H} + \mathbb{L}$ | 37 |
| 5.6 | Module Translation: $\tau_3 : \mathbb{H} + \mathbb{H} \rightarrow \mathbb{H}$ | 42 |
| 6 | Locality-Breaking Translations | 44 |
| 6.1 | Proof of Soundness of Locality-Breaking Translations | 46 |

| | | |
|-----|--|----|
| 6.2 | Including the Conjunction Rule | 49 |
| 6.3 | Module Translation: $\tau_4 : \mathbb{L} \rightarrow \mathbb{H}$ | 49 |
| 7 | Conclusions | 53 |
| 7.1 | On Locality-Preserving versus Locality-Breaking | 53 |
| 7.2 | On Abstract Predicates | 55 |
| 7.3 | On Data Refinement | 55 |
| 7.4 | On Concurrency | 56 |

1. Introduction

Hoare logic (?) is an important tool for formally proving correctness properties of programs. It takes advantage of modularity by treating program fragments in terms of provable specifications. However, heap programs tend to break this type of modular reasoning by permitting pointer aliasing. For instance, the specification that a program reverses one list does not imply that it leaves a second list alone. To achieve this disjointness property, it is necessary to establish disjointness conditions throughout the proof.

O’Hearn, Reynolds, and Yang (2001) introduced separation logic for reasoning *locally* about heap programs, in order to address this problem. The fundamental principle of local reasoning is that, if we know how a local computation behaves on some state, then we can infer the behaviour when the state is extended: it simply leaves the additional state unchanged. A program is specified in terms of its *footprint* — the resource necessary for it to operate — and a *frame rule* is used to infer that any additional resource is indeed unchanged. For example, given a proof that a program reverses a list, the frame rule can directly establish that the program leaves a second disjoint list alone. Consequently, separation logic enables modular reasoning about heap programs. Calcagno, O’Hearn, and Yang (2007b) developed *abstract* separation logic to provide a general theory and semantic basis for separation logics based on variants of the heap model.

Abstraction (see *e.g.* ??) and refinement (see *e.g.* Hoare, 1972; de Roeper and Engelhart, 1999) are also essential for modular reasoning. Abstraction takes a concrete program and produces an abstract specification; refinement takes an abstract specification and produces a correct implementation. Both approaches result in a program that correctly implements an abstract specification. Such a result is essential for modularity because it means that a program can be replaced by any other program that meets the same specification. Abstraction and refinement are well-established techniques in program verification, but have so far not been fully understood in the context of local reasoning.

Parkinson and Bierman (2005) introduced abstract predicates in separation logic to provide abstract reasoning. An abstract predicate is, to the client, an opaque object that encapsulates the unknown representation of an abstract datatype. They inherit some of the benefits of locality from separation logic: an operation on one abstract predicate leaves others alone. However, the client cannot take advantage of local behaviour that is provided by the abstraction itself.

Consider a set module. The operation of removing, say, the value 3 from the set is local at the abstract level; it is independent of whether any other value is in the set.

Yet, consider an implementation of the set as a sorted, singly-linked list in the heap, starting from address h . The operation of removing 3 from the set must traverse the list from h . The footprint, therefore, comprises the entire list segment from h up to the node with value 3. With abstract predicates, the abstract footprint corresponds to the concrete footprint and hence, in this case, includes all the elements of the set less than or equal to 3. Consequently, abstract predicates cannot be used to present a local abstract specification for removing 3.

Calcagno, Gardner, and Zarfaty (2005) introduced context logic, a generalisation of separation logic, to provide *abstract local reasoning* about abstract data structures. Context logic has been used to reason about programs that manipulate data structures, such as sequences, multisets and trees (Calcagno et al., 2007a). In particular, it has been successfully applied to reason about the W3C DOM tree update library (Gardner et al., 2008). Thus far, context logic reasoning has always been justified with respect to an operational semantics defined at the same level of abstraction as the reasoning. In this paper, we combine abstract local reasoning with data refinement, to refine abstract module specifications into correct implementations.

? (?; Mijajlović et al. 2004) have previously considered data refinement for local reasoning, in the context of heap programs. They observed that a client can violate a module’s abstraction boundary by dereferencing pointers to the module’s internal state, and thereby break the refinement between abstract and concrete module implementations. In their motivating example, a simple memory allocator, a client can violate the concrete allocator’s free list through pointers to memory that has been deallocated; the abstract allocator, which maintains a free set, is unaffected by such an access, hence the refinement breaks. Their solution was to “blame the client” by introducing a modified operational semantics that treats such access violations as faulting executions. Using special simulation relations, they were able to recover soundness of data refinement. They illustrate their approach with the example of a toy memory manager. Their techniques adapt to different data store models, however, both module and client use the same model.

In this work, we apply data refinement to local reasoning, demonstrating that abstract local reasoning is sound for module implementations. By contrast with ?, we work with the axiomatic semantics, rather than operational semantics, of the language, defining proof transformations that establish that concrete implementations simulate abstract specifications. This avoids having to consider badly behaved clients, since the proof system only makes guarantees about well behaved clients. Furthermore, the abstract and concrete levels in our refinements typically have *different* data store models, meaning that the concept of locality itself is different at each level.

The motivating example of this paper is the stepwise refinement of a tree module \mathbb{T} , illustrated in Fig. 1. We present two different refinements from the tree module \mathbb{T} to the familiar heap module of separation logic \mathbb{H} . The first, τ_1 , described in §5.4, uses a direct implementation of trees in the heap in which each tree node is represented by a contiguous block of heap cells.

The second refinement uses an abstract list module \mathbb{L} as an intermediate step in the refinement. We first show how the tree module \mathbb{T} may be correctly implemented in the combination of the heap and list modules, $\mathbb{H} + \mathbb{L}$. This step, translation τ_2 , is described




Fig. 1. Module translations presented in this paper

in §5.5. We also show that the list module \mathbb{L} can be correctly implemented using the heap module \mathbb{H} . This is translation τ_4 , which is described in §6.3. Because our approach is modular, this translation can be lifted to a translation from the combined heap and list module $\mathbb{H} + \mathbb{L}$ to the combination of two heap modules $\mathbb{H} + \mathbb{H}$. (This is illustrated by the dotted arrow in Fig. 1.) To complete the refinement, in §?? we show that the double-heap module $\mathbb{H} + \mathbb{H}$ can be trivially implemented by the heap module \mathbb{H} .

Our development introduces two general techniques for verifying module implementations with respect to their local specifications, using the data refinement technique known as *forward simulation* (*L-simulation* in de Roeve and Engelhardt, 1999). We introduce *locality-preserving* and *locality-breaking* translations. Locality-preserving translations, broadly speaking, relate locality at the abstract level with locality of the implementation. However, implementations typically operate on a larger state than the abstract footprint, for instance, by performing pointer surgery on the surrounding state. We introduce the notion of *crust* to capture this additional state. This crust intrudes on the context, and so breaks the disjointness that exists at the abstract level. We therefore relate abstract locality with implementation-level locality through a *fiction of disjointness*.

With locality-breaking translations, locality at the abstract level does not correspond to locality of the implementation. Even in this case, we can think about a locality-preserving translation using possibly the whole data structure as the crust. Instead, we prove soundness by establishing that the specifications of the module commands are preserved under translation in any abstract context, showing the soundness of the abstract frame rule. We thus establish a *fiction of locality* at the abstract level.

1.1. Dedication

This paper is dedicated to the memory of Robin Milner. Milner was an early pioneer of program verification, creating the LCF system (?) for assisting a user in proving properties about programs. Milner also achieved substantial research on modularity and

abstraction: he emphasised the use of program modules to provide abstract data types using ML modules (?); he developed full abstraction (?), providing a strong link between denotational and operational semantics; he introduced the theory of bisimulation for concurrent processes, for example demonstrating that a process corresponding to a specification has the same behaviour as a process describing an implementation. More recently, he was one of the leaders of the UK Grand Challenge *Global Ubiquitous Computing: Design and Science*, in which he highlighted a key aim: ‘to develop a coherent informatic science whose concepts, calculi, theories and automated tools allow descriptive and predictive analysis of global ubiquitous computing at each level of abstraction.’ Our work resonates well with this aim, by introducing a coherent theory of abstract local reasoning which refines abstract local specifications about modules into correct implementations.

1.2. Acknowledgements

We acknowledge the support of EPSRC Programme Grant “Resource Reasoning”. Gardner acknowledges the support of a Microsoft/RAEng Senior Research Fellowship. Dinsdale-Young and Wheelhouse acknowledge the support of EPSRC DTA awards. We thank Mohammad Raza and Uri Zarfaty for detailed discussions of this work. In particular, some of the technical details in our locality-preserving translations come from an unpublished technical report *Reasoning about High-level Tree Update and its Low-level Implementation* written by Gardner and Zarfaty in 2008.

2. Preliminaries

We introduce the syntax and axiomatic semantics for a basic imperative programming language, which includes mutable variables and standard control-flow constructs such as while loops and procedure calls. As well as variables, programs operate on a mutable data store. The programming language is parametrised by the operations on the data store, so that it can be tailored to different domains. For example, standard commands for manipulating heaps, lists and trees.

2.1. Programming Language

We define the syntax of our programming language, which is parametrised by a set of basic commands Cmd , ranged over by φ . The choice of these basic commands depends on the domain over which the language is to be used: for instance, to work with the heap, commands for allocation, mutation, lookup and disposal of heap cells would be necessary; to work with a list, commands for lookup, insertion and removal of elements would be necessary; and to work with a tree, commands for lookup, node insertion and subtree deletion would be necessary.

We assume a syntax for *value expressions* $E \in \text{Expr}$ that includes variable look-up, constant values and basic arithmetic operations. Similarly, we assume a syntax for *boolean*

expressions $B \in \mathbf{BExp}$ that includes equality and inequality on value expressions, and standard boolean operators.

Definition 2.1 (Programming Language Syntax). Given a set of basic commands \mathbf{Cmd} , ranged over by φ , the language $\mathcal{L}_{\mathbf{Cmd}}$ is the set of syntactically valid programs, making use of these basic commands, ranged over by $\mathbb{C}, \mathbb{C}_1, \dots$, defined as:

$$\begin{aligned} \mathbb{C} ::= & \varphi \mid \text{skip} \mid \mathbf{x} := E \mid \mathbb{C}_1; \mathbb{C}_2 \\ & \mid \text{if } B \text{ then } \mathbb{C}_1 \text{ else } \mathbb{C}_2 \mid \text{while } B \text{ do } \mathbb{C} \\ & \mid \text{procs } \vec{r}_1 := \mathbf{f}_1(\vec{x}_1) \{ \mathbb{C}_1 \}, \dots, \vec{r}_k := \mathbf{f}_k(\vec{x}_k) \{ \mathbb{C}_k \} \text{ in } \mathbb{C} \\ & \mid \text{call } r_1, \dots, r_i := \mathbf{f}(E_1, \dots, E_j) \mid \text{local } \mathbf{x} \text{ in } \mathbb{C} \end{aligned}$$

where $\mathbf{x}, r_1, \dots \in \mathbf{Var}$ range over program variables, $\vec{x}_i, \vec{r}_i \in \mathbf{Var}^*$ range over vectors of program variables, $E, E_1, \dots \in \mathbf{Expr}$ range over expressions, $B \in \mathbf{BExp}$ ranges over Boolean expressions, and $\mathbf{f}, \mathbf{f}_1, \dots \in \mathbf{PName}$ range over procedure names. The names $\mathbf{f}_1, \dots, \mathbf{f}_k$ of procedures defined in a single **procs** – **in** block are required to be pairwise distinct. The parameter and return variables are required to be pairwise distinct within each procedure definition.

2.2. State Model

We work with multiple data structures at multiple levels of abstraction. To handle these structures in a uniform way, we model our program states using context algebras. Context algebras are a generalisation of separation algebras (Calcagno et al., 2007b) to more complex data structures. Separation algebras are based on a commutative combination of resource, whereas context algebras are instead based on non-commutative resource combinations, which is necessary to handle structured data, such as lists and trees. We will see that many interesting state models fit the pattern of a context algebra.

Definition 2.2 (Context Algebra). A *context algebra* $\mathcal{A} = (\mathbf{C}, \mathbf{D}, \circ, \bullet, \mathbf{I}, \mathbf{0})$ consists of

- a non-empty set of contexts, \mathbf{C} ,
- a non-empty set of abstract data structures, \mathbf{D} ,
- a context composition operator, $\circ : \mathbf{C} \times \mathbf{C} \rightarrow \mathbf{C}$,
- an application operator, $\bullet : \mathbf{C} \times \mathbf{D} \rightarrow \mathbf{D}$,
- a distinguished set of contexts, $\mathbf{I} \subseteq \mathbf{C}$, and
- a distinguished set of abstract data structures, $\mathbf{0} \subseteq \mathbf{D}$,

for which the following properties hold: for all $c, c', c'' \in \mathbf{C}$, $d \in \mathbf{D}$, and $i' \in \mathbf{I}$

- $c \circ (c' \circ c'') = (c \circ c') \circ c''$ (that is, composition is associative);
- $c \circ (c' \bullet d) = (c \circ c') \bullet d$ (that is, composition associates with application);
- $i \circ c$ is defined for some $i \in \mathbf{I}$, and whenever $i' \circ c$ is defined, $i' \circ c = c$ (that is, \mathbf{I} is a left identity of \circ);
- $c \circ i$ is defined for some $i \in \mathbf{I}$, and whenever $c \circ i'$ is defined, $c \circ i = c$ (that is, \mathbf{I} is a right identity of \circ);
- $i \bullet d$ is defined for some $i \in \mathbf{I}$, and whenever $i' \bullet d$ is defined, $i' \bullet d = d$ (that is, \mathbf{I} is a left identity of \bullet); and

— the relation $\{(c, d) \mid \exists o \in \mathbf{0}. c \bullet o = d\}$ is a total surjective function.

(Undefined terms are considered equal.)

Example 1. The following are examples of context algebras:

(a) *Heaps* $h \in \mathbf{Heap}$ are defined as:

$$h ::= \text{emp} \mid a \mapsto v \mid h * h$$

where $a \in \mathbf{Addr}$ ranges over unique *heap addresses*, $v \in \mathbf{Val}$ ranges over *values*, and $*$ is associative and commutative with identity emp . (Heaps are thus finite partial functions from addresses to values.) Heaps form the *heap context algebra*, $\mathcal{A}_{\mathbb{H}} = (\mathbf{Heap}, \mathbf{Heap}, *, *, \{\text{emp}\}, \{\text{emp}\})$. All separation algebras (Calcagno et al., 2007b) can be viewed as context algebras in this way.

(b) *Variable scopes* $\rho \in \mathbf{Scope}$ are defined as:

$$\rho ::= \text{emp} \mid \mathbf{x} \Rightarrow v \mid \rho * \rho$$

where $\mathbf{x} \in \mathbf{Var}$ ranges over unique *program variables*, $v \in \mathbf{Val}$ ranges over values, and $*$ is associative and commutative with identity emp . Variable scopes form the *variable scope context algebra*, $\mathcal{A}_{\mathbf{Scope}} = (\mathbf{Scope}, \mathbf{Scope}, *, *, \{\text{emp}\}, \{\text{emp}\})$. The variable scope context algebra allows us to treat variables as resource (?).

(c) *Lists* $l \in \mathbf{Lst}$ and *list contexts* $lc \in \mathbf{CLst}$ are defined as:

$$\begin{aligned} l & ::= \emptyset \mid v \mid l_1 \cdot l_2 \\ lc & ::= - \mid lc \cdot l \mid l \cdot lc \end{aligned}$$

where values $v \in \mathbf{Val}$ are taken to occur uniquely in each list or list context and \cdot is taken to be associative with identity \emptyset . Context application $\bullet : \mathbf{CLst} \times \mathbf{Lst} \rightarrow \mathbf{Lst}$ and composition $\circ : \mathbf{CLst} \times \mathbf{CLst} \rightarrow \mathbf{CLst}$ are defined in terms of substitution as:

$$\begin{aligned} lc \bullet l & \stackrel{\text{def}}{=} lc[l/-] \text{ provided that } lc[l/-] \in \mathbf{Lst} \\ lc \circ lc' & \stackrel{\text{def}}{=} lc[lc'/-] \text{ provided that } lc[lc'/-] \in \mathbf{CLst}. \end{aligned}$$

Both operators are clearly non-commutative. Lists and list contexts form the *list context algebra*, $\mathcal{A}_{\mathbf{Lst}} = (\mathbf{CLst}, \mathbf{Lst}, \circ, \bullet, \{-\}, \{\emptyset\})$.

(d) *Trees* $t \in \mathbf{Tree}$ and *tree contexts* $c \in \mathbf{C}$ are defined as:

$$\begin{aligned} t & ::= \emptyset \mid \mathbf{a}[t] \mid t \otimes t \\ c & ::= - \mid \mathbf{a}[c] \mid t \otimes c \mid c \otimes t \end{aligned}$$

where $\mathbf{a} \in \Sigma$ ranges over unique *node identifiers*, and \otimes is associative with identity \emptyset . Context application $\bullet : \mathbf{CTree} \times \mathbf{Tree} \rightarrow \mathbf{Tree}$ and composition $\circ : \mathbf{CTree} \times \mathbf{CTree} \rightarrow \mathbf{CTree}$ are defined in terms of substitution as:

$$\begin{aligned} c \bullet t & \stackrel{\text{def}}{=} c[t/-] \text{ provided that } c[t/-] \in \mathbf{Tree} \\ c \circ c' & \stackrel{\text{def}}{=} c[c'/-] \text{ provided that } c[c'/-] \in \mathbf{CTree}. \end{aligned}$$

Both operators are clearly non-commutative. Trees and tree contexts form the *tree context algebra*, $\mathcal{A}_{\mathbb{T}} = (\mathbf{C}, \mathbf{Tree}, \circ, \bullet, \{-\}, \{\emptyset\})$.

- (e) Let $\mathcal{A}_1 = (\mathbf{C}_1, \mathbf{D}_1, \circ_1, \bullet_1, \mathbf{I}_1, \mathbf{O}_1)$ and $\mathcal{A}_2 = (\mathbf{C}_2, \mathbf{D}_2, \circ_2, \bullet_2, \mathbf{I}_2, \mathbf{O}_2)$ be context algebras. Then their *direct product* $\mathcal{A}_1 \times \mathcal{A}_2 = (\mathbf{C}_1 \times \mathbf{C}_2, \mathbf{D}_1 \times \mathbf{D}_2, \circ_1 \times \circ_2, \bullet_1 \times \bullet_2, \mathbf{I}_1 \times \mathbf{I}_2, \mathbf{O}_1 \times \mathbf{O}_2)$ is also a context algebra[†]. For example, $\mathcal{A}_{\text{Scope}} \times \mathcal{A}_{\text{H}}$ and $\mathcal{A}_{\text{Scope}} \times \mathcal{A}_{\text{T}}$ combine variable scopes with heaps and trees, respectively.

2.3. Axiomatic Semantics

We define an axiomatic semantics for \mathcal{L}_{Cmd} as a program logic based on local Hoare reasoning. This semantics treats the space of program states as a context algebra $\mathcal{A}_{\text{State}} = (\mathbf{C}_{\text{State}}, \text{State}, \circ, \bullet, \mathbf{I}, \mathbf{O})$, which is the product of the variable scope context algebra $\mathcal{A}_{\text{Scope}}$, from Example 1(b), with some data store context algebra $\mathcal{A}_{\text{Store}}$; *i.e.* $\mathcal{A}_{\text{State}} = \mathcal{A}_{\text{Scope}} \times \mathcal{A}_{\text{Store}}$. The semantics is parametrised by the choice of $\mathcal{A}_{\text{Store}}$ and the axioms given for basic commands. This gives us a fixed way of treating program variables, but allows for a flexible choice of the structure in the data store. The treatment of variables as resource (?) also allows us to avoid side conditions in our proof rules.

The Hoare logic judgements of our proof system make assertions about program state and have the form $\Gamma \vdash \{P\} \mathbb{C} \{Q\}$, where $P, Q \in \text{State}$ are predicates, $\mathbb{C} \in \mathcal{L}_{\text{Cmd}}$ is a program, and Γ is a procedure specification environment. A *procedure specification environment* associates procedure names with pairs of pre- and postconditions (parameterised by the arguments and return values of the procedure respectively). The interpretation of judgements is that, in the presence of procedures satisfying Γ , when executed from a state satisfying P , the program \mathbb{C} will either diverge or terminate in a state satisfying Q .

For simplicity, we use semantic predicates as assertions (so assertions describe sets of states), rather than using logical formulae; this reflects the practice of Calcagno et al. (2007b).

Definition 2.3 (Predicates). The set of *state predicates* $\mathcal{P}(\text{State})$, ranged over by P, Q, R, P', P_1, \dots , is the set of sets of states. The set of *state-context predicates* $\mathcal{P}(\mathbf{C}_{\text{State}})$, ranged over by K, K', K_1, \dots , is the set of sets of state contexts.

To express predicates, we use standard logical notation for conjunction, disjunction, negation and quantification, which are interpreted with the usual semantics (intersection, union, *etc.*). We also lift operations on states and contexts to predicates: for instance, $a \mapsto v$ denotes the predicate $\{a \mapsto v\}$; $\exists v. a \mapsto v$ denotes $\{a \mapsto v \mid v \in \text{Val}\}$; $P * Q$ denotes $\{d_1 * d_2 \mid d_1 \in P \text{ and } d_2 \in Q\}$; the separating application $K \bullet P$ denotes $\{c \bullet d \mid c \in K \text{ and } d \in P\}$; and so on. We use $a \mapsto -$ as shorthand for $\exists v. a \mapsto v$ and similarly for $\mathbf{x} \Rightarrow -$. We use \prod^* to denote iterated $*$, and $K \multimap K'$ to denote $\{c_1 \mid \text{for all } c_2 \in K \text{ and } c \in c_1 \circ c_2, c \in K'\}$. We use \equiv to indicate equality between predicates, \subseteq to indicate that one predicate is contained within another (the first entails the second), and \in to indicate that a state or context belongs to a predicate.

Recall from our definition of \mathcal{L}_{Cmd} 2.1 that procedures have the form $\vec{r} := \mathbf{f}(\vec{\mathbf{x}})$.

[†] The product of partial functions is defined pointwise in the natural fashion.

Definition 2.4 (Procedure Specification Environment). A procedure specification $\mathbf{f} : \mathbf{P} \rightarrow \mathbf{Q}$ comprises

- a procedure name $\mathbf{f} \in \mathbf{PName}$,
- a parametrised precondition $\mathbf{P} : \mathbf{Val}^n \rightarrow \mathcal{P}(\mathbf{State})$, and
- a parametrised postcondition $\mathbf{Q} : \mathbf{Val}^m \rightarrow \mathcal{P}(\mathbf{State})$,

where $\vec{\mathbf{r}} := \mathbf{f}(\vec{\mathbf{x}})$, $|\vec{\mathbf{r}}| = n$ and $|\vec{\mathbf{x}}| = m$.

A *procedure specification environment* is a set of procedure specifications. The metavariables Γ, Γ', \dots range over procedure specification environments.

Notation.

In proof judgements, Γ, Γ' stands for the union $\Gamma \cup \Gamma'$.

We assume a denotational semantics for value expressions,

$$\mathcal{E} \llbracket (\cdot) \rrbracket : \mathbf{Expr} \rightarrow (\mathbf{Scope} \rightarrow \mathbf{Val}),$$

and for boolean expressions,

$$\mathcal{B} \llbracket (\cdot) \rrbracket : \mathbf{BExp} \rightarrow (\mathbf{Scope} \rightarrow \mathbf{Bool}).$$

where $\mathbf{Bool} \stackrel{\text{def}}{=} \{\mathbf{T}, \mathbf{F}\}$ is the standard two-valued set; \mathbf{T} represents that a boolean expression holds for a given variable scope, and \mathbf{F} represents that it does not. Note that the semantics of an expression can be undefined on a given variable scope, for instance if some variable in the expression is not assigned in the scope, or if evaluating the expression requires division by zero.

To simplify the presentation of our inference rules we define a predicate-valued semantics for boolean expressions. This semantics interprets a boolean expression as the set of states in which that boolean expression holds.

Definition 2.5 (Predicate-Valued Semantics of Boolean Expressions). The *predicate-valued semantics of Boolean expressions*, $\mathcal{P} \llbracket (\cdot) \rrbracket : \mathbf{BExp} \rightarrow \mathcal{P}(\mathbf{State})$, is defined in terms of their truth-valued semantics as follows:

$$\mathcal{P} \llbracket B \rrbracket = \{(\rho, \chi) \mid \mathcal{B} \llbracket B \rrbracket \rho = \mathbf{T}\}.$$

Since the semantics of expressions is partial, it is also convenient to define safety predicates, which simply assert that the state permits the evaluation of a given expression.

Definition 2.6 (Safety Predicates). Given a value expression, $E \in \mathbf{Expr}$, the *expression safety predicate for E*, $vsafe(E)$, is defined as:

$$vsafe(E) = \{(\rho, \chi) \mid \mathcal{E} \llbracket E \rrbracket \rho \text{ is defined}\}.$$

Similarly, given a Boolean expression, $B \in \mathbf{BExp}$, the *expression safety predicate for B*, $bsafe(B)$, is defined as:

$$bsafe(B) = \{(\rho, \chi) \mid \mathcal{B} \llbracket B \rrbracket \rho \text{ is defined}\}.$$

Finally, in order to define the axiomatic semantics, we need axioms for the basic commands of the language.

Assumption 1 (Axioms for Basic Commands). Assume a set of axioms for the basic commands,

$$\text{Ax} \llbracket (\cdot) \rrbracket : \text{Cmd} \rightarrow \mathcal{P}(\mathcal{P}(\text{State}) \times \mathcal{P}(\text{State})).$$

Definition 2.7 (Inference Rules). The Hoare logic rules for \mathcal{L}_{Cmd} are given in Fig. 2.

The AXIOM rule allows us to use the given specifications of our basic commands and the FRAME rule is the natural generalisation of the frame rule for separation algebras to context algebras. The CONS, DISJ, SKIP, SEQ, IF and WHILE rules are standard. The ASSGN, LOCAL, PDEF, PCALL and PWK rules are standard, but adapted to our treatment of variables as resource. The ASSGN rule not only requires the resource $\mathbf{x} \Rightarrow v$, but also the resource ρ containing the other variables used in E . For the LOCAL rule, recall that the predicate P specifies a set of pairs consisting of resource from $\mathcal{D}_{\text{Store}}$ and variable resource. The predicate $(\mathbf{x} \Rightarrow - \times \mathbf{I}_{\text{Store}}) \bullet P$ therefore extends the variable component with variable \mathbf{x} of indeterminate initial value. If a local variable block is used to re-declare a variable that is already in scope, the FRAME rule must be used add the variable's outer scope after the LOCAL rule is applied. For the PDEF and PCALL rules, the procedure \mathbf{f} has parametrised predicates $\mathbf{P} = \lambda \vec{x}. P$ and $\mathbf{Q} = \lambda \vec{r}. Q$ as its pre- and postcondition, with $\mathbf{P}(\vec{v}) = P[\vec{v}/\vec{x}]$ and $\mathbf{Q}(\vec{w}) = Q[\vec{w}/\vec{r}]$; the parameters carry the call and return values of the procedure. The PWK rule simply allows for the procedure specification environment to be weakened (i.e. more procedure specifications can be added).

The conjunction rule is notably absent from Fig. 2.:

$$\frac{I \neq \emptyset \quad \text{for all } i \in I, \Gamma \vdash \{P_i\} \mathbb{C} \{Q_i\}}{\Gamma \vdash \{\bigwedge_{i \in I} P_i\} \mathbb{C} \{\bigwedge_{i \in I} Q_i\}} \text{CONJ}.$$

We give conditions under which conjunction is admissible. To justify our ideas, we first study an example where conjunction is not admissible. Consider the command `allocEither()` that operates on a double-heap data store ($\mathcal{A}_{\text{Store}} = \mathcal{A}_{\text{H}} \times \mathcal{A}_{\text{H}}$), allocating a single cell in either of the two heaps. We can specify this command by:

$$\text{Ax} \llbracket \mathbf{x} := \text{allocEither}() \rrbracket \stackrel{\text{def}}{=} \left\{ \begin{array}{l} ((\mathbf{x} \Rightarrow - \times \text{emp} \times \text{emp}), (\exists a. \mathbf{x} \Rightarrow a \times a \mapsto - \times \text{emp})), \\ ((\mathbf{x} \Rightarrow - \times \text{emp} \times \text{emp}), (\exists a. \mathbf{x} \Rightarrow a \times \text{emp} \times a \mapsto -)) \end{array} \right\}$$

Note that the choice of heap in which the cell is allocated is not at the discretion of the implementation, but rather at the discretion of the prover. The command exhibits *angelic nondeterminism*: that is, it is possible to prove both that the command allocates the cell in the left heap and that the command allocates the cell in the right heap. This seems paradoxical, since the program must somehow correctly guess the prover's choice.

Remark. One way of resolving this paradox is that, from the program's perspective, the two cases are actually *the same* and the distinction is only a logical abstraction. This is exactly the case for the implementation of a double heap in a single heap that we consider in §5.6.

$$\begin{array}{c}
\frac{(P, Q) \in \text{Ax}[\llbracket \varphi \rrbracket]}{\Gamma \vdash \{P\} \varphi \{Q\}} \text{AXIOM} \quad \frac{\Gamma \vdash \{P\} \mathbb{C} \{Q\}}{\Gamma \vdash \{K \bullet P\} \mathbb{C} \{K \bullet Q\}} \text{FRAME} \\
\\
\frac{P \subseteq P' \quad \Gamma \vdash \{P'\} \mathbb{C} \{Q'\} \quad Q' \subseteq Q}{\Gamma \vdash \{P\} \mathbb{C} \{Q\}} \text{CONS} \\
\\
\frac{\text{for all } i \in I, \Gamma \vdash \{P_i\} \mathbb{C} \{Q_i\}}{\Gamma \vdash \{\bigvee_{i \in I} P_i\} \mathbb{C} \{\bigvee_{i \in I} Q_i\}} \text{DISJ} \\
\\
\frac{}{\Gamma \vdash \{\mathbf{0}\} \text{skip} \{\mathbf{0}\}} \text{SKIP} \\
\\
\frac{\Gamma \vdash \{P\} \mathbb{C}_1 \{R\} \quad \Gamma \vdash \{R\} \mathbb{C}_2 \{Q\}}{\Gamma \vdash \{P\} \mathbb{C}_1; \mathbb{C}_2 \{Q\}} \text{SEQ} \\
\\
\frac{P \subseteq \text{bsafe}(B) \quad \Gamma \vdash \{P \wedge \mathcal{P} \llbracket B \rrbracket\} \mathbb{C}_1 \{Q\} \quad \Gamma \vdash \{P \wedge \neg \mathcal{P} \llbracket B \rrbracket\} \mathbb{C}_2 \{Q\}}{\Gamma \vdash \{P\} \text{if } B \text{ then } \mathbb{C}_1 \text{ else } \mathbb{C}_2 \{Q\}} \text{IF} \\
\\
\frac{P \subseteq \text{bsafe}(B) \quad \Gamma \vdash \{P \wedge \mathcal{P} \llbracket B \rrbracket\} \mathbb{C} \{P\}}{\Gamma \vdash \{P\} \text{while } B \text{ do } \mathbb{C} \{P \wedge \neg \mathcal{P} \llbracket B \rrbracket\}} \text{WHILE} \\
\\
\frac{(\mathbf{x} \Rightarrow v * \rho, \chi) \in \text{vsafe}(E)}{\Gamma \vdash \{(\mathbf{x} \Rightarrow v * \rho, \chi)\} \mathbf{x} := E \{(\mathbf{x} \Rightarrow \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) * \rho, \chi)\}} \text{ASSGN} \\
\\
\frac{P \wedge \text{vsafe}(\mathbf{x}) \equiv \emptyset \quad \Gamma \vdash \frac{\{(\mathbf{x} \Rightarrow - \times \mathbf{I}_{\text{Store}}) \bullet P\}}{\mathbb{C}} \{(\mathbf{x} \Rightarrow - \times \mathbf{I}_{\text{Store}}) \bullet Q\}}{\Gamma \vdash \{P\} \text{local } \mathbf{x} \text{ in } \mathbb{C} \{Q\}} \text{LOCAL} \\
\\
\text{for all } (\mathbf{f}_i : P \mapsto Q) \in \Gamma, \Gamma', \Gamma \vdash \frac{\{\exists \vec{v}. \vec{x}_i \Rightarrow \vec{v} * \vec{r}_i \Rightarrow - \times P(\vec{v})\}}{\mathbb{C}_i} \{\exists \vec{w}. \vec{x}_i \Rightarrow - * \vec{r}_i \Rightarrow \vec{w} \times Q(\vec{w})\}} \\
\text{for all } \mathbf{f} : P \mapsto Q \in \Gamma, \text{ there exists } i \text{ s.t. } \mathbf{f} = \mathbf{f}_i \\
\text{for all } \mathbf{f} : P \mapsto Q \in \Gamma', \text{ for all } i, \mathbf{f} \neq \mathbf{f}_i \\
\Gamma', \Gamma \vdash \{P\} \mathbb{C} \{Q\} \\
\frac{}{\Gamma' \vdash \{P\} \text{procs } \vec{r}_1 := \mathbf{f}_1(\vec{x}_1) \{C_1\}, \dots, \vec{r}_k := \mathbf{f}_k(\vec{x}_k) \{C_k\} \text{ in } \mathbb{C} \{Q\}} \text{PDEF} \\
\\
\frac{\{(\vec{r} \Rightarrow \vec{v} * \rho)\} \times \text{Store} \subseteq \text{vsafe}(\vec{E})}{\Gamma, \mathbf{f} : P \mapsto Q \vdash \frac{\{(\vec{r} \Rightarrow \vec{v} * \rho)\} \times P(\vec{\mathcal{E}} \llbracket E \rrbracket (\vec{r} \Rightarrow \vec{v} * \rho))\}}{\text{call } \vec{r} := \mathbf{f}(\vec{E})} \{ \exists \vec{w}. \{(\vec{r} \Rightarrow \vec{w} * \rho)\} \times Q(\vec{w}) \}} \text{PCALL} \\
\\
\frac{\Gamma \vdash \{P\} \mathbb{C} \{Q\}}{\Gamma, \Gamma' \vdash \{P\} \mathbb{C} \{Q\}} \text{PWK}
\end{array}$$

Fig. 2. Local Hoare logic rules for \mathcal{L}_{Cmd}

The conjunction rule is not compatible with angelic nondeterminism, as illustrated by the following derivation:

$$\frac{\frac{\frac{\{x \Rightarrow - \times \text{emp} \times \text{emp}\}}{x := \text{allocEither}()} \text{AXIOM}}{\{\exists a. x \Rightarrow a \times a \mapsto - \times \text{emp}\}} \text{AXIOM}}{\{x \Rightarrow - \times \text{emp} \times \text{emp}\} x := \text{allocEither}() \{\text{false}\}} \text{CONJ}$$

With CONJ, it must be the case that `allocEither` diverges. Without CONJ, we cannot draw the same conclusion.

The following two conditions on basic command $\varphi \in \text{Cmd}$ are sufficient to establish that the command does not exhibit angelic nondeterminism:

- for all $(P, Q), (P', Q') \in \text{Ax} \llbracket \varphi \rrbracket$ with $(P, Q) \neq (P', Q')$, $P \wedge P' \equiv \emptyset$; and
- the predicate $\bigvee \{P \mid (P, Q) \in \text{Ax} \llbracket \varphi \rrbracket\}$ is precise[‡].

These conditions imply that at most one axiom describes the behaviour of the command from any given state, and hence the conjunction rule cannot be used to derive a stronger postcondition for any of the basic commands. These conditions hold for all of the basic module commands considered in this paper; they are not difficult to check. None of the program constructions introduces angelic nondeterminism[§]. CONJ is therefore admissible for all of the modules considered here. Since nothing is gained by the inclusion of the conjunction rule, its omission is justified. When we later discuss our reasoning techniques in §5 and §6, we briefly consider the consequences of including the conjunction rule in our theory.

3. Abstract Modules

An abstract module is a collection of operations on some abstract state model. For example, a heap module typically provides operations that allocate and dispose blocks of heap cells, and that fetch and mutate values stored in heap cells. Similarly, a list module provides operations for adding, removing and querying list elements, while a tree module provides operations for traversing the tree structure, and adding, removing and moving nodes or subtrees.

The programming language that we introduced in the previous section can be instantiated for such abstract modules by the choice of basic commands `Cmd`, data store context algebra $\mathcal{A}_{\text{Store}} = (\text{C}_{\text{Store}}, \text{Store}, \circ, \bullet, \mathbf{I}_{\text{Store}}, \mathbf{O}_{\text{Store}})$, and axiomatisation of the basic commands $\text{Ax} \llbracket (\cdot) \rrbracket$. Together, these three parameters constitute the notion of an abstract module in our formalism.

Definition 3.1 (Abstract Module). An *abstract module*

$$\mathbb{A} = (\text{Cmd}_{\mathbb{A}}, \mathcal{A}_{\mathbb{A}}, \text{Ax} \llbracket (\cdot) \rrbracket_{\mathbb{A}})$$

[‡] Predicate P is *precise* if, for every $d \in \text{State}$ there is at most one $c \in \text{C}_{\text{State}}$ and $d' \in P$ such that $d = c \bullet d'$.

[§] Proving this statement is rather involved, especially for recursion and looping. The proof invokes Tarski's fixed-point theorem and therefore requires a proof of monotonicity of programs.

consists of:

- a set of basic commands, $\text{Cmd}_{\mathbb{A}}$;
- a context algebra $\mathcal{A}_{\mathbb{A}} = (\text{C}_{\mathbb{A}}, \text{D}_{\mathbb{A}}, \circ_{\mathbb{A}}, \bullet_{\mathbb{A}}, \mathbf{I}_{\mathbb{A}}, \mathbf{O}_{\mathbb{A}})$; and
- an axiomatisation for the basic commands

$$\text{Ax} \llbracket (\cdot) \rrbracket_{\mathbb{A}} : \text{Cmd}_{\mathbb{A}} \rightarrow \mathcal{P}(\mathcal{P}(\text{State}_{\mathbb{A}}) \times \mathcal{P}(\text{State}_{\mathbb{A}})),$$

where $\text{State}_{\mathbb{A}} = \text{Scope} \times \text{D}_{\mathbb{A}}$.

Notation.

We denote the language determined by the abstract module \mathbb{A} as $\mathcal{L}_{\mathbb{A}}$ (this is in fact $\mathcal{L}_{\text{Cmd}_{\mathbb{A}}}$). We denote the axiomatic semantic judgement determined by the abstract module \mathbb{A} as $\vdash_{\mathbb{A}}$. When the abstract module \mathbb{A} can be inferred from context, the subscript \mathbb{A} may be dropped.

In this section, we consider a number of different abstract modules, including heaps (the basis of separation logic (Ishtiaq and O’Hearn, 2001; Reynolds, 2002)), trees (the original motivation for context logic (Calcagno, Gardner, and Zarfaty, 2005)) and lists. A number of other abstract modules that are easily encoded in our formalism have been considered in the literature, such as heap with free set (Raza and Gardner, 2009), a fragment of the W3C DOM (Gardner, Smith, Wheelhouse, and Zarfaty, 2008), and tree segments (Gardner and Wheelhouse, 2009).

3.1. Heap Module

The abstract heap module $\mathbb{H} \stackrel{\text{def}}{=} (\text{Cmd}_{\mathbb{H}}, \mathcal{A}_{\mathbb{H}}, \text{Ax} \llbracket (\cdot) \rrbracket_{\mathbb{H}})$ should be familiar to aficionados of separation logic; its commands consist of heap allocation, disposal, mutation and lookup. Heaps are modelled as finite partial functions from heap addresses (Addr) to values (Val). The address set is assumed to be the positive integers, *i.e.* $\text{Addr} = \mathbb{Z}^+$, and contained within the value set, *i.e.* $\text{Addr} \subseteq \text{Val}$. This enables program variables and heap cells to hold pointers to other heap cells and arithmetic operations to be performed on heap addresses (pointer arithmetic).

The constant value **nil**, the *null reference*, is used in situations where a heap reference could occur, to indicate the absence of such a reference. It is therefore necessary that $\mathbf{nil} \notin \text{Addr}$, so that it cannot be confused with a valid heap reference, and that $\mathbf{nil} \in \text{Val}$, so that it may be stored in variables, heap cells, *etc.*. In particular, we take $\mathbf{nil} = 0$. The set $\text{Addr}_{\mathbf{nil}} \stackrel{\text{def}}{=} \text{Addr} \cup \{\mathbf{nil}\}$ consists of all valid references and the null reference.

Definition 3.2 (Heap Update Commands). The set of *heap update commands* $\text{Cmd}_{\mathbb{H}}$, ranged over by φ , is defined as follows:

$$\varphi ::= \mathbf{x} := \text{alloc}(E) \mid \text{dispose}(E_1, E_2) \mid [E_1] := E_2 \mid \mathbf{x} := [E]$$

where $\mathbf{x} \in \text{Var}$ ranges over variables and $E, E_1, E_2 \in \text{Expr}$ range over value expressions.

The intuitive meaning of these commands, which will be realised by their axiomatic semantics, is as follows:

- $\mathbf{x} := \text{alloc}(E)$ allocates a contiguous block of cells in the heap of length E , returning the address of the first cell in \mathbf{x} ;
- $\text{dispose}(E_1, E_2)$ deallocates a contiguous block of cells in the heap at address E_1 of length E_2 ;
- $[E_1] := E_2$ stores the value E_2 in the heap cell at address E_1 ; and
- $\mathbf{x} := [E]$ loads the contents of the heap cell at address E_1 into \mathbf{x} .

The data model for heaps, which we have already seen in Example 1(a), represents heaps as, effectively, partial functions from heap addresses to values. In this model, heaps are interpreted as resources: the cells that have values specified in the heap are the resources available to the program. Loads and stores can only be performed on heap cells that are available to the program; allocation makes new heap cells available; and deallocation makes available heap cells unavailable.

Definition 3.3 (Heap Context Algebra). The *heap context algebra* is:

$$\mathcal{A}_{\mathbb{H}} \stackrel{\text{def}}{=} (\text{Heap}, \text{Heap}, *, *, \{\text{emp}\}, \{\text{emp}\})$$

as given in Example 1(a).

Definition 3.4 (Heap Axiomatisation).

The *heap axiomatisation*, $\text{AX} \llbracket (\cdot) \rrbracket_{\mathbb{H}} : \text{Cmd}_{\mathbb{H}} \rightarrow \mathcal{P}(\mathcal{P}(\text{Scope} \times \text{Heap}) \times \mathcal{P}(\text{Scope} \times \text{Heap}))$ is defined as follows:

$$\text{AX} \llbracket \mathbf{x} := \text{alloc}(E) \rrbracket_{\mathbb{H}} \stackrel{\text{def}}{=} \left\{ \left(\left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times \text{emp} \wedge w \geq 1), \\ \left(\begin{array}{l} \exists a. \mathbf{x} \Rightarrow a * \rho \times \\ a \mapsto - * \dots * (a + w) \mapsto - \end{array} \right) \end{array} \right) \right) \mid w = \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) \right\}$$

$$\text{AX} \llbracket \text{dispose}(E_1, E_2) \rrbracket_{\mathbb{H}} \stackrel{\text{def}}{=} \left\{ ((\rho \times a \mapsto - * \dots * (a + v) \mapsto -), (\rho \times \text{emp})) \mid \begin{array}{l} a = \mathcal{E} \llbracket E_1 \rrbracket \rho \text{ and} \\ v = \mathcal{E} \llbracket E_2 \rrbracket \rho \end{array} \right\}$$

$$\text{AX} \llbracket [E_1] := E_2 \rrbracket_{\mathbb{H}} \stackrel{\text{def}}{=} \left\{ ((\rho \times a \mapsto -), (\rho \times a \mapsto v)) \mid a = \mathcal{E} \llbracket E_1 \rrbracket \rho \text{ and } v = \mathcal{E} \llbracket E_2 \rrbracket \rho \right\}$$

$$\text{AX} \llbracket \mathbf{x} := [E] \rrbracket_{\mathbb{H}} \stackrel{\text{def}}{=} \left\{ ((\mathbf{x} \Rightarrow v * \rho \times a \mapsto w), (\mathbf{x} \Rightarrow w * \rho \times a \mapsto w)) \mid a = \mathcal{E} \llbracket E \rrbracket \rho \right\}$$

3.2. Tree Module

The abstract tree module $\mathbb{T} \stackrel{\text{def}}{=} (\text{Cmd}_{\mathbb{T}}, \mathcal{A}_{\mathbb{T}}, \text{AX} \llbracket (\cdot) \rrbracket_{\mathbb{T}})$ should be familiar to adherents of context logic; its commands consist of node-relative traversal, node creation and subtree deletion. The tree model consists of *uniquely-labelled* trees, where each label may only

occur once in any given tree or context. This ensures that nodes in a tree are uniquely addressable by their labels. It is therefore assumed that the set of tree labels, Σ , is contained within the value set, Val , *i.e.* $\Sigma \subseteq \text{Val}$. It is also assumed that $\text{Addr} \subseteq \Sigma$, in order that heap addresses can be used to implement node addresses.

Definition 3.5 (Tree Update Commands). The set of *tree update commands* $\text{Cmd}_{\mathbb{T}}$, ranged over by φ , is defined as follows:

$$\begin{aligned} \varphi ::= & \mathbf{x} := \text{getUp}(E) \mid \mathbf{x} := \text{getLeft}(E) \mid \mathbf{x} := \text{getRight}(E) \\ & \mid \mathbf{x} := \text{getFirst}(E) \mid \mathbf{x} := \text{getLast}(E) \\ & \mid \text{newNodeAfter}(E) \mid \text{deleteTree}(E) \end{aligned}$$

where $\mathbf{x} \in \text{Var}$ ranges over variables and $E \in \text{Expr}$ ranges over value expressions.

The intuitive meaning of these commands is as follows:

- $\text{getUp}(E)$, $\text{getLeft}(E)$ and $\text{getRight}(E)$ retrieve, respectively, the identifier of the immediate parent, left sibling and right sibling (if any) of the node identified by E ;
- $\text{getFirst}(E)$ and $\text{getLast}(E)$ retrieve, respectively, the identifiers of the first and last children (if any) of the node identified by E ;
- $\text{newNodeAfter}(E)$ creates a new node with some fresh identifier and no children, which is inserted into the tree as the right sibling of the node identified by E ; and
- $\text{deleteTree}(E)$ deletes the entire subtree rooted at the node identified by E .

The data model for trees, which we have already seen in Example 1(d), represents trees as uniquely-labelled abstract forests. These trees are divisible into contexts and subtrees, with subtrees always being trees in their own right. We thus give the axiomatisation of the tree update commands in terms of such complete trees, although they may actually be subtrees of bigger trees. (Gardner and Wheelhouse (2009) specify tree commands in terms of *tree segments*, which are fragments of trees that need not be trees in their own right. We work with subtrees here for simplicity of exposition.)

Definition 3.6 (Uniquely-Labelled Tree Context Algebra). The *tree context algebra* is:

$$\mathcal{A}_{\mathbb{T}} \stackrel{\text{def}}{=} (\mathbf{C}_{\text{Tree}}, \text{Tree}, \circ, \bullet, \{-\}, \{\mathbf{0}\})$$

as given in Example 1(d).

Definition 3.7 (Tree Axiomatisation).

The *tree axiomatisation*, $\text{AX} \llbracket (\cdot) \rrbracket_{\mathbb{T}} : \text{Cmd}_{\mathbb{T}} \rightarrow \mathcal{P}(\mathcal{P}(\text{Scope} \times \text{Tree}) \times \mathcal{P}(\text{Scope} \times \text{Tree}))$ is defined as follows:

$$\text{AX} \llbracket \mathbf{x} := \text{getUp}(E) \rrbracket_{\mathbb{T}} \stackrel{\text{def}}{=} \left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times \mathbf{a}[t_1] \otimes \mathbf{b}[t_2] \otimes t_3), \\ (\mathbf{x} \Rightarrow \mathbf{a} * \rho \times \mathbf{a}[t_1] \otimes \mathbf{b}[t_2] \otimes t_3) \end{array} \right) \mid \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) = \mathbf{b} \right\}$$

$$\begin{aligned} \text{AX} \llbracket \mathbf{x} := \text{getLeft}(E) \rrbracket_{\mathbb{T}} &\stackrel{\text{def}}{=} \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times \mathbf{a}[t_1] \otimes \mathbf{b}[t_2]), \\ (\mathbf{x} \Rightarrow \mathbf{a} * \rho \times \mathbf{a}[t_1] \otimes \mathbf{b}[t_2]) \end{array} \right) \mid \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) = \mathbf{b} \right\} \cup \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times \mathbf{a}[\mathbf{b}[t_1] \otimes t_2]), \\ (\mathbf{x} \Rightarrow \mathbf{nil} * \rho \times \mathbf{a}[\mathbf{b}[t_1] \otimes t_2]) \end{array} \right) \mid \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) = \mathbf{b} \right\} \end{aligned}$$

$$\begin{aligned} \text{AX} \llbracket \mathbf{x} := \text{getRight}(E) \rrbracket_{\mathbb{T}} &\stackrel{\text{def}}{=} \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times \mathbf{b}[t_1] \otimes \mathbf{a}[t_2]), \\ (\mathbf{x} \Rightarrow \mathbf{a} * \rho \times \mathbf{b}[t_1] \otimes \mathbf{a}[t_2]) \end{array} \right) \mid \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) = \mathbf{b} \right\} \cup \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times \mathbf{a}[t_1] \otimes \mathbf{b}[t_2]), \\ (\mathbf{x} \Rightarrow \mathbf{nil} * \rho \times \mathbf{a}[t_1] \otimes \mathbf{b}[t_2]) \end{array} \right) \mid \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) = \mathbf{b} \right\} \end{aligned}$$

$$\begin{aligned} \text{AX} \llbracket \mathbf{x} := \text{getFirst}(E) \rrbracket_{\mathbb{T}} &\stackrel{\text{def}}{=} \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times \mathbf{a}[\mathbf{b}[t_1] \otimes t_2]), \\ (\mathbf{x} \Rightarrow \mathbf{b} * \rho \times \mathbf{a}[\mathbf{b}[t_1] \otimes t_2]) \end{array} \right) \mid \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) = \mathbf{a} \right\} \cup \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times \mathbf{a}[\mathbf{0}]), \\ (\mathbf{x} \Rightarrow \mathbf{nil} * \rho \times \mathbf{a}[\mathbf{0}]) \end{array} \right) \mid \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) = \mathbf{a} \right\} \end{aligned}$$

$$\begin{aligned} \text{AX} \llbracket \mathbf{x} := \text{getLast}(E) \rrbracket_{\mathbb{T}} &\stackrel{\text{def}}{=} \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times \mathbf{a}[t_1] \otimes \mathbf{b}[t_2]), \\ (\mathbf{x} \Rightarrow \mathbf{b} * \rho \times \mathbf{a}[t_1] \otimes \mathbf{b}[t_2]) \end{array} \right) \mid \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) = \mathbf{a} \right\} \cup \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times \mathbf{a}[\mathbf{0}]), \\ (\mathbf{x} \Rightarrow \mathbf{nil} * \rho \times \mathbf{a}[\mathbf{0}]) \end{array} \right) \mid \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) = \mathbf{a} \right\} \end{aligned}$$

$$\begin{aligned} \text{AX} \llbracket \text{newNodeAfter}(E) \rrbracket_{\mathbb{T}} &\stackrel{\text{def}}{=} \\ &\{ (\rho \times \mathbf{a}[t]), (\exists \mathbf{b}. \rho \times \mathbf{a}[t] \times \mathbf{b}[\mathbf{0}]) \mid \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) = \mathbf{a} \} \end{aligned}$$

$$\text{AX} \llbracket \text{deleteTree}(E) \rrbracket_{\mathbb{T}} \stackrel{\text{def}}{=} \{ (\rho \times \mathbf{a}[t]), (\rho \times \mathbf{0}) \mid \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) = \mathbf{a} \}$$

3.3. List Module

This list module $\mathbb{L} \stackrel{\text{def}}{=} (\text{Cmd}_{\mathbb{L}}, \mathcal{A}_{\mathbb{L}}, \text{AX} \llbracket (\cdot) \rrbracket_{\mathbb{L}})$ is an example of a somewhat more exotic abstract module. The module provides an addressable set of lists of unique elements, called a *list store*. Each list can be manipulated independently in a number of ways, new lists can be constructed and existing lists can be deleted.

Definition 3.8 (List Update Commands). The set of *list commands* $\text{Cmd}_{\mathbb{L}}$, ranged over by φ , is defined as follows:

$$\begin{aligned} \varphi ::= & \mathbf{x} := E.\text{getHead}() \mid \mathbf{x} := E.\text{getTail}() \\ & \mid \mathbf{x} := E_1.\text{getNext}(E_2) \mid \mathbf{x} := E_1.\text{getPrev}(E_2) \\ & \mid \mathbf{x} := E.\text{pop}() \mid E_1.\text{push}(E_2) \\ & \mid E_1.\text{remove}(E_2) \mid E_1.\text{insert}(E_2, E_3) \\ & \mid \mathbf{x} := \text{newList}() \mid \text{deleteList}(E) \end{aligned}$$

where $\mathbf{x} \in \text{Var}$ ranges over variables and $E, E_1, \dots \in \text{Expr}$ range over value expressions.

The intuitive meaning of these commands is as follows:

- $E.\text{getHead}()$ and $E.\text{getTail}()$ retrieve, respectively, the first and last elements (if any) of the list identified by E ;
- $E_1.\text{getNext}(E_2)$ and $E_1.\text{getPrev}(E_2)$ retrieve, respectively, the elements (if any) following and preceding the element E_2 in the list identified by E_1 (if E_2 is not in the list, the behaviour of these commands is undefined);
- $E.\text{pop}()$ retrieves and removes the first element of the list identified by E (if the list is empty, the behaviour is undefined);
- $E_1.\text{push}(E_2)$ adds the element E_2 to the start of the list identified by E_1 ;
- $E_1.\text{remove}(E_2)$ removes the element E_2 from the list identified by E_1 (if E_2 is not in the list, the behaviour is undefined);
- $E_1.\text{insert}(E_2, E_3)$ inserts the element E_3 immediately following E_2 in the list identified by E_1 (if E_2 is not in the list, the behaviour is undefined);
- $\text{newList}()$ creates a new list, initially empty, and returns its address; and
- $\text{deleteList}(E)$ deletes the list identified by E .

We require that elements occur at most once in any given list. Thus getNext , getPrevious and insert are unambiguous and the behaviour of push and insert is undefined if they are used to insert elements that are already present in the list.

The list context algebra of Example 1(c) is not adequate as a data model for the list module, since it only models a single list. Instead, we define a *list-store context algebra*, which models multiple independent lists. List stores are similar to heaps in the sense that they are finite maps from addresses to values, except that now the values have intrinsic structure: they are lists of unique elements.

Each list in the store is a finite sequence of values, each of which occurs only once in the list. Lists may be extended by applying contexts, for instance, $v_1 \cdot - \cdot v_2 \bullet w_1 \cdot w_2 = v_1 \cdot w_1 \cdot w_2 \cdot v_2$. However, lists may also be *completed*, which means they cannot be extended by applying a context. Completed lists are indicated by square brackets, and the result of applying a context is undefined. For example, $v_1 \cdot - \cdot v_2 \bullet [w_1 \cdot w_2]$ is undefined.

It is necessary to deal with complete lists in order to specify a number of the update and lookup commands on lists. For example, getFirst returns the first item in a list. Given the partial list $v_1 \cdot v_2$, it is not clear that v_1 is the first element of the list. Indeed, in the context $w \cdot -$ it is certainly not the first element. However, given the completed list $[v_1 \cdot v_2]$ it is completely certain that v_1 is the first element.

Remark. It is quite possible to define list stores to allow the list to be open at one end but not the other. For example, $[v_1$ would allow additional list elements to be added by context application to the right of v_1 , but not to the left. We have chosen not to take this approach in order to avoid further complicating an already-complex definition.

Recall the definition of lists and list contexts from Example 1(c):

Definition 3.9 (List and Contexts). The set of *lists* Lst , ranged over by l, l_1, \dots and the set of *list contexts* C_{Lst} , ranged over by lc, lc_1, \dots , are defined inductively as follows:

$$\begin{aligned} l & ::= \emptyset \mid v \mid l_1 \cdot l_2 \\ lc & ::= - \mid lc \cdot l \mid l \cdot lc \end{aligned}$$

where values $v \in \text{Val}$ are taken to occur uniquely in each list or list context and \cdot is taken to be associative with identity \emptyset .

We defined list stores and their contexts in terms lists and list contexts.

Definition 3.10 (List Stores and Contexts). The set of *list stores* LStore , ranged over by ls, ls_1, \dots and the set of *list store contexts* C_{LStore} , ranged over by lsc, lsc_1, \dots , are defined inductively as follows:

$$\begin{aligned} ls & ::= \text{emp} \mid a \Rightarrow l \mid a \Rightarrow [l] \mid ls_1 * ls_2 \\ lsc & ::= \text{emp} \mid a \Rightarrow l \mid a \Rightarrow [l] \mid a \Rightarrow lc \mid a \Rightarrow [lc] \mid lsc_1 * lsc_2 \end{aligned}$$

where addresses $a \in \text{Addr}$ are taken to occur uniquely in each list store or list store context and $*$ is taken to be associative and commutative with identity emp .

Since list stores are structured like heaps we need a notion of context application that will allow list stores to be split in the same fashion as heaps, for instance

$$(a_1 \Rightarrow v_1 \cdot v_2 \cdot v_3) * (a_2 \Rightarrow w_1 \cdot v_1) = (a_1 \Rightarrow v_1 \cdot v_2 \cdot v_3) \bullet (a_2 \Rightarrow w_1 \cdot v_1).$$

We also want context application to allow splitting within lists themselves, as was possible with the list context algebra, so, for example

$$a_1 \Rightarrow v_1 \cdot v_2 \cdot v_3 = a_1 \Rightarrow v_1 \cdot - \cdot v_3 \bullet a_1 \Rightarrow v_2.$$

Since such applications can be nested, the associativity requirements on context algebras mean that it must be possible to split within multiple lists at the same time, as in

$$\begin{aligned} a_1 \Rightarrow v_1 \cdot v_2 \cdot v_3 * a_2 \Rightarrow w_1 \cdot v_1 & \\ &= a_1 \Rightarrow v_1 \cdot - \cdot v_3 \bullet (a_1 \Rightarrow v_2 * a_2 \Rightarrow w_1 \cdot v_1) \\ &= a_1 \Rightarrow v_1 \cdot - \cdot v_3 \bullet (a_2 \Rightarrow - \cdot v_1 \bullet (a_1 \Rightarrow v_2 * a_2 \Rightarrow w_1)) \\ &= (a_1 \Rightarrow v_1 \cdot - \cdot v_3 \circ a_2 \Rightarrow - \cdot v_1) \bullet (a_1 \Rightarrow v_2 * a_2 \Rightarrow w_1) \\ &= (a_1 \Rightarrow v_1 \cdot - \cdot v_3 * a_2 \Rightarrow - \cdot v_1) \bullet (a_1 \Rightarrow v_2 * a_2 \Rightarrow w_1). \end{aligned}$$

(In a sense, the context $a_1 \Rightarrow v_1 \cdot - \cdot v_3 * a_2 \Rightarrow - \cdot v_1$ above can be seen as a multi-holed context — after all, it certainly has multiple holes! Yet in applying it to a list store, both of the holes must be filled up at the same time. We may see this as the context having multiple ‘list holes’ but a single ‘list store hole’.)

Completed lists cannot be extended by application, but they can be split, as in

$$a_2 \mapsto [w_1 \cdot w_2] = a_2 \mapsto [w_1 \cdot -] \bullet a_2 \mapsto w_2.$$

All of these properties are embodied in the definition of the application and composition operators.

Definition 3.11 (Application and Composition). The application of list store contexts to list stores $\bullet : \mathbf{C}_{\text{LStore}} \times \text{LStore} \rightarrow \text{LStore}$ is defined inductively by:

$$\begin{aligned} \text{emp} \bullet ls &\stackrel{\text{def}}{=} ls \\ (lsc * a \mapsto l) \bullet ls &\stackrel{\text{def}}{=} (lsc \bullet ls) * a \mapsto l \\ (lsc * a \mapsto [l]) \bullet ls &\stackrel{\text{def}}{=} (lsc \bullet ls) * a \mapsto [l] \\ (lsc * a \mapsto lc) \bullet (ls * a \mapsto l) &\stackrel{\text{def}}{=} (lsc \bullet ls) * a \mapsto lc[l/-] \\ (lsc * a \mapsto [lc]) \bullet (ls * a \mapsto l) &\stackrel{\text{def}}{=} (lsc \bullet ls) * a \mapsto [lc[l/-]] \end{aligned}$$

where $lc[l/-]$ denotes the substitution of l for the context hole in lc . The result of the application is undefined when either the right-hand side is badly formed or no case applies.

The composition of list store contexts $\circ : \mathbf{C}_{\text{LStore}} \times \mathbf{C}_{\text{LStore}} \rightarrow \mathbf{C}_{\text{LStore}}$ is defined inductively by:

$$\begin{aligned} \text{emp} \circ lsc' &\stackrel{\text{def}}{=} lsc' \\ (lsc * a \mapsto l) \circ lsc' &\stackrel{\text{def}}{=} (lsc \circ lsc') * a \mapsto l \\ (lsc * a \mapsto [l]) \circ lsc' &\stackrel{\text{def}}{=} (lsc \circ lsc') * a \mapsto [l] \\ (lsc * a \mapsto lc) \circ lsc' &\stackrel{\text{def}}{=} \begin{cases} (lsc \circ lsc') * a \mapsto lc & \text{if } a \notin \text{dom}(lsc') \\ (lsc \circ lsc'') * a \mapsto lc[l/-] & \text{if } lsc' = lsc'' * a \mapsto l \\ (lsc \circ lsc'') * a \mapsto lc[lc'/-] & \text{if } lsc' = lsc'' * a \mapsto lc' \end{cases} \\ (lsc * a \mapsto [lc]) \circ lsc' &\stackrel{\text{def}}{=} \begin{cases} (lsc \circ lsc') * a \mapsto [lc] & \text{if } a \notin \text{dom}(lsc') \\ (lsc \circ lsc'') * a \mapsto [lc[l/-]] & \text{if } lsc' = lsc'' * a \mapsto l \\ (lsc \circ lsc'') * a \mapsto [lc[lc'/-]] & \text{if } lsc' = lsc'' * a \mapsto lc'. \end{cases} \end{aligned}$$

Again, the result of the composition is undefined when either the right-hand side is badly formed or no case applies.

Definition 3.12 (List-Store Context Algebra). The *list-store context algebra*:

$$\mathcal{A}_{\mathbb{L}} \stackrel{\text{def}}{=} (\mathbf{C}_{\text{LStore}}, \text{LStore}, \circ, \bullet, \{\text{emp}\}, \{\text{emp}\})$$

is given by the above definitions.

Definition 3.13 (List Axiomatisation).

The *list axiomatisation*, $\text{AX} \llbracket (\cdot) \rrbracket_{\mathbb{L}} : \text{Cmd}_{\mathbb{L}} \rightarrow \mathcal{P}(\mathcal{P}(\text{Scope} \times \text{LStore}) \times \mathcal{P}(\text{Scope} \times \text{LStore}))$

is defined as follows:

$$\begin{aligned} \text{AX} \llbracket \mathbf{x} := E.\text{getHead}() \rrbracket_{\mathbb{L}} &\stackrel{\text{def}}{=} \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times a \mapsto [w \cdot l]), \\ (\mathbf{x} \Rightarrow w * \rho \times a \mapsto [w \cdot l]) \end{array} \right) \mid a = \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) \right\} \cup \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times a \mapsto [\emptyset]), \\ (\mathbf{x} \Rightarrow \mathbf{nil} * \rho \times a \mapsto [\emptyset]) \end{array} \right) \mid a = \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) \right\} \end{aligned}$$

$$\begin{aligned} \text{AX} \llbracket \mathbf{x} := E.\text{getTail}() \rrbracket_{\mathbb{L}} &\stackrel{\text{def}}{=} \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times a \mapsto [l \cdot w]), \\ (\mathbf{x} \Rightarrow w * \rho \times a \mapsto [l \cdot w]) \end{array} \right) \mid a = \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) \right\} \cup \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times a \mapsto [\emptyset]), \\ (\mathbf{x} \Rightarrow \mathbf{nil} * \rho \times a \mapsto [\emptyset]) \end{array} \right) \mid a = \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) \right\} \end{aligned}$$

$$\begin{aligned} \text{AX} \llbracket \mathbf{x} := E_1.\text{getNext}(E_2) \rrbracket_{\mathbb{L}} &\stackrel{\text{def}}{=} \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times a \mapsto w \cdot u), \\ (\mathbf{x} \Rightarrow u * \rho \times a \mapsto w \cdot u) \end{array} \right) \mid \begin{array}{l} a = \mathcal{E} \llbracket E_1 \rrbracket (\mathbf{x} \Rightarrow v * \rho) \text{ and} \\ w = \mathcal{E} \llbracket E_2 \rrbracket (\mathbf{x} \Rightarrow v * \rho) \end{array} \right\} \cup \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times a \mapsto [l \cdot w]), \\ (\mathbf{x} \Rightarrow \mathbf{nil} * \rho \times a \mapsto [l \cdot w]) \end{array} \right) \mid \begin{array}{l} a = \mathcal{E} \llbracket E_1 \rrbracket (\mathbf{x} \Rightarrow v * \rho) \text{ and} \\ w = \mathcal{E} \llbracket E_2 \rrbracket (\mathbf{x} \Rightarrow v * \rho) \end{array} \right\} \end{aligned}$$

$$\begin{aligned} \text{AX} \llbracket \mathbf{x} := E_1.\text{getPrev}(E_2) \rrbracket_{\mathbb{L}} &\stackrel{\text{def}}{=} \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times a \mapsto u \cdot w), \\ (\mathbf{x} \Rightarrow u * \rho \times a \mapsto u \cdot w) \end{array} \right) \mid \begin{array}{l} a = \mathcal{E} \llbracket E_1 \rrbracket (\mathbf{x} \Rightarrow v * \rho) \text{ and} \\ w = \mathcal{E} \llbracket E_2 \rrbracket (\mathbf{x} \Rightarrow v * \rho) \end{array} \right\} \cup \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times a \mapsto [w \cdot l]), \\ (\mathbf{x} \Rightarrow \mathbf{nil} * \rho \times a \mapsto [w \cdot l]) \end{array} \right) \mid \begin{array}{l} a = \mathcal{E} \llbracket E_1 \rrbracket (\mathbf{x} \Rightarrow v * \rho) \text{ and} \\ w = \mathcal{E} \llbracket E_2 \rrbracket (\mathbf{x} \Rightarrow v * \rho) \end{array} \right\} \end{aligned}$$

$$\begin{aligned} \text{AX} \llbracket \mathbf{x} := E.\text{pop}() \rrbracket_{\mathbb{L}} &\stackrel{\text{def}}{=} \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times a \mapsto [w \cdot l]), \\ (\mathbf{x} \Rightarrow w * \rho \times a \mapsto [l]) \end{array} \right) \mid a = \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) \right\} \end{aligned}$$

$$\begin{aligned} \text{AX} \llbracket E_1.\text{push}(E_2) \rrbracket_{\mathbb{L}} &\stackrel{\text{def}}{=} \\ &\left\{ ((\rho \times a \mapsto [l] \wedge v \notin l), (\rho \times a \mapsto [v \cdot l])) \mid \begin{array}{l} a = \mathcal{E} \llbracket E_1 \rrbracket \rho \text{ and} \\ v = \mathcal{E} \llbracket E_2 \rrbracket \rho \end{array} \right\} \end{aligned}$$

$$\begin{aligned} \text{AX} \llbracket E_1.\text{remove}(E_2) \rrbracket_{\mathbb{L}} &\stackrel{\text{def}}{=} \\ &\left\{ ((\rho \times a \mapsto v), (\rho \times a \mapsto \emptyset)) \mid a = \mathcal{E} \llbracket E_1 \rrbracket \rho \text{ and } v = \mathcal{E} \llbracket E_2 \rrbracket \rho \right\} \end{aligned}$$

$$\begin{aligned} \text{AX} \llbracket E_1.\text{insert}(E_2, E_3) \rrbracket_{\mathbb{L}} &\stackrel{\text{def}}{=} \\ &\left\{ \left(\begin{array}{l} (\rho \times a \Rightarrow [l_1 \cdot v \cdot b_2] \wedge w \notin l_1 \cdot v \cdot b_2), \\ (\rho \times a \Rightarrow [l_1 \cdot v \cdot w \cdot b_2]) \end{array} \right) \mid \begin{array}{l} a = \mathcal{E} \llbracket E_1 \rrbracket \rho \text{ and} \\ v = \mathcal{E} \llbracket E_2 \rrbracket \rho \text{ and} \\ w = \mathcal{E} \llbracket E_3 \rrbracket \rho \end{array} \right\} \\ \text{AX} \llbracket \mathbf{x} := \text{newList}() \rrbracket_{\mathbb{L}} &\stackrel{\text{def}}{=} \{ (\mathbf{x} \Rightarrow - \times \text{emp}), (\exists a. \mathbf{x} \Rightarrow a \times a \Rightarrow [\emptyset]) \} \\ \text{AX} \llbracket \text{deleteList}(E) \rrbracket_{\mathbb{L}} &\stackrel{\text{def}}{=} \{ (\rho \times a \Rightarrow [l]), (\rho \times \text{emp}) \mid a = \mathcal{E} \llbracket E \rrbracket \rho \} \end{aligned}$$

3.4. Combining Abstract Modules

It is often useful when programming to be able to make use of several modules. Just as it is natural to combine context algebras, as in Example 1(e), it is also natural to have a mechanism for combining abstract modules. The most intuitive approach is to allow arbitrary interleavings of their commands, whilst interpreting these commands over the product of their data-store context algebras. If the modules want to share any information, this must be done through the common variable store.

Definition 3.14 (Abstract Module Combination). Given abstract modules $\mathbb{A}_1 = (\text{Cmd}_{\mathbb{A}_1}, \mathcal{A}_{\mathbb{A}_1}, \text{AX} \llbracket (\cdot) \rrbracket_{\mathbb{A}_1})$ and $\mathbb{A}_2 = (\text{Cmd}_{\mathbb{A}_2}, \mathcal{A}_{\mathbb{A}_2}, \text{AX} \llbracket (\cdot) \rrbracket_{\mathbb{A}_2})$, their combination

$$\mathbb{A}_1 + \mathbb{A}_2 \stackrel{\text{def}}{=} (\text{Cmd}_{\mathbb{A}_1} \oplus \text{Cmd}_{\mathbb{A}_2}, \mathcal{A}_{\mathbb{A}_1} \times \mathcal{A}_{\mathbb{A}_2}, \text{AX} \llbracket (\cdot) \rrbracket_{\mathbb{A}_1 + \mathbb{A}_2})$$

is an abstract module, where

- $\text{Cmd}_{\mathbb{A}_1} \oplus \text{Cmd}_{\mathbb{A}_2} \stackrel{\text{def}}{=} (\text{Cmd}_{\mathbb{A}_1} \times \{1\}) \cup (\text{Cmd}_{\mathbb{A}_2} \times \{2\})$ is the discriminated union of the command sets,
- $\mathcal{A}_{\mathbb{A}_1} \times \mathcal{A}_{\mathbb{A}_2}$ is the direct product of the context algebras, and
- $\text{AX} \llbracket (\cdot) \rrbracket_{\mathbb{A}_1 + \mathbb{A}_2} : \text{Cmd}_{\mathbb{A}_1} \oplus \text{Cmd}_{\mathbb{A}_2} \rightarrow \mathcal{P}(\mathcal{P}(\text{Scope} \times \text{D}_{\mathbb{A}_1} \times \text{D}_{\mathbb{A}_2}) \times \mathcal{P}(\text{Scope} \times \text{D}_{\mathbb{A}_1} \times \text{D}_{\mathbb{A}_2}))$ is defined as

$$\begin{aligned} \text{AX} \llbracket (\varphi, 1) \rrbracket_{\mathbb{A}_1 + \mathbb{A}_2} &\stackrel{\text{def}}{=} \{ (\pi_1(P), \pi_1(Q) \mid (P, Q) \in \text{AX} \llbracket \varphi \rrbracket_{\mathbb{A}_1} \} \\ \text{AX} \llbracket (\varphi, 2) \rrbracket_{\mathbb{A}_1 + \mathbb{A}_2} &\stackrel{\text{def}}{=} \{ (\pi_2(P), \pi_2(Q) \mid (P, Q) \in \text{AX} \llbracket \varphi \rrbracket_{\mathbb{A}_2} \} \end{aligned}$$

where

$$\begin{aligned} \pi_1(P) &\stackrel{\text{def}}{=} \{ (\rho, \chi_1, o_2) \mid (\rho, \chi_1) \in P \text{ and } o_2 \in \mathbf{0}_2 \} \\ \pi_2(P) &\stackrel{\text{def}}{=} \{ (\rho, o_1, \chi_2) \mid (\rho, \chi_2) \in P \text{ and } o_1 \in \mathbf{0}_1 \}. \end{aligned}$$

Notation.

When the command sets $\text{Cmd}_{\mathbb{A}_1}$ and $\text{Cmd}_{\mathbb{A}_2}$ are disjoint, we will drop the tags when referring to the commands in the combined abstract module. When the tags are necessary, we will indicate them with an appropriately place subscript.

In §5.5 we shall use a combination of the heap module \mathbb{H} and the list module \mathbb{L} to provide an implementation of the tree module \mathbb{T} .

4. Module Translations

We now define what it means to correctly implement one module in terms of another, using *module translations*.

Definition 4.1 (Module Translation). A module translation $\tau : \mathbb{A} \rightarrow \mathbb{B}$ from abstract module \mathbb{A} to abstract module \mathbb{B} comprises:

- an *abstraction relation* $\alpha_\tau \subseteq D_{\mathbb{B}} \times D_{\mathbb{A}}$, and
- a *substitutive implementation function* $\llbracket (\cdot) \rrbracket_\tau : \mathcal{L}_{\mathbb{A}} \rightarrow \mathcal{L}_{\mathbb{B}}$ which uniformly substitutes each basic command of $\text{Cmd}_{\mathbb{A}}$ with a call to a procedure written in $\mathcal{L}_{\mathbb{B}}$.

Notation.

The abstraction relation α_τ is lifted to a *predicate translation* $\llbracket (\cdot) \rrbracket_\tau : \mathcal{P}(\text{State}_{\mathbb{A}}) \rightarrow \mathcal{P}(\text{State}_{\mathbb{B}})$ as follows:

$$\llbracket P \rrbracket = \{(\rho, \chi_{\mathbb{B}}) \mid \text{there exists } \chi_{\mathbb{A}} \text{ s.t. } (\rho, \chi_{\mathbb{A}}) \in P \text{ and } \chi_{\mathbb{B}} \alpha_\tau \chi_{\mathbb{A}}\}.$$

When the translation τ is implicit from context, the subscripts on the abstraction relation, implementation function and predicate translation may be dropped.

In the context of a translation $\tau : \mathbb{A} \rightarrow \mathbb{B}$, \mathbb{A} is called the *abstract* or *high-level* module and \mathbb{B} is called the *concrete* or *low-level* module. (Of course, there is no reason why a module should not be abstract with respect to one translation and concrete with respect to another, or even both the abstract and concrete module with respect to a single translation.)

Definition 4.2 (Sound Module Translation). A module translation $\tau : \mathbb{A} \rightarrow \mathbb{B}$ is said to be sound, if for all $P, Q \in \mathcal{P}(\text{State}_{\mathbb{A}})$ and $C \in \mathcal{L}_{\mathbb{A}}$,

$$\vdash_{\mathbb{A}} \{P\} C \{Q\} \quad \Longrightarrow \quad \vdash_{\mathbb{B}} \{\llbracket P \rrbracket\} \llbracket C \rrbracket \{\llbracket Q \rrbracket\}.$$

Intuitively, a sound module translation appears to be a reasonable correctness condition for a module implementation: everything that can be proved about the abstract module also holds for its implementation. There are, however, a few caveats.

Firstly, since we have elected to work with partial correctness Hoare triples, it is acceptable for an implementation to simply loop forever. If termination guarantees are required, they could either be made separately or a logic based on total correctness could be used. We have chosen to work with partial correctness for simplicity and on the basis that partial correctness is generally used in the separation logic and context logic literature (Ishtiaq and O’Hearn, 2001; Reynolds, 2002; Calcagno et al., 2005).

Secondly, it is possible for the abstraction relation to lose information. For instance, if all predicates were unsatisfiable under translation then it would be possible to soundly implement every abstract command with `skip`; such an implementation is useless. One way of mitigating this would be to consider a set of *initial predicates* that must be satisfiable under translation. A triple whose precondition is such an initial predicate is then meaningful under translation, since it does not hold vacuously. A more stringent approach would be to require the abstraction relation α_τ to be surjective, and therefore

every satisfiable predicate to be satisfiable under translation. This condition is, however, not met by one of the natural implementations considered here.

In this section, we present two techniques for constructing sound module translations: locality-preserving translations and locality-breaking translations. *Locality-preserving translations*, discussed in §5, closely preserve the structure of the abstract module’s context algebra through the translation, which leads to an elegant inductive proof transformation from the abstract to the concrete. In particular, context application at the abstract level corresponds to context application at the concrete level, and so the abstract frame rule is transformed to the concrete frame rule. *Locality-breaking translations*, discussed in §6, on the other hand do not necessarily preserve the structure of the abstract module’s context algebra. Even so, certain properties of the proof theory can be used to simplify the problem of establishing the soundness of such a translation.

4.1. Modularity

An important property of module translations is that they are composable: given translations $\tau_1 : \mathbb{A}_1 \rightarrow \mathbb{A}_2$ and $\tau_2 : \mathbb{A}_2 \rightarrow \mathbb{A}_3$, the translation $\tau_2 \circ \tau_1 : \mathbb{A}_1 \rightarrow \mathbb{A}_3$ can be defined in the natural fashion. If its constituent translations are sound then so is the composition. Therefore, it is possible to construct module translations stepwise.

A translation $\tau : \mathbb{A}_1 \rightarrow \mathbb{A}_2$ can be naturally lifted to a translation $\tau + \mathbb{B} : \mathbb{A}_1 + \mathbb{B} \rightarrow \mathbb{A}_2 + \mathbb{B}$, for any module \mathbb{B} . If τ is a sound translation, we might also expect $\tau + \mathbb{B}$ to be sound. However, it is not obvious that this is the case in general. The techniques for constructing sound translations in this paper do, however, admit such a lifting. This is because they transform high-level proofs to low-level proofs in a fashion that preserves any additional module component. Being able to combine translations for independent modules means that these techniques are modular.

5. Locality-Preserving Translations

Sometimes, there is a close correspondence between the locality exhibited by a high-level module and the locality exhibited by the low-level implementation of the module. In this section, we expand on this intuition and formalise the concept of a *locality-preserving translation*. In §5.1, we establish that locality-preserving translations give sound module translations, and in §5.4, §5.5 and §5.6 we give locality-preserving translations $\tau_1 : \mathbb{T} \rightarrow \mathbb{H}$, $\tau_2 : \mathbb{T} \rightarrow \mathbb{H} + \mathbb{L}$ and $\tau_3 : \mathbb{H} + \mathbb{H} \rightarrow \mathbb{H}$.

So what exactly does it mean for there to be a correspondence between locality at the high level and locality at the low level? Consider Fig. 3, which depicts a typical tree from the tree module \mathbb{T} (a), together with possible representations of that tree in the heap module \mathbb{H} (b), and in the combined heap-and-list module $\mathbb{H} + \mathbb{L}$ (c). In Fig. 3(b), each tree node is represented by a memory block comprising four pointer fields (depicted by a circle with outgoing arrows) which record the addresses of the memory blocks representing the left sibling, parent, right sibling and first child. Where there is no such node (for example, when a node has no children) the pointer field holds the *nil* value (depicted by the absence of an arrow). In Fig. 3(c), each tree node is represented by a




Fig. 3. An abstract tree from \mathbb{T} (a), and representations of the tree in \mathbb{H} (b), and in $\mathbb{H} + \mathbb{L}$ (c).

memory block that comprises a pointer to the node’s parent and a pointer to the child list of the node. The child list (depicted by a box with dots for each value in the list) is a list of pointers to the node’s children.

What these examples exhibit is a simple inductive transformation from the abstract data structure to its concrete representations. This suggests the possibility of a simple inductive transformation from high-level *proofs* to low-level proofs: in particular, it should be possible to transform high-level frames into low-level frames. Such a transformation may well be said to preserve locality.

Essential to constructing such a transformation is that the concrete representation of abstract data preserves the abstract application structure. In particular, if we split the abstract data structure into a context and subdata, its concrete representation can also be split into the representations of this context and subdata. In Fig. 3, the dashed line indicates such a splitting of the abstract tree, and the corresponding splittings of the two representations. In general, for all contexts c and abstract data structures χ , we would like their intuitive representations, denoted $\langle\langle c \rangle\rangle$ and $\langle\langle \chi \rangle\rangle$, to satisfy the *application preservation property* $\langle\langle c \bullet \chi \rangle\rangle \equiv \langle\langle c \rangle\rangle \bullet \langle\langle \chi \rangle\rangle$.

The reality is, however, more subtle than this intuition, since the concrete representations of context and data must mesh together correctly. In the example representations we have considered, this means that the pointers between the context and subdata must link up correctly. Thus, while an abstract context is agnostic to the data that is placed in its hole, the representation of the context needs to “know” some information about the representation of the data, and vice-versa. Thus the representations of context and data need to be parametrised by *interfaces* which record this “knowledge” that the context and data need about each other. Specifically, the representation of abstract data structure χ is given by the concrete predicate $\langle\langle \chi \rangle\rangle^I$, and the representation of abstract context c is given by the concrete predicate $\langle\langle c \rangle\rangle_{I'}$. The context representation requires two interfaces: the I' between it and the data in its hole and the I between it and its own surrounding context. Taking interfaces into account, the *application preservation property* is now, essentially,

$$\langle\langle c \bullet \chi \rangle\rangle^I \equiv \exists I'. \langle\langle c \rangle\rangle_{I'}^I \bullet \langle\langle \chi \rangle\rangle^{I'}$$

The interface is divided into two parts: the knowledge that a context representation needs about the data to be put in its hole is called the *in-interface*, while the knowledge that a data representation needs about its surrounding context is called the *out-interface*. In Fig. 3 the dashed lines depict a splitting of the abstract tree, and its representations, into a context and a subtree. In Fig. 3(b), the in-interface consists of the addresses of the left- and right-most nodes at the top level of the subtree, while the out-interface consists of the addresses of the nodes immediately left, above and to the right of the context hole. In Fig. 3(c), the in-interface consists of the addresses of all of the nodes at the top level of the subtree, while the out-interface consists of the address of the parent node of the context hole. Typically, as in these examples, the interfaces are a collection of pointers, which are identified in the diagrams by the arrows that cross the dividing line between context and data. The direction of the arrow determines whether it belongs to the in- or out-interface.

Having established an application-preserving representation of data, we aim to transform high-level proofs about an abstract program into corresponding low-level proofs about its implementation, by simply replacing the high-level predicates with their low-level representations. Most of the proof rules should transform simply to their low-level counterparts; application preservation should allow us to transform the FRAME rule; and to deal with AXIOM we simply need to prove that the implementations of the basic commands satisfy the low-level representations of their specifications. However, consider the operation of disposing the subtree indicated in Fig. 3. At the abstract level, it is clear that the resource required to run the command is just the subtree that is to be deleted. This is reflected in the axiom:

$$\text{Ax } \llbracket \text{deleteTree}(E) \rrbracket_{\mathbb{T}} \stackrel{\text{def}}{=} \{ ((\rho \times \mathbf{a}[t]), (\rho \times \mathbf{0})) \mid \mathcal{E} \llbracket E \rrbracket (\mathbf{x} \Rightarrow v * \rho) = \mathbf{a} \}$$

Yet in both implementations, something more than the representation of the subtree is required: for the heap implementation in Fig. 3(b), the pointers from the context into the deleted subtree must be updated; for the heap-and-list implementation in Fig. 3(c), the pointers to the subtree’s top-level nodes must be removed from their parent’s child list. In both cases, the low-level footprint of dispose is larger than the intuitive representation. The axiom for dispose cannot therefore simply be translated.

We could decide that our chosen representation is unsuitable. In our two tree implementations however, the representations seem pretty intuitive. Instead, we demonstrate that it is possible to repair our existing approach. We do this by introducing the concept of *crust*, a predicate that corresponds to the minimal extra resource (taken from the surrounding context) that is required by the command implementations. In Fig. 3, the crust for each subtree representation is depicted by the shaded portions of the context representation. For the heap implementation in Fig. 3(b), this is the nodes (in the example, the parent and right sibling) that have pointers into the subtree; for the heap-and-list implementation in Fig. 3(c), this is the parent node, its child list and its other children. In general, the crust is represented by the predicate \mathfrak{m}_I^F , parametrised by an interface I and some additional parameter F that together fully determine it. (Since the crust corresponds to some portion of the surrounding context, some information about that context is contained within the crust; F is the part of this information that is not

already provided by I . In the heap implementation of trees, for example, F provides all the values of pointers in the nodes belonging to the crust which are not in I .)

We must now ensure that the *axiom correctness property* holds with this extra crust. This means we must check that the implementations of the basic commands satisfy the representations of their specifications, with the addition of crust. Loosely:

$$\vdash \{ \exists in. \mathfrak{m}_{in,out}^F \bullet \langle\langle P \rangle\rangle^{in,out} \} \llbracket \varphi \rrbracket \{ \exists in. \mathfrak{m}_{in,out}^F \bullet \langle\langle Q \rangle\rangle^{in,out} \}$$

for every $(P, Q) \in \text{Ax} \llbracket \varphi \rrbracket$.

With crust now added to the representation of data, the application preservation property no longer immediately provides a way to translate high-level frames into low-level frames. If we simply took the low-level frame to be the representation of the high-level frame plus its outer crust, this would duplicate the data's crust. This *inner crust* must therefore be removed from the frame before it is applied. For example, consider again disposing the subtree indicated in Fig. 3. In the heap implementation of Fig. 3(b), the command operates on the representation of the subtree (as indicated by the dashed line) plus the crust (as indicated by the shaded nodes). If the context is added by frame at the high-level, what is added at the low-level is the representation of the context (and its outer crust) *minus* the inner crust (the shaded nodes). The possibility of removing the inner crust in this way is established by the *crust inclusion property*, which (in part) states that the representation of a context together with its outer crust contains the inner crust. Loosely:

$$\left(\exists in'. \mathfrak{m}_{in',out'}^F \circ \langle\langle c \rangle\rangle_I^{in',out'} \right) \equiv K \circ \mathfrak{m}_I^{F'}.$$

This crust inclusion property enables us to transform a high-level frame into a low-level frame, and hence construct a *locality-preserving* module translation. We introduce crust to deal with the fact that the footprint of the low-level operations overlaps with the representation of the context. Hence we say that such a module translation provides a *fiction of disjointness*.

Having fleshed out the intuition behind locality-preserving translations, we now introduce their formal definition. We first define the concept of *pre-locality-preserving translations*, which have the appropriate form, and then restrict locality-preserving translations to being those that exhibit the properties of application preservation, crust inclusion and axiom correctness. We prove a general result that locality-preserving translations are sound module translations.

Definition 5.1 (Pre-Localty-Preserving Translation). A *pre-locality-preserving translation* $\tau : \mathbb{A} \rightarrow \mathbb{B}$ comprises:

- a set of *in-interfaces* \mathcal{I}_{in} and a set of *out-interfaces* \mathcal{I}_{out} , whose Cartesian product constitutes the set of *interfaces* $\mathcal{I} = \mathcal{I}_{\text{in}} \times \mathcal{I}_{\text{out}}$;
- a *data representation function* $\langle\langle (\cdot) \rangle\rangle_{(\cdot)}^{(\cdot)} : \mathbb{D}_{\mathbb{A}} \times \mathcal{I} \rightarrow \mathcal{P}(\mathbb{D}_{\mathbb{B}})$;
- a *context representation function* $\langle\langle (\cdot) \rangle\rangle_{(\cdot)}^{(\cdot)} : \mathbb{C}_{\mathbb{A}} \times \mathcal{I} \times \mathcal{I} \rightarrow \mathcal{P}(\mathbb{C}_{\mathbb{B}})$;
- a set of *crust parameters* \mathcal{F} ;
- a crust predicate $\mathfrak{m}_I^F \in \mathcal{P}(\mathbb{C}_{\mathbb{B}})$, parametrised by interface $I \in \mathcal{I}$ and crust parameter $F \in \mathcal{F}$; and

— a *substitutive implementation function* $\llbracket (\cdot) \rrbracket_\tau : \mathcal{L}_\mathbb{A} \rightarrow \mathcal{L}_\mathbb{B}$.

We construct a module translation from a pre-locality-preserving translation by defining the abstraction relation in terms of the data representation function and the crust predicate. For any given $out \in \mathcal{I}_{out}$ and $F \in \mathcal{F}$, the abstraction relation α is defined to be

$$\alpha = \{(\chi_\mathbb{B}, \chi_\mathbb{A}) \mid \chi_\mathbb{B} \in \exists in. \mathbb{M}_{in, out}^F \bullet \langle\langle \chi_\mathbb{A} \rangle\rangle^{in, out}\}.$$

Frequently, there will be a natural choice of out and F , but in general any choice is permissible.

We define an *intermediate predicate translation* $\langle\langle (\cdot) \rangle\rangle^{(\cdot)}$, which is closely related to the predicate translation $\llbracket (\cdot) \rrbracket$, except that it is explicitly parametrised by the choice of out and F .

Definition 5.2 (Intermediate Predicate Translation). Given a pre-locality-preserving translation, the *intermediate predicate translation function*

$$\langle\langle (\cdot) \rangle\rangle^{(\cdot)} : \mathcal{P}(\text{State}_\mathbb{A}) \times (\mathcal{I}_{out} \times \mathcal{F}) \rightarrow \mathcal{P}(\text{State}_\mathbb{B})$$

is defined as:

$$\langle\langle P \rangle\rangle^{out, F} \stackrel{\text{def}}{=} \bigvee_{(\rho, \chi_\mathbb{A}) \in P} \{\rho\} \times (\exists in. \mathbb{M}_{in, out}^F \bullet \langle\langle \chi_\mathbb{A} \rangle\rangle^{in, out}).$$

For a module translation defined by a pre-locality-preserving translation with respect to $out \in \mathcal{I}_{out}$ and $F \in \mathcal{F}$, the predicate translation can be expressed simply in terms of the intermediate predicate translation: $\llbracket P \rrbracket = \langle\langle P \rangle\rangle^{out, F}$.

Pre-locality-preserving translations do not embody the intuition of what it means for a module translation to soundly preserve locality. This is reserved for locality-preserving translations, which require the following three properties.

Property 1 (Application Preservation). Context application is preserved by the representation functions. That is, for all $c \in \mathcal{C}_\mathbb{A}$, $\chi \in \mathcal{D}_\mathbb{A}$ and $I \in \mathcal{I}$,

$$\langle\langle c \bullet_\mathbb{A} \chi \rangle\rangle^I \equiv \exists I'. \langle\langle c \rangle\rangle_{I'}^I \bullet_\mathbb{B} \langle\langle \chi \rangle\rangle^{I'}.$$

Property 2 (Crust Inclusion). For all $out, out' \in \mathcal{I}_{out}$, $F \in \mathcal{F}$ and $c \in \mathcal{C}_\mathbb{A}$, there exist $K \in \mathcal{P}(\mathcal{C}_\mathbb{B})$ and $F' \in \mathcal{F}$ such that, for all $in \in \mathcal{I}_{in}$,

$$\left(\exists in'. \mathbb{M}_{in', out'}^{F'} \circ \langle\langle c \rangle\rangle_{in, out}^{in', out'} \right) \equiv K \circ \mathbb{M}_{in, out}^{F'}.$$

Property 3 (Axiom Correctness). For all $\varphi \in \text{Cmd}_\mathbb{A}$, $(P, Q) \in \text{Ax} \llbracket \varphi \rrbracket_\mathbb{A}$, $out \in \mathcal{I}_{out}$ and $F \in \mathcal{F}$,

$$\vdash_\mathbb{B} \left\{ \langle\langle P \rangle\rangle^{out, F} \right\} \llbracket \varphi \rrbracket \left\{ \langle\langle Q \rangle\rangle^{out, F} \right\}.$$

Definition 5.3 (Locality Preserving Translation). A *locality-preserving translation* is a pre-locality-preserving translation that satisfies Properties 1, 2 and 3.

Remark. The crust inclusion property is stronger than just asserting that inner crust is contained within the context plus its outer crust. This is to allow for the fact that

a command may modify the crust, and we must be able to use this modified crust to reconstitute the representation of the same context. However, the crust is only ever altered with respect to the in-interface, so it is only necessary that K in Property 2 be independent of the choice of in .

Theorem 5.1 (Soundness of Locality-Preserving Translations). A locality-preserving translation is a sound module translation (for any fixed choice of $out \in \mathcal{I}_{out}$ and $F \in \mathcal{F}$).

5.1. Proof of Soundness of Locality-Preserving Translations

In this section, assume that $\tau : \mathbb{A} \rightarrow \mathbb{B}$ is a locality preserving translation. The following proposition is sufficient to establish that τ gives rise to a sound module translation.

Proposition 5.2. For all $out \in \mathcal{I}_{out}$ and $F \in \mathcal{F}$, and for all $P, K \in \mathcal{P}(\text{State}_{\mathbb{A}})$ and $\mathbb{C} \in \mathcal{L}_{\mathbb{A}}$,

$$\Gamma \vdash_{\mathbb{A}} \{P\} \mathbb{C} \{Q\} \implies \llbracket \Gamma \rrbracket \vdash_{\mathbb{B}} \left\{ \llbracket P \rrbracket^{out, F} \right\} \mathbb{C} \left\{ \llbracket Q \rrbracket^{out, F} \right\},$$

where

$$\llbracket \Gamma \rrbracket = \left\{ \mathfrak{f} : \llbracket P' \rrbracket^{out', F'} \rightsquigarrow \llbracket Q' \rrbracket^{out', F'} \mid \begin{array}{l} (\mathfrak{f} : P' \rightsquigarrow Q') \in \Gamma \text{ and} \\ out' \in \mathcal{I}_{out} \text{ and } F' \in \mathcal{F} \end{array} \right\}.$$

Before embarking on the proof of this proposition, two auxiliary lemmata are required. The following lemma gives an alternative characterisation of the crust inclusion property.

Lemma 5.3 (Crust Inclusion II). For all $K \in \mathcal{P}(\mathbb{C}_{\mathbb{A}})$, $in \in \mathcal{I}_{in}$, $out, out' \in \mathcal{I}_{out}$ and $F \in \mathcal{F}$,

$$\begin{aligned} & \left(\exists in'. \mathfrak{M}_{in', out'}^F \circ \llbracket K \rrbracket_{in, out}^{in', out'} \right) \\ & \subseteq \exists F'. \left(\forall in''. \mathfrak{M}_{in'', out}^{F'} \multimap \left(\exists in'. \mathfrak{M}_{in', out'}^F \circ \llbracket K \rrbracket_{in'', out}^{in', out'} \right) \right) \circ \mathfrak{M}_{in, out}^{F'}, \end{aligned}$$

where

$$\llbracket K \rrbracket_I^{I'} = \bigvee_{c \in K} \llbracket c \rrbracket_I^{I'}$$

is the point-wise lift of the context representation function to predicates.

Note that the converse of this property is trivially true, hence the entailment holds in both directions.

Proof. Fix arbitrary $K \in \mathcal{P}(\mathbb{C}_{\mathbb{A}})$, $in \in \mathcal{I}_{in}$, $out, out' \in \mathcal{I}_{out}$ and $F \in \mathcal{F}$. Fix $c' \in \mathbb{C}_{\mathbb{A}}$ with

$$c' \in \left(\exists in'. \mathfrak{M}_{in', out'}^F \circ \llbracket K \rrbracket_{in, out}^{in', out'} \right) \equiv \bigvee_{c \in K} \left(\exists in'. \mathfrak{M}_{in', out'}^F \circ \llbracket c \rrbracket_{in, out}^{in', out'} \right).$$

There exists $c'' \in K$ such that

$$c' \in \left(\exists in'. \mathfrak{M}_{in', out'}^F \circ \llbracket c'' \rrbracket_{in, out}^{in', out'} \right).$$

By the Crust Inclusion Property, there exist $K' \in \mathcal{P}(\mathbb{C}_{\mathbb{B}})$ and $F' \in \mathcal{F}$ such that, for all $in'' \in \mathcal{I}_{in}$,

$$\left(\exists in'. \mathbb{M}_{in',out'}^F \circ \langle\langle c'' \rangle\rangle_{in'',out'}^{in',out'} \right) \equiv K' \circ \mathbb{M}_{in'',out}^{F'}. \quad (1)$$

Hence, $c' \in K' \circ \mathbb{M}_{in,out}^{F'}$, and so there are $c_1 \in K'$ and $c_2 \in \mathbb{M}_{in,out}^{F'}$ with $c' = c_1 \circ c_2$. Fix $in'' \in \mathcal{I}_{in}$ and $c'_2 \in \mathbb{M}_{in'',out}^{F'}$. Since $c_1 \circ c_2 \in K' \circ \mathbb{M}_{in'',out}^{F'}$, it follows by (1) that

$$c_1 \circ c'_2 \in \left(\exists in'. \mathbb{M}_{in',out'}^F \circ \langle\langle c'' \rangle\rangle_{in'',out'}^{in',out'} \right) \subseteq \exists in'. \mathbb{M}_{in',out'}^F \circ \langle\langle K \rangle\rangle_{in'',out'}^{in',out'}.$$

The choice of c'_2 was arbitrary, and so

$$\text{for all } c'_2, c'' \in \mathbb{C}_{\mathbb{B}}, c'_2 \in \mathbb{M}_{in'',out}^{F'} \text{ and } c'' = c_1 \circ c'_2 \implies c'' \in \exists in'. \mathbb{M}_{in',out'}^F \bullet \langle\langle K \rangle\rangle_{in'',out'}^{in',out'}.$$

Hence

$$c_1 \in \mathbb{M}_{in'',out}^{F'} \multimap \left(\exists in'. \mathbb{M}_{in',out'}^F \bullet \langle\langle K \rangle\rangle_{in'',out'}^{in',out'} \right)$$

and since the choice of in'' was arbitrary,

$$c_1 \in \forall in''. \mathbb{M}_{in'',out}^{F'} \multimap \left(\exists in'. \mathbb{M}_{in',out'}^F \bullet \langle\langle K \rangle\rangle_{in'',out'}^{in',out'} \right)$$

Since $c' = c_1 \circ c_2$,

$$c' \in \exists F'. \left(\forall in''. \mathbb{M}_{in'',out}^{F'} \multimap \left(\exists in'. \mathbb{M}_{in',out'}^F \bullet \langle\langle K \rangle\rangle_{in'',out'}^{in',out'} \right) \right) \circ \mathbb{M}_{in,out}^{F'}$$

Since the choice of c' was arbitrary, it follows that

$$\begin{aligned} & \left(\exists in'. \mathbb{M}_{in',out'}^F \circ \langle\langle K \rangle\rangle_{in,out}^{in',out'} \right) \\ & \subseteq \exists F'. \left(\forall in''. \mathbb{M}_{in'',out}^{F'} \multimap \left(\exists in'. \mathbb{M}_{in',out'}^F \circ \langle\langle K \rangle\rangle_{in'',out'}^{in',out'} \right) \right) \circ \mathbb{M}_{in,out}^{F'} \end{aligned}$$

as required. \square

We define the notion of an *intermediate frame translation*, which translates abstract frames (context predicates) to concrete frames. The intermediate frame translation combines the context representation with the outer crust, but removes the inner crust. When the frame as applied at the concrete level, the inner crust will already be present: removing the inner crust from the context is necessary to avoid duplication.

Definition 5.4 (Intermediate Frame Translation). The *intermediate frame translation function*

$$\langle\langle \cdot \rangle\rangle_{(\cdot)}^{(\cdot)} : \mathcal{P}(\mathbb{C}_{\text{State}_{\mathbb{A}}}) \times (\mathcal{I}_{\text{out}} \times \mathcal{F}) \times (\mathcal{I}_{\text{out}} \times \mathcal{F}) \rightarrow \mathcal{P}(\mathbb{C}_{\text{State}_{\mathbb{B}}})$$

is defined as:

$$\langle\langle K \rangle\rangle_{out',F'}^{out,F} = \bigvee_{(\rho, c_{\mathbb{A}}) \in K} \{ \rho \} \times \left(\forall in''. \mathbb{M}_{in'',out'}^F \multimap \left(\exists in. \mathbb{M}_{in,out}^F \circ \langle\langle c_{\mathbb{A}} \rangle\rangle_{in',out'}^{in,out} \right) \right).$$

The fact that the intermediate frame translation produces concrete frames that exactly correspond to abstract frames is critical for dealing with the FRAME rule in the proof of Proposition 5.2. This fact is captured formally in the following lemma.

Lemma 5.4 (Application Preservation II). For all $K \in \mathcal{P}(\mathbb{C}_{\text{State}_{\mathbb{A}}})$, all $P \in \mathcal{P}(\text{State}_{\mathbb{A}})$, $out \in \mathcal{I}_{\text{out}}$ and $F \in \mathcal{F}$,

$$\langle\langle K \bullet_{\mathbb{A}} P \rangle\rangle^{out, F} \equiv \exists out', F'. \langle\langle K \rangle\rangle_{out', F'}^{out, F} \bullet_{\mathbb{B}} \langle\langle P \rangle\rangle^{out', F'}.$$

Proof.

$$\langle\langle K \bullet P \rangle\rangle^{out, F} \equiv \bigvee_{\substack{(\rho', c_{\mathbb{A}}) \in K \\ (\rho, \chi_{\mathbb{A}}) \in P}} \{\rho' * \rho\} \times (\exists in. \mathbb{M}_{in, out}^F \bullet \langle\langle c_{\mathbb{A}} \bullet \chi_{\mathbb{A}} \rangle\rangle^{in, out})$$

(Property 1)

$$\begin{aligned} &\equiv \bigvee_{\substack{(\rho', c_{\mathbb{A}}) \in K \\ (\rho, \chi_{\mathbb{A}}) \in P}} \{\rho' * \rho\} \times (\exists in. \mathbb{M}_{in, out}^F \bullet \exists in', out'. \langle\langle c_{\mathbb{A}} \rangle\rangle_{in', out'}^{in, out} \bullet \langle\langle \chi_{\mathbb{A}} \rangle\rangle^{in', out'}) \\ &\equiv \bigvee_{\substack{(\rho', c_{\mathbb{A}}) \in K \\ (\rho, \chi_{\mathbb{A}}) \in P}} \{\rho' * \rho\} \times (\exists in', out'. \exists in. \mathbb{M}_{in, out}^F \circ \langle\langle c_{\mathbb{A}} \rangle\rangle_{in', out'}^{in, out} \bullet \langle\langle \chi_{\mathbb{A}} \rangle\rangle^{in', out'}) \end{aligned}$$

(Lemma 5.3)

$$\begin{aligned} &\equiv \bigvee_{\substack{(\rho', c_{\mathbb{A}}) \in K \\ (\rho, \chi_{\mathbb{A}}) \in P}} \{\rho' * \rho\} \times \left(\begin{array}{l} \exists in', out'. \exists F'. \\ \left(\forall in''. \mathbb{M}_{in'', out'}^{F'} \multimap \left(\exists in. \mathbb{M}_{in, out}^F \circ \langle\langle c_{\mathbb{A}} \rangle\rangle_{in'', out'}^{in, out} \right) \right) \\ \quad \circ \mathbb{M}_{in', out'}^{F'} \bullet \langle\langle \chi_{\mathbb{A}} \rangle\rangle^{in', out'} \end{array} \right) \\ &\equiv \exists out', F'. \\ &\quad \left(\bigvee_{(\rho', c_{\mathbb{A}}) \in K} \{\rho'\} \times \left(\forall in''. \mathbb{M}_{in'', out'}^{F'} \multimap \left(\exists in. \mathbb{M}_{in, out}^F \circ \langle\langle c_{\mathbb{A}} \rangle\rangle_{in'', out'}^{in, out} \right) \right) \right) \\ &\quad \bullet \left(\bigvee_{(\rho, \chi_{\mathbb{A}}) \in P} \{\rho\} \times \left(\exists in'. \mathbb{M}_{in', out'}^{F'} \bullet \langle\langle \chi_{\mathbb{A}} \rangle\rangle^{in', out'} \right) \right) \\ &\equiv \exists out', F'. \langle\langle K \rangle\rangle_{out', F'}^{out, F} \bullet \langle\langle P \rangle\rangle^{out', F'}. \end{aligned}$$

□

The proof of Proposition 5.2 inductively transforms a proof in \mathbb{A} into a proof in \mathbb{B} .

Proof. The proof is by induction on the structure of the proof of $\vdash_{\mathbb{A}} \{P\} \mathbb{C} \{Q\}$, considering the cases for the last rule applied in the proof. Assume as the inductive hypothesis that the translated premises have proofs in \mathbb{B} . We show how to derive a proof of the translated conclusions from these translated premises. (We omit the procedure specification environment when it plays no role in the derivation.)

Fix arbitrary $out \in \mathcal{I}_{\text{out}}$, $F \in \mathcal{F}$.

AXIOM case:

This case is immediate by the Axiom Correctness Property (Property 3).

FRAME case:

$$\begin{array}{c}
\text{for all } out', F', \{ \langle P \rangle^{out', F'} \} \mathbb{C} \{ \langle Q \rangle^{out', F'} \} \\
\hline
\{ \langle K \rangle_{out', F'}^{out, F} \bullet \langle P \rangle^{out', F'} \} \\
\text{for all } out', F', \mathbb{C} \\
\{ \langle K \rangle_{out', F'}^{out, F} \bullet \langle Q \rangle^{out', F'} \} \\
\hline
\{ \exists out', F'. \langle K \rangle_{out', F'}^{out, F} \bullet \langle P \rangle^{out', F'} \} \\
\mathbb{C} \\
\{ \exists out', F'. \langle K \rangle_{out', F'}^{out, F} \bullet \langle Q \rangle^{out', F'} \} \\
\hline
\{ \langle K \bullet P \rangle^{out, F} \} \mathbb{C} \{ \langle K \bullet Q \rangle^{out, F} \} \quad \text{Lemma 5.4}
\end{array}$$

CONS case:

$$\begin{array}{c}
\frac{P \subseteq P' \quad \{ \langle P' \rangle^{out, F} \} \mathbb{C} \quad \frac{Q' \subseteq Q}{\langle Q' \rangle^{out, F} \subseteq \langle Q \rangle^{out, F}}}{\langle P \rangle^{out, F} \subseteq \langle P' \rangle^{out, F} \quad \{ \langle Q' \rangle^{out, F} \} \mathbb{C} \quad \langle Q' \rangle^{out, F} \subseteq \langle Q \rangle^{out, F}} \quad \text{CONS} \\
\hline
\{ \langle P \rangle^{out, F} \} \mathbb{C} \{ \langle Q \rangle^{out, F} \}
\end{array}$$

DISJ case:

$$\begin{array}{c}
\text{for all } i \in I, \{ \langle P_i \rangle^{out, F} \} \mathbb{C} \{ \langle Q_i \rangle^{out, F} \} \\
\hline
\{ \bigvee_{i \in I} \langle P_i \rangle^{out, F} \} \mathbb{C} \{ \bigvee_{i \in I} \langle Q_i \rangle^{out, F} \} \\
\hline
\{ \langle \bigvee_{i \in I} P_i \rangle^{out, F} \} \mathbb{C} \{ \langle \bigvee_{i \in I} Q_i \rangle^{out, F} \} \quad \text{DISJ}
\end{array}$$

PDEF case:

$$\begin{array}{c}
\text{for all } (\mathbf{f}_i : P \multimap Q) \in \Gamma, \quad \{ \left(\left(\exists \vec{v}. \{ \vec{x}_i \Rightarrow \vec{v} * \vec{r}_i \Rightarrow - \} \right) \times P(\vec{v}) \right) \}^{out', F'} \\
\text{out}' \in \mathcal{I}_{out}, F' \in \mathcal{F}, \quad \llbracket \Gamma', \Gamma \rrbracket \vdash_{\mathbb{B}} \quad \mathbb{C}_i \\
\left\{ \left(\left(\exists \vec{w}. \{ \vec{x}_i \Rightarrow - * \vec{r}_i \Rightarrow \vec{w} \} \right) \times Q(\vec{w}) \right) \right\}^{out', F'} \\
\hline
\text{for all } (\mathbf{f}_i : P \multimap Q) \in \Gamma, \quad \{ \left(\left(\exists \vec{v}. \{ \vec{x}_i \Rightarrow \vec{v} * \vec{r}_i \Rightarrow - \} \right) \times (P(\vec{v}))^{out', F'} \right) \} \\
\text{out}' \in \mathcal{I}_{out}, F' \in \mathcal{F}, \quad \llbracket \Gamma', \Gamma \rrbracket \vdash_{\mathbb{B}} \quad \mathbb{C}_i \\
\left\{ \left(\left(\exists \vec{w}. \{ \vec{x}_i \Rightarrow - * \vec{r}_i \Rightarrow \vec{w} \} \right) \times (Q(\vec{w}))^{out', F'} \right) \right\} \\
\hline
\text{for all } (\mathbf{f}_i : P \multimap Q) \in \llbracket \Gamma \rrbracket, \quad \{ \left(\left(\exists \vec{v}. \{ \vec{x}_i \Rightarrow \vec{v} * \vec{r}_i \Rightarrow - \} \right) \times P(\vec{v}) \right) \} \\
\llbracket \Gamma', \Gamma \rrbracket \vdash_{\mathbb{B}} \quad \mathbb{C}_i \\
\left\{ \left(\left(\exists \vec{w}. \{ \vec{x}_i \Rightarrow - * \vec{r}_i \Rightarrow \vec{w} \} \right) \times Q(\vec{w}) \right) \right\} \\
\hline
(\star)
\end{array}$$

$$\begin{array}{c}
(\star) \quad \llbracket \Gamma', \Gamma \rrbracket \vdash_{\mathbb{B}} \left\{ \langle P \rangle^{out, F} \right\} \mathbb{C} \left\{ \langle Q \rangle^{out, F} \right\} \\
\hline
\text{PDEF} \\
\llbracket \Gamma' \rrbracket \vdash_{\mathbb{B}} \text{procs } \vec{r}_1 := \mathbf{f}_1(\vec{x}_1) \{ \mathbb{C}_1 \}, \dots, \vec{r}_k := \mathbf{f}_k(\vec{x}_k) \{ \mathbb{C}_k \} \text{ in } \mathbb{C} \\
\left\{ \langle P \rangle^{out, F} \right\} \\
\left\{ \langle R \rangle^{out, F} \right\}
\end{array}$$

The two further premises of the PDEF rule, which are not shown in the above derivation, are easily dispatched since $\llbracket (\cdot) \rrbracket$ preserves the procedure names in a procedure specification environment.

The cases for the remaining rules follow by the point-wise and variable-preserving nature of the translation functions. \square

This completes the proof of Theorem 5.1.

5.2. Eliminating Crust

Our theory includes the explicit concept of crust in order to provide additional resource that was necessary for the low-level operations, but not included in the representation of the abstract resource. However, in general it is possible to choose a different representation that avoids the need for crust. In fact, such a choice is embodied in a key element of the soundness proof, namely the intermediate translation functions (Definitions 5.2 and 5.4).

The intermediate translation functions essentially add the outer crust to representations of data structures and contexts while removing the inner crust from representations of contexts. Because of the crust inclusion and application preservation properties of the original representations, this modified representation also preserves application (Lemma 5.4). (Note that the interface set for this modified representation is $\mathcal{I}_{out} \times \mathcal{F}$, which can all be considered as the out part.) This modified representation incorporates all of the resource required at the low level by its very construction, and so no further crust is required.

5.3. Including the Conjunction Rule

If we wish to add the conjunction rule to the locality-preserving theory, we can add a case to the proof of Proposition 5.2. The CONJ rule can be dealt with in the same fashion as the DISJ rule, provided that $\llbracket (\cdot) \rrbracket^{(\cdot)}$ distributes over conjunctions. Together, the following two conditions are sufficient to establish this:

- for all $\chi, \chi' \in \mathbb{D}$ with $\chi \neq \chi'$, and all $I \in \mathcal{I}$, $\langle \chi \rangle^I \wedge \langle \chi' \rangle^I \equiv \emptyset$; and
- for all $out \in \mathcal{I}_{out}$ and $F \in \mathcal{F}$, the predicate $\bigvee \{ \mathbb{m}_{in, out}^F \mid in \in \mathcal{I}_{in} \}$ is precise.

Remark. It is not a coincidence that these conditions are similar to those that prevent a command from behaving angelically given in §2: in both cases, the conditions are constraining the predicate transformers corresponding to the abstraction relation or command to being *conjunctive*.

5.4. Module Translation: $\tau_1 : \mathbb{T} \rightarrow \mathbb{H}$

We now present a locality-preserving translation τ_1 from the tree module \mathbb{T} into the heap module \mathbb{H} . This translation represents each tree node \mathbf{a} as a block of four cells in the heap, $\mathbf{a} \mapsto l, u, d, r$, which contain pointers to the node's left sibling (l), parent (u), first child (d) and right sibling (r). This is the representation illustrated in Fig. 3(b).

As expected, the in-interface consists of the addresses of the left- and right-most nodes at the root level of a tree, while the out-interface consists of the addresses of the tree's parent node and of the nodes immediately adjacent to the tree on the left and right.

Note that for the empty tree \emptyset , the addresses constituting the in-interface are not simply *nil*; rather the “address of the left-most node” should actually be the address of the node immediately adjacent to the tree on the right and the “address of the right-most node” should be the address of the node adjacent on the left. If this were not the case, it would disrupt the continuous list of nodes. On the other hand, if the nodes described in the out-interface do not exist, their addresses will be *nil*.

For contexts, addresses in the inner interface and outer interface can be the same. This is particularly evident in the case of the context hole, $-$, for which the two interfaces must match exactly.

Definition 5.5 ($\tau_1 : \mathbb{T} \rightarrow \mathbb{H}$). The pre-locality-preserving translation $\tau_1 : \mathbb{T} \rightarrow \mathbb{H}$ is constructed as follows:

- the in-interfaces $\mathcal{I}_{\text{in}} = (\text{Addr}_{\text{nil}})^2$ are pairs of addresses;
- the out-interfaces $\mathcal{I}_{\text{out}} = (\text{Addr}_{\text{nil}})^3$ are triples of addresses;
- the data representation function is defined inductively by

$$\begin{aligned} \langle\langle \emptyset \rangle\rangle^{(i,j),(l,u,r)} &\stackrel{\text{def}}{=} \text{emp} \wedge i = r \wedge j = l \\ \langle\langle \mathbf{a}[t] \rangle\rangle^{(i,j),(l,u,r)} &\stackrel{\text{def}}{=} \exists i', j'. \mathbf{a} \mapsto l, u, i', r * \langle\langle t \rangle\rangle^{(i',j'),(\text{nil},\mathbf{a},\text{nil})} \wedge i = j = \mathbf{a} \\ \langle\langle t_1 \otimes t_2 \rangle\rangle^{(i,j),(l,u,r)} &\stackrel{\text{def}}{=} \exists i', j'. \langle\langle t_1 \rangle\rangle^{(i,j'),(l,u,i')} * \langle\langle t_2 \rangle\rangle^{(i',j),(j',u,r)}; \end{aligned}$$

- the context representation function is defined inductively by

$$\begin{aligned} \langle\langle - \rangle\rangle_{I'}^I &\stackrel{\text{def}}{=} \text{emp} \wedge I = I' \\ \langle\langle \mathbf{a}[c] \rangle\rangle_{I'}^{(i,j),(l,u,r)} &\stackrel{\text{def}}{=} \exists i', j'. \mathbf{a} \mapsto l, u, i', r * \langle\langle c \rangle\rangle_{I'}^{(i',j'),(\text{nil},\mathbf{a},\text{nil})} \wedge i = j = \mathbf{a} \\ \langle\langle t \otimes c \rangle\rangle_{I'}^{(i,j),(l,u,r)} &\stackrel{\text{def}}{=} \exists i', j'. \langle\langle t \rangle\rangle^{(i,j'),(l,u,i')} * \langle\langle c \rangle\rangle_{I'}^{(i',j),(j',u,r)} \\ \langle\langle c \otimes t \rangle\rangle_{I'}^{(i,j),(l,u,r)} &\stackrel{\text{def}}{=} \exists i', j'. \langle\langle c \rangle\rangle_{I'}^{(i,j'),(l,u,i')} * \langle\langle t \rangle\rangle^{(i',j),(j',u,r)}; \end{aligned}$$

- the crust parameters $\mathcal{F} = (\text{Addr}_{\text{nil}})^7$ are 7-tuples of addresses;
- the crust predicate is defined as

$$\begin{aligned} \overline{\mathfrak{M}}_{(i,j),(l,u,r)}^{\vec{f}} &\stackrel{\text{def}}{=} (l \mapsto f_1, u, f_2, i \vee (l = \text{nil} \wedge (u \mapsto f_3, f_4, i, f_5 \vee (u = \text{nil} \wedge \text{emp})))) \\ &\quad * (r \mapsto j, u, f_6, f_7 \vee (r = \text{nil} \wedge \text{emp})); \text{ and} \end{aligned}$$

- the substitutive implementation function is given by replacing each tree-module command with a call to the correspondingly-named procedure given in Fig. 4.

$$\begin{aligned}
E.\text{left} &\stackrel{\text{def}}{=} E \\
E.\text{up} &\stackrel{\text{def}}{=} E + 1 \\
E.\text{down} &\stackrel{\text{def}}{=} E + 2 \\
E.\text{right} &\stackrel{\text{def}}{=} E + 3 \\
n := \text{newNode}() &\stackrel{\text{def}}{=} n := \text{alloc}(4) \\
\text{disposeNode}(E) &\stackrel{\text{def}}{=} \text{dispose}(E, 4)
\end{aligned}$$

```

m := getUp(n) {
  m := [n.up]
}

m := getLast(n) {
  local x in
    m := [n.down];
    if m ≠ nil then
      x := [m.right];
      while x ≠ nil do
        m := x;
        x := [m.right]
}

newNodeAfter(n) {
  local x, y, z in
    y := [n.right];
    z := [n.up];
    z := newNode();
    [x.left] := n;
    [x.up] := z;
    [x.down] := nil;
    [x.right] := y;
    [n.right] := x;
    if y ≠ nil then
      [y.left] := x
}

m := getFirst(n) {
  m := [n.down]
}

m := getRight(n) {
  m := [n.right]
}

m := getLeft(n) {
  m := [n.left]
}

deleteTree(n) {
  local x, y, z, w in
    x := [n.right];
    y := [n.left];
    z := [n.up];
    w := [n.down];
    call disposeForest(w);
    if x ≠ nil then
      [x.left] := y ;
    if y ≠ nil then
      [y.right] := x
    else
      if z ≠ nil then
        [z.down] := x
}

disposeForest(n) {
  local r, d in
    if n ≠ nil then
      r := [n.right];
      call disposeForest(r);
      d := [n.down];
      call disposeForest(d);
      disposeNode(n)
}

```

Fig. 4. Procedures for the heap-based implementation of trees

Theorem 5.5 (Soundness of τ_1). The pre-locality-preserving translation τ_1 is a locality-preserving translation.

5.4.1. Soundness of $\tau_1 : \mathbb{T} \rightarrow \mathbb{H}$

Lemma 5.6 (τ_1 Application Preservation). For all $c \in \mathbb{C}_{\text{Tree}}$, $t \in \text{Tree}$ and $I \in \mathcal{I}$,

$$\langle\langle c \bullet t \rangle\rangle^I \equiv \exists I'. \langle\langle c \rangle\rangle_{I'}^I \bullet \langle\langle t \rangle\rangle^{I'}.$$

Proof. The proof is by induction on the structure of the context c , assuming some fixed $t \in \text{Tree}$.

$c = -$ case:

$$\begin{aligned} \exists I'. \langle\langle - \rangle\rangle_{I'}^I \bullet \langle\langle t \rangle\rangle^{I'} &\equiv \exists I'. I = I' \wedge \langle\langle t \rangle\rangle^{I'} \\ &\equiv \langle\langle t \rangle\rangle^I \\ &\equiv \langle\langle - \bullet t \rangle\rangle^I. \end{aligned}$$

$c = \mathbf{a}[c']$ case:

$$\begin{aligned} \exists I'. \langle\langle \mathbf{a}[c'] \rangle\rangle_{I'}^{(i,j),(l,u,r)} \bullet \langle\langle t \rangle\rangle^{I'} &\equiv \exists I'. \left(\exists i', j'. \mathbf{a} \mapsto l, u, i', r * \langle\langle c' \rangle\rangle_{I'}^{(i',j'),(nil,\mathbf{a},nil)} \wedge i = j = \mathbf{a} \right) \bullet \langle\langle t \rangle\rangle^{I'} \\ &\equiv \exists i', j'. \mathbf{a} \mapsto l, u, i', r * \left(\exists I'. \langle\langle c' \rangle\rangle_{I'}^{(i',k'),(nil,\mathbf{a},nil)} \bullet \langle\langle t \rangle\rangle^{I'} \right) \wedge i = j = \mathbf{a} \\ &\equiv \exists i', j'. \mathbf{a} \mapsto l, u, i', r * \langle\langle c' \bullet t \rangle\rangle_{I'}^{(i',k'),(nil,\mathbf{a},nil)} \wedge i = j = \mathbf{a} \\ &\equiv \langle\langle \mathbf{a}[c' \bullet t] \rangle\rangle_{I'}^{(i,j),(l,u,r)} \\ &\equiv \langle\langle \mathbf{a}[c'] \bullet t \rangle\rangle_{I'}^{(i,j),(l,u,r)}. \end{aligned}$$

$c = c' \otimes t'$ case:

$$\begin{aligned} \exists I'. \langle\langle c' \otimes t' \rangle\rangle_{I'}^{(i,j),(l,u,r)} \bullet \langle\langle t \rangle\rangle^{I'} &\equiv \exists I'. \left(\exists i', j'. \langle\langle c' \rangle\rangle_{I'}^{(i,j'),(l,u,i')} * \langle\langle t' \rangle\rangle_{I'}^{(i',j),(j',u,r)} \right) \bullet \langle\langle t \rangle\rangle^{I'} \\ &\equiv \exists i', j'. \left(\exists I'. \langle\langle c' \rangle\rangle_{I'}^{(i,j'),(l,u,i')} \bullet \langle\langle t \rangle\rangle^{I'} \right) * \langle\langle t' \rangle\rangle_{I'}^{(i',j),(j',u,r)} \\ &\equiv \exists i', j'. \langle\langle c' \bullet t \rangle\rangle_{I'}^{(i,j'),(l,u,i')} * \langle\langle t' \rangle\rangle_{I'}^{(i',j),(j',u,r)} \\ &\equiv \langle\langle (c' \bullet t) \otimes t' \rangle\rangle_{I'}^{(i,j),(l,u,r)} \\ &\equiv \langle\langle (c' \otimes t') \bullet t \rangle\rangle_{I'}^{(i,j),(l,u,r)}. \end{aligned}$$

The case where $c = t' \otimes c'$ follows the same pattern as the $c' \otimes t'$ case. \square

The following lemma is used to prove the crust inclusion property. It essentially captures that a given tree and out-interface uniquely determine an in-interface. When the tree is non-empty, the in-interface is the identifiers of the left-most and right-most nodes at the top level of the tree; when the tree is empty, the out-interface provides the correct identifiers.

Lemma 5.7. For all $out \in \mathcal{I}_{out}$ and $t \in \text{Tree}$, there exists some $in \in \mathcal{I}_{in}$ such that

$$\langle\langle t \rangle\rangle^{in, out} \equiv \exists in'. \langle\langle t \rangle\rangle^{in', out}.$$

The proof is by a straightforward induction on the structure of the tree t .

Lemma 5.8 (τ_1 Crust Inclusion). For all $out, out' \in \mathcal{I}_{out}$, $F \in \mathcal{F}$ and $c \in \mathcal{C}_{Tree}$, there exist $K \in \mathcal{P}(\text{Heap})$ and $F' \in \mathcal{F}$ such that, for all $in \in \mathcal{I}_{in}$,

$$\left(\exists in'. \mathbb{M}_{in', out'}^F \circ \langle\langle c \rangle\rangle_{in, out}^{in', out'} \right) \equiv K \circ \mathbb{M}_{in, out}^{F'}.$$

Proof. The proof is by induction on the structure of the context c , assuming some fixed $out, out' \in \mathcal{I}_{out}$ and $F \in \mathcal{F}$.

$c = -$ case:

Choose $K = \text{emp}$ if $out' = out$ and $K = \text{False}$ otherwise. If $out' \neq out$ then both sides of the equation are equivalent to False ; otherwise, observe that

$$\begin{aligned} \exists in'. \mathbb{M}_{in', out'}^F \circ \langle\langle - \rangle\rangle_{in, out}^{in', out'} &\equiv \mathbb{M}_{in, out}^F \\ &\equiv K \circ \mathbb{M}_{in, out}^F. \end{aligned}$$

$c = \mathbf{a}[c']$ case:

Let $out'' = (\mathbf{nil}, \mathbf{a}, \mathbf{nil})$ and $F'' = (\mathbf{nil}, \mathbf{nil}, l, u, r, \mathbf{nil}, \mathbf{nil})$, where $(l, u, r) = out'$. By the inductive hypothesis, there exist K' and F' such that, for all in ,

$$\exists i', j'. \mathbb{M}_{(i', j'), out''}^{F''} \circ \langle\langle c' \rangle\rangle_{in, out}^{(i', j'), out''} \equiv K' \circ \mathbb{M}_{in, out}^{F'}.$$

Choose $K = \mathbb{M}_{(\mathbf{a}, \mathbf{a}), out'}^F * K'$; choose F' as given above. Observe that, for all in ,

$$\begin{aligned} \exists in'. \mathbb{M}_{in', out'}^F \circ \langle\langle c \rangle\rangle_{in, out}^{in', out'} &\equiv \exists in'. \mathbb{M}_{in', out'}^F * \exists i', j'. \mathbf{a} \mapsto l, u, d, r * \langle\langle c' \rangle\rangle_{I'}^{(i', j'), (\mathbf{nil}, \mathbf{a}, \mathbf{nil})} \wedge in' = (\mathbf{a}, \mathbf{a}) \\ &\equiv \mathbb{M}_{(\mathbf{a}, \mathbf{a}), out'}^F * \exists i', j'. \mathbb{M}_{(i', j'), out''}^{F''} \circ \langle\langle c' \rangle\rangle_{in, out}^{(i', j'), out''} \\ &\equiv \mathbb{M}_{(\mathbf{a}, \mathbf{a}), out'}^F * K' \circ \mathbb{M}_{in, out}^{F'} \\ &\equiv K \circ \mathbb{M}_{in, out}^{F'}. \end{aligned}$$

$c = c' \otimes t$ case:

We can assume that $t \neq \emptyset$, since otherwise $c = c'$ and the result holds by the inductive assumption. Thus $t = \mathbf{a}[t_1] \otimes t_2$ for some \mathbf{a}, t_1, t_2 . In order to apply the inductive hypothesis, we must determine the crust for the c' part of c . The left or parent node part of the crust will be the same as for c ; the right node will be \mathbf{a} .

Suppose that $out' = (l, u, r)$ and $F = (f_1, \dots, f_7)$. Let f'_6 be such that, for some j''' ,

$$\langle\langle t_1 \rangle\rangle_{(f'_6, j'''), (\mathbf{nil}, \mathbf{a}, \mathbf{nil})} \equiv \exists in'''. \langle\langle t_1 \rangle\rangle_{in''', (\mathbf{nil}, \mathbf{a}, \mathbf{nil})}.$$

Let f'_7 be such that, for some j ,

$$\langle\langle t_2 \rangle\rangle_{(f'_7, j), (\mathbf{a}, u, r)} \equiv \exists in''. \langle\langle t_2 \rangle\rangle_{in'', (\mathbf{a}, u, r)}.$$

Let $F' = (f_1, \dots, f_5, f'_6, f'_7)$. By the inductive hypothesis, let K' and F'' be such that,

for all in ,

$$\exists i, j'. \mathfrak{M}_{(i,j'),(l,u,\mathbf{a})}^{F'} \circ \langle\langle c' \rangle\rangle_{in,out}^{(i,j'),(l,u,\mathbf{a})} \equiv K' \circ \mathfrak{M}_{in,out}^{F''}$$

Now, by construction, for arbitrary in ,

$$\begin{aligned} & \exists i, j. \mathfrak{M}_{(i,j),out'}^F \circ \langle\langle c' \otimes t \rangle\rangle_{in,out}^{(i,j),(l,u,r)} \\ & \equiv \exists i, j. \mathfrak{M}_{(i,j),out'}^F \circ \exists i', j'. \langle\langle c' \rangle\rangle_{in,out}^{(i,j'),(l,u,i')} * \langle\langle \mathbf{a}[t_1] \otimes t_2 \rangle\rangle^{(i',j),(j',u,r)} \\ & \equiv \exists i, j. \mathfrak{M}_{(i,j),out'}^F \circ \exists i', j'. \langle\langle c' \rangle\rangle_{in,out}^{(i,j'),(l,u,i')} * \\ & \quad \exists i'', j''. \langle\langle \mathbf{a}[t_1] \rangle\rangle^{(i',j''),(j',u,i'')} * \langle\langle t_2 \rangle\rangle^{(i'',j),(j'',u,r)} \\ & \equiv \exists i, j. \mathfrak{M}_{(i,j),out'}^F \circ \exists j'. \langle\langle c' \rangle\rangle_{in,out}^{(i,j'),(l,u,\mathbf{a})} * \\ & \quad \exists i'', i''', j'''. \mathbf{a} \mapsto j', u, i''', i'' * \langle\langle t_1 \rangle\rangle^{(i'',j'''),(nil,\mathbf{a},nil)} * \langle\langle t_2 \rangle\rangle^{(i'',j),(a,u,r)} \\ & \equiv \exists i, j. \mathfrak{M}_{(i,j),out'}^F \circ \exists j'. \langle\langle c' \rangle\rangle_{in,out}^{(i,j'),(l,u,\mathbf{a})} * \\ & \quad \exists j'''. \mathbf{a} \mapsto j', u, f'_6, f'_7 * \langle\langle t_1 \rangle\rangle^{(f'_6,j'''),(nil,\mathbf{a},nil)} * \langle\langle t_2 \rangle\rangle^{(f'_7,j),(a,u,r)} \\ & \equiv \exists i, j, j'. (l \mapsto f_1, u, f_2, i \vee (l = \mathbf{nil} \wedge (u \mapsto f_3, f_4, i, f_5 \vee (u = \mathbf{nil} \wedge \text{emp})))) \\ & \quad * (r \mapsto j, u, f_6, f_7 \vee (r = \mathbf{nil} \wedge \text{emp})) \circ \langle\langle c' \rangle\rangle_{in,out}^{(i,j'),(l,u,\mathbf{a})} * \\ & \quad \mathbf{a} \mapsto j', u, f'_6, f'_7 * \exists j'''. \langle\langle t_1 \rangle\rangle^{(f'_6,j'''),(nil,\mathbf{a},nil)} * \langle\langle t_2 \rangle\rangle^{(f'_7,j),(a,u,r)} \\ & \equiv \exists i, j'. (l \mapsto f_1, u, f_2, i \vee (l = \mathbf{nil} \wedge (u \mapsto f_3, f_4, i, f_5 \vee (u = \mathbf{nil} \wedge \text{emp})))) \\ & \quad * \mathbf{a} \mapsto j', u, f'_6, f'_7 \circ \langle\langle c' \rangle\rangle_{in,out}^{(i,j'),(l,u,\mathbf{a})} \\ & \quad * \exists j, j''. (r \mapsto j, u, f_6, f_7 \vee (r = \mathbf{nil} \wedge \text{emp})) * \\ & \quad \langle\langle t_1 \rangle\rangle^{(f'_6,j'''),(nil,\mathbf{a},nil)} * \langle\langle t_2 \rangle\rangle^{(f'_7,j),(a,u,r)} \\ & \equiv \exists i, j'. \mathfrak{M}_{(i,j'),(l,u,\mathbf{a})}^{F'} \circ \langle\langle c' \rangle\rangle_{in,out}^{(i,j'),(l,u,\mathbf{a})} \\ & \quad * \exists j, j''. (r \mapsto j, u, f_6, f_7 \vee (r = \mathbf{nil} \wedge \text{emp})) * \\ & \quad \langle\langle t_1 \rangle\rangle^{(f'_6,j'''),(nil,\mathbf{a},nil)} * \langle\langle t_2 \rangle\rangle^{(f'_7,j),(a,u,r)} \\ & \equiv K' * \exists j, j''. (r \mapsto j, u, f_6, f_7 \vee (r = \mathbf{nil} \wedge \text{emp})) * \\ & \quad \langle\langle t_1 \rangle\rangle^{(f'_6,j'''),(nil,\mathbf{a},nil)} * \langle\langle t_2 \rangle\rangle^{(f'_7,j),(a,u,r)} \circ \mathfrak{M}_{in,out}^{F''} \end{aligned}$$

Let

$$\begin{aligned} K & \equiv K' * \exists j, j''. (r \mapsto j, u, f_6, f_7 \vee (r = \mathbf{nil} \wedge \text{emp})) * \\ & \quad \langle\langle t_1 \rangle\rangle^{(f'_6,j'''),(nil,\mathbf{a},nil)} * \langle\langle t_2 \rangle\rangle^{(f'_7,j),(a,u,r)}. \end{aligned}$$

K and F'' therefore have the desired property.

The case where $c = t \otimes c'$ is similar. \square

Lemma 5.9. For all $\varphi \in \text{Cmd}_{\mathbb{T}}$, $(P, Q) \in \text{Ax} \llbracket \varphi \rrbracket_{\mathbb{T}}$, $out \in \mathcal{I}_{out}$ and $F \in \mathcal{F}$,

$$\vdash_{\mathbb{H}} \left\{ \llbracket P \rrbracket^{out,F} \right\} \llbracket \varphi \rrbracket \left\{ \llbracket Q \rrbracket^{out,F} \right\}.$$

We omit the full proof details here, which can be found in (Dinsdale-Young et al., 2010b). Figs. 5 and 6 show proof outlines for the `getLast` axioms. The two cases cover when the node whose last child is to be fetched has and does not have children. The correctness of the axioms for `getLast` follow from these proofs by the procedure definition and procedure call rules.

This completes the proof of the soundness of the translation $\tau_1 : \mathbb{T} \rightarrow \mathbb{H}$ (Theorem 5.5).

5.5. Module Translation: $\tau_2 : \mathbb{T} \rightarrow \mathbb{H} + \mathbb{L}$

In this section, we give the formal definition of a locality-preserving translation τ_2 from the tree module \mathbb{T} into the combination of the heap and list modules $\mathbb{H} + \mathbb{L}$.

Notation.

A number of notational conventions are used to simplify predicates over $\mathcal{A}_{\mathbb{H}} + \mathcal{A}_{\mathbb{L}}$. Pure heap predicates are implicitly lifted to the combined domain with the list-store component `emp`; list-store predicates are lifted similarly. The operator $*$ on the combined domain is the product of the $*$ operators from each of the two domains.

Definition 5.6 ($\tau_2 : \mathbb{T} \rightarrow \mathbb{H} + \mathbb{L}$). The pre-locality-preserving translation $\tau_2 : \mathbb{T} \rightarrow \mathbb{H} + \mathbb{L}$ is constructed as follows:

- the in-interfaces $\mathcal{I}_{\text{in}} = (\text{Addr})^*$ are sequences of addresses;
- the out-interfaces $\mathcal{I}_{\text{out}} = \text{Addr}$ are addresses;
- the data representation function is defined inductively by

$$\begin{aligned} \langle\langle \emptyset \rangle\rangle^{in, out} &\stackrel{\text{def}}{=} \text{emp} \wedge in = \emptyset \\ \langle\langle \mathbf{a}[t] \rangle\rangle^{in, out} &\stackrel{\text{def}}{=} in = \mathbf{a} \wedge \exists a, l. \mathbf{a} \mapsto out, a * a \mapsto [l] * \langle\langle t \rangle\rangle^{l, \mathbf{a}} \\ \langle\langle t_1 \otimes t_2 \rangle\rangle^{in, out} &\stackrel{\text{def}}{=} \exists l_1, l_2. (in = l_1 \cdot l_2) \wedge \langle\langle t_1 \rangle\rangle^{l_1, out} * \langle\langle t_2 \rangle\rangle^{l_2, out}; \end{aligned}$$

- the context representation function is defined inductively by

$$\begin{aligned} \langle\langle - \rangle\rangle_{in', out'}^{in, out} &\stackrel{\text{def}}{=} \text{emp} \wedge (in = in') \wedge (out = out') \\ \langle\langle \mathbf{a}[c] \rangle\rangle_{I'}^{in, out} &\stackrel{\text{def}}{=} in = \mathbf{a} \wedge \exists a, l. \mathbf{a} \mapsto out, a * a \mapsto [l] * \langle\langle c \rangle\rangle_{I'}^{l, \mathbf{a}} \\ \langle\langle c \otimes t \rangle\rangle_{I'}^{in, out} &\stackrel{\text{def}}{=} \exists l_1, l_2. (in = l_1 \cdot l_2) \wedge \langle\langle c \rangle\rangle_{I'}^{l_1, out} * \langle\langle t \rangle\rangle_{I'}^{l_2, out} \\ \langle\langle t \otimes c \rangle\rangle_{I'}^{in, out} &\stackrel{\text{def}}{=} \exists l_1, l_2. (in = l_1 \cdot l_2) \wedge \langle\langle t \rangle\rangle_{I'}^{l_1, out} * \langle\langle c \rangle\rangle_{I'}^{l_2, out}; \end{aligned}$$

- the crust parameters $\mathcal{F} = (\text{Addr})^* \times (\text{Addr})^* \times \text{Addr}_{\text{nil}}$ are tuples of two sequences of addresses and a further address or `nil`;
- the crust predicate is defined as

$$\mathbb{m}_{in, out}^{l_1, l_2, a} \stackrel{\text{def}}{=} \exists a'. out \mapsto a, a' * a' \mapsto [l_1 \cdot in \cdot l_2] * \prod_{\mathbf{a} \in l_1 \cdot l_2}^* \mathbf{a} \mapsto out; \text{ and}$$

- the substitutive implementation function is given by replacing each tree-module command with a call to the correspondingly-named procedure given in Fig. 7.

Theorem 5.10 (Soundness of τ_2). The pre-locality-preserving translation τ_2 is a locality-preserving translation.

```

m := getLast(n) {
  { n ⇒ a * m ⇒ - × ∩(a,a),(l,u,r)F * ⟨⟨a[t1 ⊗ b[t2]]⟩⟩(a,a),(l,u,r) }
  local x in
    { ∃i, j. n ⇒ a * m ⇒ - * x ⇒ - × ∩(a,a),(l,u,r)F * a ↦ l, u, i, r * ⟨⟨t1 ⊗ b[t2]]⟩⟩(i,j),(nil,a,nil) }
    m := [n.down];
    { ∃i, j. n ⇒ a * m ⇒ i * x ⇒ - × ∩(a,a),(l,u,r)F * a ↦ l, u, i, r * ⟨⟨t1 ⊗ b[t2]]⟩⟩(i,j),(nil,a,nil) }
    { m ⇒ i * x ⇒ - × ⟨⟨t1 ⊗ b[t2]]⟩⟩(i,j),(nil,a,nil) }
    // t1 ⊗ b[t2] ≠ ∅ ⇒ i ≠ nil
    if m ≠ nil then
      { ∃d, x, t', t''. (i[t'] ⊗ t'' = t1 ⊗ b[t2]) ∧ m ⇒ i * x ⇒ -
        × i ↦ nil, a, d, x * ⟨⟨t'⟩⟩(d,-),(nil,i,nil) * ⟨⟨t''⟩⟩(x,j),(i,a,nil) }
      x := [m.right];
      { ∃m, j1, d, x, t', t''. (t ⊗ m[t'] ⊗ t'' = t1 ⊗ b[t2]) ∧ m ⇒ m * x ⇒ x
        × ⟨⟨t⟩⟩(i,j1),(nil,a,m) * m ↦ j1, a, d, x * ⟨⟨t'⟩⟩(d,-),(nil,m,nil) * ⟨⟨t''⟩⟩(x,j),(m,a,nil) }
      while x ≠ nil do
        { ∃m, j1, d, x, t', t''. (t ⊗ m[t'] ⊗ t'' = t1 ⊗ b[t2]) ∧ x ≠ nil ∧ m ⇒ m * x ⇒ x
          × ⟨⟨t⟩⟩(i,j1),(nil,a,m) * m ↦ j1, a, d, x
            * ⟨⟨t'⟩⟩(d,-),(nil,m,nil) * ⟨⟨t''⟩⟩(x,j),(m,a,nil) }
        { ∃m, d, x, x', t', t''. (t ⊗ x[t'] ⊗ t'' = t1 ⊗ b[t2]) ∧ m ⇒ m * x ⇒ x
          × ⟨⟨t⟩⟩(i,m),(nil,a,x) * x ↦ m, a, d, x' * ⟨⟨t'⟩⟩(d,-),(nil,x,nil) * ⟨⟨t''⟩⟩(x,j),(m,a,nil) }
        m := x;
        x := [m.right]
        { ∃m, j1, d, x, t', t''. (t ⊗ m[t'] ⊗ t'' = t1 ⊗ b[t2]) ∧ m ⇒ m * x ⇒ x
          × ⟨⟨t⟩⟩(i,j1),(nil,a,m) * m ↦ j1, a, d, x
            * ⟨⟨t'⟩⟩(d,-),(nil,m,nil) * ⟨⟨t''⟩⟩(x,j),(m,a,nil) }
        { ∃m, j1, d, t, t', t''. (t ⊗ m[t'] ⊗ t'' = t1 ⊗ b[t2]) ∧ m ⇒ m * x ⇒ nil
          × ⟨⟨t⟩⟩(i,j1),(nil,a,m) * m ↦ j1, a, d, nil
            * ⟨⟨t'⟩⟩(d,-),(nil,m,nil) * ⟨⟨t''⟩⟩(nil,j),(m,a,nil) }
        { ∃m, j1, d, t, t'. (t ⊗ m[t'] = t1 ⊗ b[t2]) ∧ m ⇒ m * x ⇒ nil
          × ⟨⟨t⟩⟩(i,j1),(nil,a,m) * m ↦ j1, a, d, nil * ⟨⟨t'⟩⟩(d,-),(nil,m,nil) }
        { ∃j1, d. m ⇒ b * x ⇒ nil
          × ⟨⟨t1⟩⟩(i,j1),(nil,a,b) * b ↦ j1, a, d, nil * ⟨⟨t2⟩⟩(d,-),(nil,b,nil) }
        { m ⇒ b * x ⇒ nil × ⟨⟨t1 ⊗ b[t2]]⟩⟩(i,j),(nil,a,nil) }
      { ∃i, j. n ⇒ a * m ⇒ b * x ⇒ nil × ∩(a,a),(l,u,r)F * a ↦ l, u, i, r * ⟨⟨t1 ⊗ b[t2]]⟩⟩(i,j),(nil,a,nil) }
      { n ⇒ a * m ⇒ b × ∩(a,a),(l,u,r)F * ⟨⟨a[t1 ⊗ b[t2]]⟩⟩(a,a),(l,u,r) }
    }
}

```

Fig. 5. Proof outline for getLast implementation (where node is present)

```

m := getLast(n) {
  { n ⇒ a * m ⇒ - × ∩(a,a),(l,u,r)F * ⟨⟨a[∅]⟩⟩(a,a),(l,u,r) }
  local x in
    { n ⇒ a * m ⇒ - * x ⇒ - × ∩(a,a),(l,u,r)F * a ↦ l, u, nil, r }
    m := [n.down];
    { n ⇒ a * m ⇒ nil * x ⇒ - × ∩(a,a),(l,u,r)F * a ↦ l, u, nil, r }
    if m ≠ nil then
      ...
      { n ⇒ a * m ⇒ nil * x ⇒ - × ∩(a,a),(l,u,r)F * a ↦ l, u, nil, r }
      { n ⇒ a * m ⇒ nil × ∩(a,a),(l,u,r)F * ⟨⟨a[∅]⟩⟩(a,a),(l,u,r) }
    }
}

```

Fig. 6. Proof outline for `getLast` implementation (where node is absent)5.5.1. Soundness of $\tau_2 : \mathbb{T} \rightarrow \mathbb{H} + \mathbb{L}$

Lemma 5.11 (τ_2 Application Preservation). For all $c \in \mathbb{C}_{\text{Tree}}$, $t \in \text{Tree}$ and $I \in \mathcal{I}$,

$$\langle\langle c \bullet t \rangle\rangle^I \equiv \exists I'. \langle\langle c \rangle\rangle_{I'}^I \bullet \langle\langle t \rangle\rangle^{I'}.$$

Proof. The proof is by induction on the structure of the context c , assuming some fixed $t \in \text{Tree}$.

$c = -$ case:

$$\begin{aligned} \exists I'. \langle\langle - \rangle\rangle_{I'}^I \bullet \langle\langle t \rangle\rangle^{I'} &\equiv \exists I'. I = I' \wedge \langle\langle t \rangle\rangle^{I'} \\ &\equiv \langle\langle t \rangle\rangle^I \\ &\equiv \langle\langle - \bullet t \rangle\rangle^I. \end{aligned}$$

$c = \mathbf{a}[c']$ case:

$$\begin{aligned} \exists I'. \langle\langle \mathbf{a}[c'] \rangle\rangle_{I'}^{in,out} \bullet \langle\langle t \rangle\rangle^{I'} &\equiv \exists I'. \left(in = \mathbf{a} \wedge \exists a, l. \mathbf{a} \mapsto out, a * a \mapsto [l] * \langle\langle c' \rangle\rangle_{I'}^{l,\mathbf{a}} \right) \bullet \langle\langle t \rangle\rangle^{I'} \\ &\equiv in = \mathbf{a} \wedge \exists a, l. \mathbf{a} \mapsto out, a * a \mapsto [l] * \left(\exists I'. \langle\langle c' \rangle\rangle_{I'}^{l,\mathbf{a}} \bullet \langle\langle t \rangle\rangle^{I'} \right) \\ &\equiv in = \mathbf{a} \wedge \exists a, l. \mathbf{a} \mapsto out, a * a \mapsto [l] * \langle\langle c' \bullet t \rangle\rangle^{l,\mathbf{a}} \\ &\equiv \langle\langle \mathbf{a}[c' \bullet t] \rangle\rangle^{in,out} \\ &\equiv \langle\langle \mathbf{a}[c'] \bullet t \rangle\rangle^{in,out}. \end{aligned}$$

$$E.\text{parent} \stackrel{\text{def}}{=} E$$

$$E.\text{children} \stackrel{\text{def}}{=} E + 1$$

$$n := \text{newNode}() \stackrel{\text{def}}{=} n := \text{alloc}(2)$$

$$\text{disposeNode}(E) \stackrel{\text{def}}{=} \text{dispose}(E, 2)$$

```

m := getUp(n) {
  local x in
    m := [n.parent];
    x := [m.parent];
    if x = nil then
      m := nil
}

m := getLast(n) {
  local x in
    x := [n.children];
    m := x.getTail()
}

newNodeAfter(n) {
  local x, y, z, w in
    x := [n.parent];
    z := [x.children];
    y := newNode();
    [y.parent] := x;
    z.insert(n, y);
    w := newList();
    [y.children] := w
}

m := getFirst(n) {
  local x in
    x := [n.children];
    m := x.getHead()
}

m := getRight(n) {
  local x, y in
    x := [n.parent];
    y := [x.children];
    m := y.getNext(n)
}

m := getLeft(n) {
  local x, y in
    x := [n.parent];
    y := [x.children];
    m := y.getPrev(n)
}

deleteTree(n) {
  local x, y, z in
    x := [n.parent];
    y := [x.children];
    y.remove(n);
    y := [n.children];
    z := y.getHead();
    while z ≠ nil do
      call deleteTree(z);
      z := y.getHead() ;
    deleteList(y);
    disposeNode(n)
}

```

Fig. 7. Procedures for the list-based implementation of trees

$c = c' \otimes t'$ case:

$$\begin{aligned}
& \exists I'. \langle\langle c' \otimes t' \rangle\rangle_{I'}^{in, out} \bullet \langle\langle t \rangle\rangle^{I'} \\
& \equiv \exists I'. \left(\exists l_1, l_2. (in = l_1 \cdot l_2) \wedge \langle\langle c' \rangle\rangle_{I'}^{l_1, out} * \langle\langle t' \rangle\rangle^{l_2, out} \right) \bullet \langle\langle t \rangle\rangle^{I'} \\
& \equiv \exists l_1, l_2. (in = l_1 \cdot l_2) \wedge \left(\exists I'. \langle\langle c' \rangle\rangle_{I'}^{l_1, out} \bullet \langle\langle t \rangle\rangle^{I'} \right) * \langle\langle t' \rangle\rangle^{l_2, out} \\
& \equiv \exists l_1, l_2. (in = l_1 \cdot l_2) \wedge \langle\langle c' \bullet t \rangle\rangle^{l_1, out} * \langle\langle t' \rangle\rangle^{l_2, out} \\
& \equiv \langle\langle (c' \bullet t) \otimes t' \rangle\rangle^{in, out} \\
& \equiv \langle\langle (c' \otimes t') \bullet t \rangle\rangle^{in, out}.
\end{aligned}$$

The case where $c = t' \otimes c'$ follows the same pattern as the $c' \otimes t'$ case. \square

Lemma 5.12 (τ_2 Crust Inclusion). For all $out, out' \in \mathcal{I}_{out}$, $F \in \mathcal{F}$ and $c \in \mathbf{C}_{Tree}$, there exist $K \in \mathcal{P}(\mathbf{C}_{\mathbb{H}+\mathbb{L}})$ and $F' \in \mathcal{F}$ such that, for all $in \in \mathcal{I}_{in}$,

$$\left(\exists in'. \mathbb{M}_{in', out'}^F \circ \langle\langle c \rangle\rangle_{in, out}^{in', out'} \right) \equiv K \circ \mathbb{M}_{in, out}^{F'}.$$

Proof. The proof is by induction on the structure of the context c , assuming some fixed $out, out' \in \mathcal{I}_{out}$ and $F \in \mathcal{F}$.

$c = -$ case:

Choose $K = \text{emp}$ if $out' = out$ and $K = \text{False}$ otherwise; choose $F' = F$. If $out' \neq out$ then both sides of the equation are equivalent to False ; otherwise, observe that

$$\begin{aligned}
\exists in'. \mathbb{M}_{in', out'}^F \circ \langle\langle - \rangle\rangle_{in, out}^{in', out'} & \equiv \mathbb{M}_{in, out}^F \\
& \equiv K \circ \mathbb{M}_{in, out}^{F'}.
\end{aligned}$$

$c = \mathbf{a}[c']$ case:

By the inductive hypothesis, there exist K' and F' such that, for all in ,

$$\exists l. \mathbb{M}_{l, \mathbf{a}}^{\emptyset, \emptyset, out'} \circ \langle\langle c' \rangle\rangle_{in, out}^{l, \mathbf{a}} \equiv K' \circ \mathbb{M}_{in, out}^{F'}.$$

Choose $K = \mathbb{M}_{\mathbf{a}, out'}^F \circ K'$; choose F' as given above. Observe that

$$\begin{aligned}
& \exists in'. \mathbb{M}_{in', out'}^F \circ \langle\langle \mathbf{a}[c'] \rangle\rangle_{in, out}^{in', out'} \\
& \equiv \exists in'. \mathbb{M}_{in', out'}^F \circ \left(in' = \mathbf{a} \wedge \exists a, l. \mathbf{a} \mapsto out', a * a \Rightarrow [l] * \langle\langle c' \rangle\rangle_{in, out}^{l, \mathbf{a}} \right) \\
& \equiv \mathbb{M}_{\mathbf{a}, out'}^F \circ \exists l. \mathbb{M}_{l, \mathbf{a}}^{\emptyset, \emptyset, out'} \circ \langle\langle c' \rangle\rangle_{in, out}^{l, \mathbf{a}} \\
& \equiv \mathbb{M}_{\mathbf{a}, out'}^F \circ K' \circ \mathbb{M}_{in, out}^{F'} \\
& \equiv K \circ \mathbb{M}_{in, out}^{F'}.
\end{aligned}$$

$c = c' \otimes t'$ case:

Observe that there is exactly one choice of $l_2 \in (\mathbf{Addr})^*$ such that $\langle\langle t' \rangle\rangle^{l_2, out'} \neq \text{false}$. Let \hat{l}_2 be that choice. Observe also that there exists some K' such that

$$\langle\langle t' \rangle\rangle^{\hat{l}_2, out'} \equiv K' * \prod_{\mathbf{a} \in \hat{l}_2}^* \mathbf{a} \mapsto out'.$$

Let $(l'_1, l'_2, a') = F$. By the inductive hypothesis, there exist K'' and F' such that, for all in ,

$$\exists l_1. \mathbb{M}_{l_1, out'}^{l'_1, \hat{l}_2, l'_2, a'} \circ \langle\langle c' \rangle\rangle_{in, out}^{l_1, out'} \equiv K'' \circ \mathbb{M}_{in, out}^{F'}$$

Choose $K = K' \circ K''$; choose F' as given above. Observe that

$$\begin{aligned} & \exists in'. \mathbb{M}_{in', out'}^{l'_1, l'_2, a'} \circ \langle\langle c' \otimes t' \rangle\rangle_{in, out}^{in', out'} \\ & \equiv \exists in'. \left(\exists a. out' \mapsto a', a * a \mapsto [l'_1 \cdot in' \cdot l'_2] * \prod_{\mathbf{a} \in l'_1 \cdot l'_2}^* \mathbf{a} \mapsto out' \right) \\ & \quad \circ \exists l_1. in' = l_1 \cdot \hat{l}_2 \wedge \langle\langle c' \rangle\rangle_{in, out}^{l_1, out'} * \langle\langle t' \rangle\rangle_{\hat{l}_2, out'} \\ & \equiv \exists l_1. \left(\exists a. out' \mapsto a', a * a \mapsto [l'_1 \cdot l_1 \cdot \hat{l}_2 \cdot l'_2] * \prod_{\mathbf{a} \in l'_1 \cdot l'_2}^* \mathbf{a} \mapsto out' \right) \\ & \quad \circ \langle\langle c' \rangle\rangle_{in, out}^{l_1, out'} * K' * \prod_{\mathbf{a} \in \hat{l}_2}^* \mathbf{a} \mapsto out' \\ & \equiv K' \circ \exists l_1. \mathbb{M}_{l_1, out'}^{l'_1, \hat{l}_2, l'_2, a'} \circ \langle\langle c' \rangle\rangle_{in, out}^{l_1, out'} \\ & \equiv K' \circ K'' \circ \mathbb{M}_{in, out}^{F'} \\ & \equiv K \circ \mathbb{M}_{in, out}^{F'} \end{aligned}$$

The case where $c = t' \otimes c'$ follows the same pattern as the $c' \otimes t'$ case. \square

Lemma 5.13 (τ_2 **Axiom Correctness**). For all $\varphi \in \text{Cmd}_{\mathbb{L}}$, $(P, Q) \in \text{AX}[\varphi]_{\mathbb{L}}$, $out \in \mathcal{I}_{out}$ and $F \in \mathcal{F}$,

$$\vdash_{\mathbb{H}+\mathbb{L}} \left\{ \langle\langle P \rangle\rangle^{out, F} \right\} \llbracket \varphi \rrbracket \left\{ \langle\langle Q \rangle\rangle^{out, F} \right\}.$$

We omit the full proof details here, which are due to Wheelhouse and can be found in (Dinsdale-Young et al., 2010b). Fig. 8 shows one of the simpler proof cases involved in establishing axiom correctness, namely for the `getUp` axiom.

This completes the proof of Theorem 5.10.

5.6. Module Translation: $\tau_3 : \mathbb{H} + \mathbb{H} \rightarrow \mathbb{H}$

Another example of a locality-preserving translation is the natural implementation of a pair of heap modules $\mathbb{H} + \mathbb{H}$ with a single heap module \mathbb{H} that simply treats the two heaps as (disjoint) portions of the same heap. This example is an important one, as it does not result in a surjective abstraction relation (different abstract heap pairs can map to the same concrete heap), but is still a sound locality preserving translation.

Definition 5.7 ($\tau_3 : \mathbb{H} + \mathbb{H} \rightarrow \mathbb{H}$). The pre-locality-preserving translation $\tau_3 : \mathbb{H} + \mathbb{H} \rightarrow \mathbb{H}$ is constructed as follows:

$$\begin{aligned}
m &:= \text{getUp}(n) \{ \\
&\left\{ m \Rightarrow - * n \Rightarrow \mathbf{b} \times \exists in. \mathfrak{M}_{in,out}^F \bullet \langle\langle \mathbf{a}[t_1] \otimes \mathbf{b}[t_2] \otimes t_3 \rangle\rangle^{in,out} \right\} \\
&\left\{ m \Rightarrow - * n \Rightarrow \mathbf{b} \times \mathfrak{M}_{\mathbf{a},out}^F \bullet \exists a, l_1, l_3. \mathbf{a} \mapsto out, a * a \Leftrightarrow [l_1 \cdot \mathbf{b} \cdot l_3] \right. \\
&\quad \left. * \langle\langle t_1 \rangle\rangle^{l_1, \mathbf{a}} * \exists a', l. \mathbf{b} \mapsto \mathbf{a}, a' * a' \Leftrightarrow [l_2] * \langle\langle t_2 \rangle\rangle^{l_2, \mathbf{b}} * \langle\langle t_3 \rangle\rangle^{l_3, \mathbf{a}} \right\} \\
&\left\{ m \Rightarrow - * n \Rightarrow \mathbf{b} \times \mathbf{a} \mapsto out * \mathbf{b} \mapsto \mathbf{a} \right\} \\
&\text{local } x \text{ in} \\
&\left\{ m \Rightarrow - * n \Rightarrow \mathbf{b} * x \Rightarrow - \times \mathbf{a} \mapsto out * \mathbf{b} \mapsto \mathbf{a} \right\} \\
&m := [n.\text{parent}]; \\
&\left\{ m \Rightarrow \mathbf{a} * n \Rightarrow \mathbf{b} * x \Rightarrow - \times \mathbf{a} \mapsto out * \mathbf{b} \mapsto \mathbf{a} \right\} \\
&x := [m.\text{parent}]; \\
&\left\{ m \Rightarrow \mathbf{a} * n \Rightarrow \mathbf{b} * x \Rightarrow out \times \mathbf{a} \mapsto out * \mathbf{b} \mapsto \mathbf{a} \right\} \\
&\text{if } x = \text{nil} \text{ then} \\
&\quad m := \text{nil} \\
&\left\{ m \Rightarrow \mathbf{a} * n \Rightarrow \mathbf{b} * x \Rightarrow out \times \mathbf{a} \mapsto out * \mathbf{b} \mapsto \mathbf{a} \right\} \\
&\left\{ m \Rightarrow \mathbf{a} * n \Rightarrow \mathbf{b} \times \mathbf{a} \mapsto out * \mathbf{b} \mapsto \mathbf{a} \right\} \\
&\left\{ m \Rightarrow \mathbf{a} * n \Rightarrow \mathbf{b} \times \mathfrak{M}_{\mathbf{a},out}^F \bullet \exists a, l_1, l_3. \mathbf{a} \mapsto out, a * a \Leftrightarrow [l_1 \cdot \mathbf{b} \cdot l_3] \right. \\
&\quad \left. * \langle\langle t_1 \rangle\rangle^{l_1, \mathbf{a}} * \exists a', l. \mathbf{b} \mapsto \mathbf{a}, a' * a' \Leftrightarrow [l_2] * \langle\langle t_2 \rangle\rangle^{l_2, \mathbf{b}} * \langle\langle t_3 \rangle\rangle^{l_3, \mathbf{a}} \right\} \\
&\left\{ m \Rightarrow \mathbf{a} * n \Rightarrow \mathbf{b} \times \exists in. \mathfrak{M}_{in,out}^F \bullet \langle\langle \mathbf{a}[t_1] \otimes \mathbf{b}[t_2] \otimes t_3 \rangle\rangle^{in,out} \right\} \\
&\} \\
&\}
\end{aligned}$$
Fig. 8. Proof outline for `getUp` implementation

- the interface sets are both unit sets,[¶] *i.e.* $\mathcal{I}_{in} = \{1\} = \mathcal{I}_{out}$, and hence they can be ignored;
- the representation function (which is the same for both data and contexts) is defined as

$$\langle\langle (h_1, h_2) \rangle\rangle \stackrel{\text{def}}{=} \{h_1\} * \{h_2\};$$

- the crust parameter set is also the unit set and the crust predicate is defined as $\mathfrak{M} = \text{emp}$; and
- the implementation function is given by replacing the commands for both heaps with their detagged versions, for example

$$\llbracket n := \text{alloc}_1(E) \rrbracket = n := \text{alloc}(E) = \llbracket n := \text{alloc}_2(E) \rrbracket.$$

Theorem 5.14 (Soundness of τ_3). The pre-locality-preserving translation τ_3 is a locality-preserving translation.

This theorem is trivial: application preservation holds by properties of $*$; crust inclusion holds since the crust is simply emp ; and axiom correctness holds because the axioms of $\mathbb{H} + \mathbb{H}$ are almost directly translated to those of \mathbb{H} , modulo a frame. Notice, however, that

[¶] It makes no difference which set, as long as it only has one element.




Fig. 9. Representation of the list store $a_1 \mapsto [v_1 \cdot v_2 \cdot v_3] * a_2 \mapsto [w_1 \cdot v_1]$ as singly linked lists in the heap

this translation does not satisfy the first of the properties for including the conjunction rule (§5.3), since $\langle\langle(1 \mapsto 0, \text{emp})\rangle\rangle = \{1 \mapsto 0\} = \langle\langle(\text{emp}, 1 \mapsto 0)\rangle\rangle$.

If a command like `allocEither` from §2 were added to $\mathbb{H} + \mathbb{H}$, then without `CONJ` it could be soundly implemented in \mathbb{H} by allocating a single cell and returning its address. With `CONJ`, however, this does not work; the implementation must diverge.

6. Locality-Breaking Translations

There is not always a close correspondence between the locality exhibited by a high-level module and the locality of the low-level module on which it is implemented. In this section, we reduce the burden of proof for a sound module translation in such cases by introducing *locality-breaking translations*. In §6.1, we establish that locality-breaking translations give sound module translations, and in §6.3 we give a locality-breaking translation $\tau_4 : \mathbb{L} \rightarrow \mathbb{H}$.

Our motivating example is an implementation of the list module (as defined in §3.3) that represents each abstract list with a singly-linked list in the heap. An example of a list store represented in this way is depicted in Fig. 9. Consider the operation of removing the value v_3 from the list at address a_1 . At the abstract level, the resource required by the operation is simply $a_1 \mapsto v_3$, but in the implementation the list at a_1 must be traversed all the way to the node labelled v_3 , in this case this is the entire list.

To apply the locality-preserving translation approach, an arbitrary amount of additional data would have to be included in the crust. This seems to defeat the purpose of the locality-preserving approach in the first place, so perhaps a more direct approach is called for.

In order to prove the soundness of a module translation, it is necessary to demonstrate a transformation from high-level proofs about programs that use an abstract module, to low-level proofs of those programs using the concrete module implementation. Inductively transforming proofs is, intuitively, a good strategy because the high-level rules are matched by the low-level rules. Since the definition of predicate translations preserves disjunctions and entailments, and also the variable scope, the majority of the proof rules can be inductively transformed to their low-level counterparts directly. The two exceptions to this are `FRAME` and `AXIOM`. For locality-preserving translations, `FRAME` was dealt with by the fact that a locality-preserving translation preserves context application,

and AXIOM was dealt with by the fact that the implementations of the basic commands satisfy the axioms under such a translation.

When it is not appropriate to consider a translation that preserves locality, it would be useful if it were only necessary to consider proofs which do not use the frame rule, or, at least, only use it in a limited fashion. Intuitively, of course this should be possible. The purpose of the frame rule is to factor out parts of the state that do not play a role (and hence, do not change) in part of the program under consideration. It is common folklore that it should be possible to transform a proof to one in which the frame rule is only applied to basic statements (*i.e.* basic commands and assignment) by factoring in the state earlier in the proof (*i.e.* at the leaves of the proof). This intuition is formalised in the following lemma.

Lemma 6.1 (Frame-Free Derivations). Let \mathbb{A} be an abstract module. If there is a proof derivation of $\vdash_{\mathbb{A}} \{P\} \mathbb{C} \{Q\}$ then there is also a derivation that only uses the frame rule in the following ways:

$$\frac{\overline{\Gamma \vdash_{\mathbb{A}} \{P'\} \mathbb{C}' \{Q'\}} \quad (\dagger)}{\Gamma \vdash_{\mathbb{A}} \{K \bullet P'\} \mathbb{C}' \{K \bullet Q'\}} \text{FRAME} \quad (2)$$

$$\frac{\overline{\Gamma \vdash_{\mathbb{A}} \{P'\} \mathbb{C}' \{Q'\}} \quad \vdots}{\Gamma \vdash_{\mathbb{A}} \{(K_{\text{Scope}} \times \mathbf{I}_{\mathbb{A}}) \bullet P'\} \mathbb{C}' \{(K_{\text{Scope}} \times \mathbf{I}_{\mathbb{A}}) \bullet Q'\}} \text{FRAME} \quad (3)$$

where (\dagger) is either AXIOM or ASSGN.

Consider a translation $\tau : \mathbb{A} \rightarrow \mathbb{B}$. Lemma 6.1 implies that it is only necessary to provide proofs of $\vdash_{\mathbb{B}} \{\llbracket P \rrbracket_{\tau}\} \llbracket \mathbb{C} \rrbracket_{\tau} \{\llbracket Q \rrbracket_{\tau}\}$ when there is a proof of $\vdash_{\mathbb{A}} \{P\} \mathbb{C} \{Q\}$ having the prescribed form. Provided that there are proofs that the implementation of each command in $\text{Cmd}_{\mathbb{A}}$ satisfies the translation of its axioms under every possible frame, the proof in \mathbb{A} can be transformed to a proof in \mathbb{B} by straightforward induction. In fact, we only need to consider *singleton* frames, as we can treat an arbitrary frame as the disjunction of singletons and apply the DISJ rule. We can further reduce our considerations to those singleton frames with no variable scope component, since the variable scope component can be added by the FRAME rule at the low-level. These conditions are formalised in the definition of a locality-breaking translation.

Definition 6.1 (Locality-Breaking Translation). A *locality-breaking translation* $\tau : \mathbb{A} \rightarrow \mathbb{B}$ is a module translation having the property that, for all $c \in \mathbb{C}_{\mathbb{A}}$, $\varphi \in \text{Cmd}_{\mathbb{A}}$, and $(P, Q) \in \text{Ax} \llbracket \varphi \rrbracket_{\mathbb{A}}$ there is a derivation of

$$\vdash_{\mathbb{B}} \{\llbracket \{(\emptyset, c)\} \bullet P \rrbracket_{\tau}\} \llbracket \varphi \rrbracket_{\tau} \{\llbracket \{(\emptyset, c)\} \bullet Q \rrbracket_{\tau}\}.$$

Theorem 6.2 (Soundness of Locality-Breaking Translations). A locality-breaking translation is a sound module translation.

A locality-breaking translation transforms proofs that use locality (in the form of the FRAME rule) at the abstract level into proofs that do not. To do so, we must effectively

prove directly that the abstract FRAME rule is sound with respect to the implementation of each module operation. Hence we say that such a module translation provides a *fiction of locality*.

Since a locality-breaking translation transforms proofs that use locality (in the form of the frame rule) into proofs that do not, we

6.1. Proof of Soundness of Locality-Breaking Translations

Proof of Lemma 6.1 The result is a special case of the more general result, that if there is a derivation of $\Gamma \vdash_{\mathbb{A}} \{P\} \mathbb{C} \{Q\}$ then there is a derivation of $F(\Gamma) \vdash_{\mathbb{A}} \{P\} \mathbb{C} \{Q\}$ with the required property, where

$$F(\Gamma) = \{\mathbf{f} : K \bullet P \mapsto K \bullet Q \mid K \in \mathcal{P}(C_{\mathbb{A}}) \text{ and } (\mathbf{f} : P \mapsto Q) \in \Gamma\}.$$

Note that $\Gamma \subseteq F(\Gamma) = F(F(\Gamma))$. Since procedure specifications are only relevant to the PDEF and PCALL rules, we will omit them when considering other rules.

The proof of the generalised statement is by induction on the depth of the derivation. If the last rule applied in the derivation is anything other than FRAME or PDEF then it is simple to transform the derivation: simply apply the induction hypothesis to transform all of the premises and then apply the last rule using $F(\gamma)$ in place of γ .

Consider the case where the last rule of the derivation is FRAME:

$$\frac{\frac{\vdots}{\{P'\} \mathbb{C} \{Q'\}} (\ddagger)}{\{K \bullet P'\} \mathbb{C} \{K \bullet Q'\}} \text{FRAME}$$

By applying the disjunction rule, this can be reduced to the case of singleton frames $\{c\}$, transforming the derivation as follows:

$$\frac{\text{for all } c \in K, \frac{\frac{\vdots}{\{P'\} \mathbb{C} \{Q'\}} (\ddagger)}{\{\{c\} \bullet P'\} \mathbb{C} \{\{c\} \bullet Q'\}} \text{FRAME}}{\{K \bullet P'\} \mathbb{C} \{K \bullet Q'\}} \text{DISJ}$$

Now consider cases for (\ddagger) , the last rule applied before FRAME.

If the rule is CONS then, since $P \subseteq P'$ implies that $\{c\} \bullet P \subseteq \{c\} \bullet P'$, the application of the FRAME can be moved earlier in the derivation, transforming it as follows:

$$\frac{\frac{\{c\} \bullet P' \subseteq \{c\} \bullet P'' \quad \frac{\{P''\} \mathbb{C} \{Q''\}}{\{\{c\} \bullet P''\} \mathbb{C} \{\{c\} \bullet Q''\}} \text{FRAME}}{\{\{c\} \bullet P'\} \mathbb{C} \{\{c\} \bullet Q'\}} \text{CONS}}{\{\{c\} \bullet P'\} \mathbb{C} \{\{c\} \bullet Q'\}} \text{CONS}$$

The application of the frame rule can then be removed by the inductive hypothesis.

If the rule is DISJ then, since \bullet right-distributes over \vee , the derivation can be trans-

formed as follows:

$$\frac{\frac{\vdots}{\{P_i\} \mathbb{C} \{Q_i\}}}{\text{for all } i \in I, \quad \{\{c\} \bullet P_i\} \mathbb{C} \{\{c\} \bullet Q_i\}} \text{ FRAME}}{\{\{c\} \bullet \bigvee_{i \in I} P_i\} \mathbb{C} \{\{c\} \bullet \bigvee_{i \in I} Q_i\}} \text{ DISJ}}$$

If the rule is LOCAL then it is possible that the frame c includes a program variable with the same name as one that is scoped by the `local` block. This means that the frame cannot in general be pushed into the local block. However, the frame can be split into scope and store components, that is, for some $\rho \in \text{Scope}$ and $c_{\mathbb{A}} \in \mathbb{C}_{\mathbb{A}}$, $c = (\rho, c_{\mathbb{A}})$ and so $\{c\} = (\{\rho\} \times \mathbf{I}_{\mathbb{A}}) \circ \{(\emptyset, c_{\mathbb{A}})\}$. Hence the derivation can be transformed as follows:

$$\frac{\frac{\frac{\vdots}{\{(\mathbf{x} \Rightarrow - \times \mathbf{I}_{\mathbb{A}}) \bullet P'\} \mathbb{C}' \{(\mathbf{x} \Rightarrow - \times \mathbf{I}_{\mathbb{A}}) \bullet Q'\}}}{\{(\mathbf{x} \Rightarrow - \times \{c_{\mathbb{A}}\}) \bullet P'\} \mathbb{C}' \{(\mathbf{x} \Rightarrow - \times \{c_{\mathbb{A}}\}) \bullet Q'\}} \text{ FRAME}}{\{\{(\emptyset, c_{\mathbb{A}})\} \bullet P'\} \text{ local } \mathbf{x} \text{ in } \mathbb{C}' \{\{(\emptyset, c_{\mathbb{A}})\} \bullet Q'\}} \text{ LOCAL}}{\{\{c\} \bullet P'\} \text{ local } \mathbf{x} \text{ in } \mathbb{C}' \{\{c\} \bullet Q'\}} \text{ FRAME}}$$

The side condition for the LOCAL rule, that $(\{(\emptyset, c_{\mathbb{A}})\} \bullet P') \wedge \text{vsafe}(\mathbf{x}) \equiv \emptyset$, follows from the original side-condition that $P' \wedge \text{vsafe}(\mathbf{x}) \equiv \emptyset$. The applications of the FRAME rule are now either of the form of (3) or can be removed by the inductive hypothesis.

If the rule is PCALL then it is again necessary to split the framing context into its components, that is, some $\rho \in \text{Scope}$ and $c_{\mathbb{A}} \in \mathbb{C}_{\mathbb{A}}$ with $c = (\rho, c_{\mathbb{A}})$. The PCALL rule uses some $\mathbf{f} : \mathbf{P} \mapsto \mathbf{Q} \in \Gamma$. By definition, $\mathbf{f} : \{c_{\mathbb{A}}\} \bullet \mathbf{P} \mapsto \{c_{\mathbb{A}}\} \bullet \mathbf{Q} \in F(\Gamma)$. Hence, the derivation can be transformed as follows:

$$\frac{\frac{\frac{\{(\vec{\tau} \Rightarrow \vec{w} * \rho')\} \times D_{\mathbb{A}} \subseteq \text{vsafe}(\vec{E})}{F(\Gamma) \vdash_{\mathbb{A}} \left\{ (\vec{\tau} \Rightarrow \vec{v} * \rho') \times \left(\{(\{c_{\mathbb{A}}\}) \bullet \mathbf{P} \left(\mathcal{E} \left[\vec{E} \right] (\vec{\tau} \Rightarrow \vec{v} * \rho') \right) \right) \right\}}}{\text{call } \vec{\tau} := \mathbf{f}(\vec{E})} \text{ PCALL}}{\frac{\{(\exists \vec{w}. \{(\vec{\tau} \Rightarrow \vec{w} * \rho')\} \times (\{(\{c_{\mathbb{A}}\}) \bullet \mathbf{Q}(\vec{w}))\})\}}{F(\Gamma) \vdash_{\mathbb{A}} \left\{ (\{(\rho, c_{\mathbb{A}})\} \bullet \left((\vec{\tau} \Rightarrow \vec{v} * \rho') \times \mathbf{P} \left(\mathcal{E} \left[\vec{E} \right] (\vec{\tau} \Rightarrow \vec{v} * \rho') \right) \right) \right\}}}{\text{call } \vec{\tau} := \mathbf{f}(\vec{E})} \text{ FRAME}} \{\{(\rho, c_{\mathbb{A}})\} \bullet (\exists \vec{w}. \{(\vec{\tau} \Rightarrow \vec{w} * \rho')\} \times \mathbf{Q}(\vec{w}))\}}$$

The application of the frame rule is now of the form of (3), with the frame being $\{\rho\} \times \mathbf{I}_{\mathbb{A}}$.

The remaining cases for the penultimate rule are straightforward.

Consider the case when PDEF is the last rule applied:

$$\begin{array}{c}
 \vdots \\
 \hline
 \text{for all } (\mathbf{f}_i : P \multimap Q) \in \Gamma, \quad \Gamma', \Gamma \vdash_{\mathbb{A}} \frac{\{\exists \vec{v}. \vec{x}_i \Rightarrow \vec{v} * \vec{r}_i \Rightarrow - \times P(\vec{v})\}}{\{\exists \vec{w}. \vec{x}_i \Rightarrow - * \vec{r}_i \Rightarrow \vec{w} \times Q(\vec{w})\}} \\
 \hline
 (\star) \\
 \vdots \\
 (\star) \quad \frac{\Gamma', \Gamma \vdash_{\mathbb{A}} \{P\} \mathbb{C}' \{Q\}}{\Gamma' \vdash_{\mathbb{A}} \{P\} \text{procs } \vec{r}_1 := \mathbf{f}_1(\vec{x}_1) \{\mathbb{C}_1\}, \dots, \vec{r}_k := \mathbf{f}_k(\vec{x}_k) \{\mathbb{C}_k\} \text{ in } \mathbb{C}' \{Q\}} \text{PDEF}
 \end{array}$$

The derivations for the function bodies can be extended by applying the frame rule to give:

$$\begin{array}{c}
 \vdots \\
 \hline
 \Gamma', \Gamma \vdash_{\mathbb{A}} \frac{\{\exists \vec{v}. \vec{x}_i \Rightarrow \vec{v} * \vec{r}_i \Rightarrow - \times P(\vec{v})\}}{\{\exists \vec{w}. \vec{x}_i \Rightarrow - * \vec{r}_i \Rightarrow \vec{w} \times Q(\vec{w})\}} \\
 \hline
 \text{for all } K \in \mathcal{P}(\mathbb{C}_{\mathbb{A}}), \quad \Gamma', \Gamma \vdash_{\mathbb{A}} \frac{\{\exists \vec{v}. \vec{x}_i \Rightarrow \vec{v} * \vec{r}_i \Rightarrow - \times (K \bullet P(\vec{v}))\}}{\{\exists \vec{w}. \vec{x}_i \Rightarrow - * \vec{r}_i \Rightarrow \vec{w} \times (K \bullet Q(\vec{w}))\}} \text{FRAME} \\
 (\mathbf{f}_i : P \multimap Q) \in \Gamma,
 \end{array}$$

These derivations and the derivation of the premise $\Gamma', \Gamma \vdash_{\mathbb{A}} \{P\} \mathbb{C}' \{Q\}$ can be transformed by the inductive hypothesis so that they only use the frame rule in the required manner and use the procedure environment $F(\Gamma', \Gamma) = F(\Gamma'), F(\Gamma)$. These derivations can then be recombined to give the required derivation as follows:

$$\begin{array}{c}
 \vdots \\
 \hline
 \text{for all } (\mathbf{f}_i : P \multimap Q) \in F(\Gamma), \quad F(\Gamma', \Gamma) \vdash_{\mathbb{A}} \frac{\{\exists \vec{v}. \vec{x}_i \Rightarrow \vec{v} * \vec{r}_i \Rightarrow - \times P(\vec{v})\}}{\{\exists \vec{w}. \vec{x}_i \Rightarrow - * \vec{r}_i \Rightarrow \vec{w} \times Q(\vec{w})\}} \\
 \hline
 (\star) \\
 \vdots \\
 (\star) \quad \frac{F(\Gamma', \Gamma) \vdash_{\mathbb{A}} \{P\} \mathbb{C}' \{Q\}}{F(\Gamma') \vdash_{\mathbb{A}} \text{procs } \vec{r}_1 := \mathbf{f}_1(\vec{x}_1) \{\mathbb{C}_1\}, \dots, \vec{r}_k := \mathbf{f}_k(\vec{x}_k) \{\mathbb{C}_k\} \text{ in } \mathbb{C}' \{Q\}} \text{PDEF}
 \end{array}$$

The two further conditions on the PDEF rule, not included above, hold for the transformed derivation because F preserves the names of the functions in procedure specifications. \square

Let $\tau : \mathbb{A} \rightarrow \mathbb{B}$ be a locality-breaking translation. To show that τ is a sound module translation, it is necessary to establish that whenever there is a derivation of $\vdash_{\mathbb{A}} \{P\} \mathbb{C} \{Q\}$ there is a derivation of $\vdash_{\mathbb{B}} \{[[P]]_{\tau}\} [[\mathbb{C}]]_{\tau} \{[[Q]]_{\tau}\}$. First, transform the high-level derivation using Lemma 6.1 into a frame-free derivation. Now transform this deriva-

tion to the required low-level derivation by replacing each subderivation of the form

$$\frac{\Gamma \vdash_{\mathbb{A}} \{P'\} \varphi \{Q'\}}{\Gamma \vdash_{\mathbb{A}} \{K \bullet P'\} \varphi \{K \bullet Q'\}} \text{FRAME} \quad \text{AXIOM}$$

with the derivation

$$\frac{\text{for all } (\rho, c_{\mathbb{A}}) \in K, \quad \frac{\frac{\vdash_{\mathbb{B}} \{[\{(\emptyset, c_{\mathbb{A}}) \bullet P'\}]\} [\varphi] \{[\{(\emptyset, c_{\mathbb{A}}) \bullet Q'\}]\}}{\vdash_{\mathbb{B}} \{[\{(\rho, c_{\mathbb{A}}) \bullet P'\}]\} [\varphi] \{[\{(\rho, c_{\mathbb{A}}) \bullet Q'\}]\}} \text{FRAME}}{\vdash_{\mathbb{B}} \{[K \bullet P']\} [\varphi] \{[K \bullet Q']\}} \text{DISJ}}}{\vdash_{\mathbb{B}} \{[K \bullet P']\} [\varphi] \{[K \bullet Q']\}} \text{DISJ} \quad (\star)$$

where (\star) stands for the framed derivation provided by the locality-breaking translation, and replacing all other rules with their low-level equivalents.

This completes the proof of Theorem 6.2.

6.2. Including the Conjunction Rule

If we wish to add the conjunction rule to the locality-breaking theory, we can add a case to the proof of Lemma 6.1 to deal with pushing FRAME over CONJ, in a similar fashion to DISJ, provided that the context algebra $\mathcal{A}_{\mathbb{A}}$ is left-cancellative.

Definition 6.2 (Left-cancellativity). A context algebra $\mathcal{A} = (\mathbb{C}, \mathbb{D}, \circ, \bullet, \mathbf{I}, \mathbf{0})$ is left-cancellative if, for all $c \in \mathbb{C}$, $d_1, d_2, d_3 \in \mathbb{D}$, $c \bullet d_1 = c \bullet d_2 = d_3$ implies $d_1 = d_2$.

Left-cancellativity ensures that $\{c\} \bullet \bigwedge_{i \in I} P_i \equiv \bigwedge_{i \in I} \{c\} \bullet P_i$. It is also necessary for the predicate translation $\llbracket (\cdot) \rrbracket$ to distribute over conjunction; this is equivalent to the condition that the abstraction relation α is functional (that is, it defines a partial function from concrete states to abstract states).

6.3. Module Translation: $\tau_4 : \mathbb{L} \rightarrow \mathbb{H}$

We return to the example of an implementation of the list-store module \mathbb{L} with singly-linked lists in the heap module \mathbb{H} to illustrate a locality-breaking translation.

Definition 6.3 ($\tau_4 : \mathbb{L} \rightarrow \mathbb{H}$). The module translation $\tau_4 : \mathbb{L} \rightarrow \mathbb{H}$ is constructed as follows:

— the abstraction relation $\alpha_{\tau_4} \subseteq \text{Heap} \times \text{LStore}$ is defined by

$$h \alpha_{\tau_4} ls \iff h \in \llbracket ls \rrbracket$$

where $\llbracket (\cdot) \rrbracket : \text{LStore} \rightarrow \mathcal{P}(\text{Heap})$ is defined inductively as follows:

$$\begin{aligned} \llbracket \text{emp} \rrbracket &\stackrel{\text{def}}{=} \text{emp} \\ \llbracket a \mapsto l * ls \rrbracket &\stackrel{\text{def}}{=} \text{false} \\ \llbracket a \mapsto [l] * ls \rrbracket &\stackrel{\text{def}}{=} \exists x. a \mapsto x * \llbracket l \rrbracket^{(x, \text{nil})} * \llbracket ls \rrbracket \end{aligned}$$

where

$$\begin{aligned} \llbracket \emptyset \rrbracket^{(x,y)} &\stackrel{\text{def}}{=} (x = y) \wedge \text{emp} \\ \llbracket v \rrbracket^{(x,y)} &\stackrel{\text{def}}{=} x \mapsto v, y \\ \llbracket l_1 \cdot l_2 \rrbracket^{(x,y)} &\stackrel{\text{def}}{=} \exists z. \llbracket l_1 \rrbracket^{(x,z)} * \llbracket l_2 \rrbracket^{(z,y)}; \text{ and} \end{aligned}$$

— the substitutive implementation function is given by replacing each list-module command with a call to the correspondingly-named procedure given in Figs. 10 and 11.

Note that the abstraction relation is not surjective, since incomplete lists do not have heap representations (they are mapped to false). The intuition behind this approach is that incomplete lists are purely a useful means to the ultimate end of reasoning about complete lists. Typically, clients of the list module will work with complete lists (only complete lists may be created or deleted, for a start) and so restricting their specifications to only describe stores of complete lists should be no significant problem. Of course, it is perfectly acceptable to use assertions and specifications that refer to incomplete lists within client proofs; the transformation of this proof to a low-level proof will complete all of these lists by making use of Lemma 6.1.

The fact that incomplete lists do not have representations can be seen as an advantage from the point of view of establishing that τ_4 is a locality-breaking translation. This is because it is only necessary to prove that the framed axioms hold under the translation for frames that complete all of the lists in the precondition. In all other cases, the precondition is false, and so the triple holds trivially.

Theorem 6.3 (Soundness of τ_4). The module translation τ_4 is a locality-breaking translation.

Again, we omit the full details of this proof, which can be found in (Dinsdale-Young et al., 2010b), but give a particular case: `getNext`. Recall the axiom for `getNext`:

$$\begin{aligned} \text{AX } \llbracket \mathbf{x} := E_1.\text{getNext}(E_2) \rrbracket_{\mathbb{L}} &\stackrel{\text{def}}{=} \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times a \mapsto w \cdot u), \\ (\mathbf{x} \Rightarrow u * \rho \times a \mapsto w \cdot u) \end{array} \right) \middle| \begin{array}{l} a = \mathcal{E} \llbracket E_1 \rrbracket (\mathbf{x} \Rightarrow v * \rho) \text{ and} \\ w = \mathcal{E} \llbracket E_2 \rrbracket (\mathbf{x} \Rightarrow v * \rho) \end{array} \right\} \cup \\ &\left\{ \left(\begin{array}{l} (\mathbf{x} \Rightarrow v * \rho \times a \mapsto [l \cdot w]), \\ (\mathbf{x} \Rightarrow \mathbf{nil} * \rho \times a \mapsto [l \cdot w]) \end{array} \right) \middle| \begin{array}{l} a = \mathcal{E} \llbracket E_1 \rrbracket (\mathbf{x} \Rightarrow v * \rho) \text{ and} \\ w = \mathcal{E} \llbracket E_2 \rrbracket (\mathbf{x} \Rightarrow v * \rho) \end{array} \right\}. \end{aligned}$$

Fix arbitrary $a \in \text{Addr}$, $w, u \in \text{Val}$, $l \in \text{Val}^*$ and $c \in \text{C}_{\text{LStore}}$. It is sufficient to establish that the procedure body of `getNext` meets the following specifications:

$$\begin{aligned} &\{ \llbracket (\mathbf{i} \Rightarrow a) * (w \Rightarrow w) * (v \Rightarrow -) \times (\{c\} \bullet a \mapsto w \cdot u) \rrbracket \\ &\quad \llbracket \mathbf{v} := \mathbf{i}.\text{getNext}(w) \rrbracket \} \quad (4) \\ &\{ \llbracket (\mathbf{i} \Rightarrow a) * (w \Rightarrow w) * (v \Rightarrow u) \times (\{c\} \bullet a \mapsto w \cdot u) \rrbracket \} \end{aligned}$$

$$\begin{aligned} &\{ \llbracket (\mathbf{i} \Rightarrow a) * (w \Rightarrow w) * (v \Rightarrow -) \times (\{c\} \bullet a \mapsto [l \cdot w]) \rrbracket \\ &\quad \llbracket \mathbf{v} := \mathbf{i}.\text{getNext}(w) \rrbracket \} \quad (5) \\ &\{ \llbracket (\mathbf{i} \Rightarrow a) * (w \Rightarrow w) * (v \Rightarrow \mathbf{nil}) \times (\{c\} \bullet a \mapsto [l \cdot w]) \rrbracket \} \end{aligned}$$

Either specification (4) holds trivially, since the precondition is equivalent to false, or $c = a \mapsto [l_1 \cdot - \cdot l_2] * ls$ for some $l_1, l_2 \in \text{Lst}$ with w and u not in either l_1 or l_2 , and some

| | |
|--|--|
| $E.value \stackrel{\text{def}}{=} E$ | $E.next \stackrel{\text{def}}{=} E + 1$ |
| $x := \text{newNode}() \stackrel{\text{def}}{=} x := \text{alloc}(2)$ | $x := \text{newRoot}() \stackrel{\text{def}}{=} x := \text{alloc}(1)$ |
| $\text{disposeNode}(x) \stackrel{\text{def}}{=} \text{dispose}(x, 2)$ | $\text{disposeRoot}(x) \stackrel{\text{def}}{=} \text{dispose}(x, 1)$ |
| <pre> v := getHead(i) { local x in x := [i]; if x = nil then v := x else v := [x.value] } v := getTail(i) { local x, y in x := [i]; if x = nil then v := x else y := [x.next]; while y ≠ nil do x := y; y := [x.next] ; v := [x.value] } v := getNext(i, w) { local x in x := [i]; v := [x.value]; while v ≠ w do x := [x.next]; v := [x.value] ; x := [x.next]; if x = nil then v := x else v := [x.value] } </pre> | <pre> v := getPrev(i, w) { local x, y in x := [i]; v := [x.value]; if v = w then v := nil else while v ≠ w do y := x; x := [y.next]; v := [x.value] ; v := [y.value] } v := pop(i) { local x, y in x := [i]; if x = nil then v := x else y := [x.next]; [i] := y; v := [x.value]; disposeNode(x) } push(i, v) { local X, y in x := newNode(); [x.value] := v; y := [i]; [x.next] := y; [i] := x } </pre> |

Fig. 10. Linked-list-based list store implementation

```

remove(i, v) {
  local u, x, y, z in
    x := [i];
    u := [x.value];
    y := [x.next];
    if u = v then
      [i] := y;
      disposeNode(x)
    else
      u := [y.value];
      while u ≠ v do
        x := y;
        y := [x.next];
        u := [y.value] ;
      z := [y.next];
      [x.next] := z;
      disposeNode(y)
  }

i := newList() {
  i := newRoot();
  [i] := nil
}

```

```

insert(i, v, w) {
  local u, x, y, z in
    x := [i];
    u := [x.value];
    while u ≠ v do
      x := [x.next];
      u := [x.value] ;
    y := [x.next];
    z := newNode();
    [z.value] := w;
    [z.next] := y;
    [x.next] := z
  }

deleteList(i) {
  local x, y in
    x := [i];
    while x ≠ nil do
      y := x;
      x := [y.next];
      disposeNode(y) ;
    disposeRoot(i)
}

```

Fig. 11. Linked-list-based list store implementation

$ls \in \text{LStore}$. A proof outline for this case is given in Fig. 13. Similarly, either specification (5) holds trivially or $c = ls$ for some $ls \in \text{LStore}$. A proof outline for this case is given in Fig. 14. In both cases the `getNext` implementation performs the same search for the value w in the list. The proof outline for this common part is given in Fig. 12.

7. Conclusions

We have shown how to refine module specifications, given by abstract local reasoning, into correct implementations. This provides a justification for the soundness of abstract local reasoning with context algebras. We have identified two general approaches for proving the correctness of an implementation with respect to an abstract specification: locality-preserving and locality-breaking translations. Locality-preserving translations relate the abstract locality of a module with the low-level locality of its implementation. This is subtle since disjoint structures at the high-level are not quite disjoint at the low-level, because of the additional crust that is required to handle the pointer surgery. Locality-preserving translations thus establish a “fiction of disjointness” at the abstract level. Meanwhile, locality-breaking translations establish a “fiction of locality”, by justifying abstract locality even though this locality is not matched by the implementation.

We now conclude with some discussion of our choices and of the possible future extensions of our work.

$$\begin{array}{l}
\left\{ \exists p. i \Rightarrow a * w \Rightarrow w * v \Rightarrow - * x \Rightarrow - \times a \mapsto p * \langle\langle l_1 \cdot w \rangle\rangle^{(p,y)} \right\} \\
\mathbf{x} := [i]; \\
\left\{ \exists p. i \Rightarrow a * w \Rightarrow w * v \Rightarrow - * x \Rightarrow p \times a \mapsto p * \langle\langle l_1 \cdot w \rangle\rangle^{(p,y)} \right\} \\
\left. \begin{array}{l}
\left\{ \exists p. w \Rightarrow w * v \Rightarrow - * x \Rightarrow p \times \right. \\
\left. \left((l_1 = \emptyset \wedge p \mapsto w, y) \vee (\exists v, l'_1, q. l_1 = v \cdot l'_1 \wedge p \mapsto w, q * \langle\langle l'_1 \cdot w \rangle\rangle^{(q,y)}) \right) \right\} \right\} \\
\mathbf{v} := [\mathbf{x.value}]; \\
\left\{ \begin{array}{l}
\exists v, x, l'_1, l''_1, q. w \Rightarrow w * v \Rightarrow v * x \Rightarrow x \times \\
l_1 \cdot w = l'_1 \cdot v \cdot l''_1 \wedge \langle\langle l'_1 \rangle\rangle^{(p,x)} * x \mapsto v, q * \langle\langle l''_1 \rangle\rangle^{(q,y)}
\end{array} \right\} \\
\mathbf{while} \ v \neq w \ \mathbf{do} \\
\left\{ \begin{array}{l}
\exists v, v', x, x', l'_1, l''_1, q. w \Rightarrow w * v \Rightarrow v * x \Rightarrow x \times \\
l_1 \cdot w = l'_1 \cdot v \cdot v' \cdot l''_1 \wedge \langle\langle l'_1 \rangle\rangle^{(p,x)} * x \mapsto v, x' * x' \mapsto v', q * \langle\langle l''_1 \rangle\rangle^{(q,y)}
\end{array} \right\} \\
\mathbf{x} := [\mathbf{x.next}]; \\
\mathbf{v} := [\mathbf{x.value}] \\
\left\{ \begin{array}{l}
\exists v, v', x, x', l'_1, l''_1, q. w \Rightarrow w * v \Rightarrow v' * x \Rightarrow x' \times \\
l_1 \cdot w = l'_1 \cdot v \cdot v' \cdot l''_1 \wedge \langle\langle l'_1 \rangle\rangle^{(p,x)} * x \mapsto v, x' * x' \mapsto v', q * \langle\langle l''_1 \rangle\rangle^{(q,y)}
\end{array} \right\} \\
\left\{ \begin{array}{l}
\exists v, x, l'_1, l''_1, q. w \Rightarrow w * v \Rightarrow v * x \Rightarrow x \times \\
l_1 \cdot w = l'_1 \cdot v \cdot l''_1 \wedge \langle\langle l'_1 \rangle\rangle^{(p,x)} * x \mapsto v, q * \langle\langle l''_1 \rangle\rangle^{(q,y)}
\end{array} \right\} \ ; \\
\left\{ \exists x. w \Rightarrow w * v \Rightarrow w * x \Rightarrow x \times \langle\langle l_1 \rangle\rangle^{(p,x)} * x \mapsto w, y \right\} \\
\mathbf{x} := [\mathbf{x.next}] \\
\left\{ w \Rightarrow w * v \Rightarrow w * x \Rightarrow y \times \langle\langle l_1 \cdot w \rangle\rangle^{(p,y)} \right\} \\
\left\{ \exists p. i \Rightarrow a * w \Rightarrow w * v \Rightarrow w * x \Rightarrow y \times a \mapsto p * \langle\langle l_1 \cdot w \rangle\rangle^{(p,y)} \right\}
\end{array}$$

Fig. 12. Proof outline for search part of getNext implementation (common part)

7.1. On Locality-Preserving versus Locality-Breaking

Despite the fact that their names imply a black-and-white distinction between locality-preserving and locality-breaking translations, there is no clear distinction between where the two techniques are applicable.

As an example, consider the implementation of the list module from §6.3. We proved that this implementation gave a sound module translation using the locality-breaking technique, since some of the basic operations have a low-level footprint that incorporates an arbitrarily large portion of the linked list. However, it is quite reasonable to identify portions of the low-level state (individual nodes in a linked list) that correspond to portions of the high level state (individual elements of the corresponding list). In fact, the locality-preserving approach can be applied to this implementation, by treating the portion of the linked list that leads up to the nodes of interest as crust.

On the other hand, consider the implementation of the tree module from §5.5. In this case, the crust could be arbitrarily large, accounting for all of the siblings of the top level of the tree. This suggests that the locality-breaking approach is also applicable.

Clearly, there is a significant overlap in the applicability of the two approaches, but it turns out that there are no cases in which one approach is applicable but the other is not.

```

v := getNext(i, w) {
  { [[i ⇒ a * w ⇒ w * v ⇒ - × a ⇒ [l1 · w · u · l2] * ls]] }
  { ∃y, z, p. i ⇒ a * w ⇒ w * v ⇒ - × a ⇒ p * ⟨⟨l1 · w⟩⟩(p,y) * y ↦ u, z * ⟨⟨l2⟩⟩(z,nil) * (ls) }
  { ∃p. i ⇒ a * w ⇒ w * v ⇒ - × a ⇒ p * ⟨⟨l1 · w⟩⟩(p,y) * y ↦ u, z }
  local x in
    { ∃p. i ⇒ a * w ⇒ w * v ⇒ - * x ⇒ - × a ⇒ p * ⟨⟨l1 · w⟩⟩(p,y) * y ↦ u, z }
    (see Fig. 12)
    { ∃p. i ⇒ a * w ⇒ w * v ⇒ w * x ⇒ y × a ⇒ p * ⟨⟨l1 · w⟩⟩(p,y) * y ↦ u, z }
  if x = nil then
    v := x
  else
    v := [x.value]
    { ∃p. i ⇒ a * w ⇒ w * v ⇒ u * x ⇒ y × a ⇒ p * ⟨⟨l1 · w⟩⟩(p,y) * y ↦ u, z }
    { ∃p. i ⇒ a * w ⇒ w * v ⇒ u × a ⇒ p * ⟨⟨l1 · w⟩⟩(p,y) * y ↦ u, z }
    { ∃y, z, p. i ⇒ a * w ⇒ w * v ⇒ u × a ⇒ p * ⟨⟨l1 · w⟩⟩(p,y) * y ↦ u, z * ⟨⟨l2⟩⟩(z,nil) * (ls) }
    { [[i ⇒ a * w ⇒ w * v ⇒ u × a ⇒ [l1 · w · u · l2] * ls]] }
}

```

Fig. 13. Proof outline for getNext implementation (success case)

A locality-preserving translation can be used to construct a locality-breaking translation and vice versa.

From locality-preserving to locality-breaking is simple. Assume that $\tau : \mathbb{A} \rightarrow \mathbb{B}$ is a locality-preserving translation. For all $c \in \mathbb{C}_{\mathbb{A}}$, $\varphi \in \text{Cmd}_{\mathbb{A}}$ and $(P, Q) \in \text{AX} \llbracket \varphi \rrbracket_{\mathbb{A}}$, there is a derivation of $\vdash_{\mathbb{A}} \{\{\emptyset, c\} \bullet P\} \varphi \{\{\emptyset, c\} \bullet Q\}$ which simply uses AXIOM followed by FRAME. By the soundness of locality-preserving translations, there must therefore also be a derivation of $\vdash_{\mathbb{B}} \{\llbracket \{\emptyset, c\} \bullet P \rrbracket_{\tau}\} \llbracket \varphi \rrbracket_{\tau} \{\llbracket \{\emptyset, c\} \bullet Q \rrbracket_{\tau}\}$. Thus τ defines a locality-breaking translation.

From locality-breaking to locality-preserving is slightly trickier. Assume that $\tau : \mathbb{A} \rightarrow \mathbb{B}$ is a locality-breaking translation. For the interface sets, choose $\mathcal{I}_{\text{in}} = \{1\}$ and $\mathcal{I}_{\text{out}} = \mathbb{C}_{\mathbb{A}}$. Define the representation functions for data $\chi_{\mathbb{A}}$ and contexts c as follows:

$$\begin{aligned} \langle\langle \chi_{\mathbb{A}} \rangle\rangle^c &= \{\chi_{\mathbb{B}} \mid \chi_{\mathbb{B}} \alpha_{\tau} (c \bullet \chi_{\mathbb{A}})\} \\ \langle\langle c \rangle\rangle_{c_2}^{c_1} &= \begin{cases} \mathbf{I}_{\mathbb{B}} & \text{if } c_2 = c_1 \bullet c \\ \emptyset & \text{otherwise.} \end{cases} \end{aligned}$$

Finally, choose $\mathcal{F} = \{1\}$ and define $\mathfrak{m}_c = \mathbf{I}_{\mathbb{B}}$. Application preservation holds by construction, as does crust inclusion. Axiom correctness simply reduces to the criterion that τ is a locality-breaking translation, *i.e.* that the axioms, with each singleton frame, hold under translation. Thus τ defines a locality-preserving translation.

$$\begin{array}{l}
v := \text{getNext}(i, w) \{ \\
\quad \left\{ \llbracket i \Rightarrow a * w \Rightarrow w * v \Rightarrow - \times a \Rightarrow [l \cdot w] * ls \rrbracket \right\} \\
\quad \left\{ \exists p. i \Rightarrow a * w \Rightarrow w * v \Rightarrow - \times a \mapsto p * \langle\langle l \cdot w \rangle\rangle^{(p, nil)} * \langle\langle ls \rangle\rangle \right\} \\
\quad \left\{ \exists p. i \Rightarrow a * w \Rightarrow w * v \Rightarrow - \times a \mapsto p * \langle\langle l \cdot w \rangle\rangle^{(p, nil)} \right\} \\
\quad \text{local } x \text{ in} \\
\quad \quad \left\{ \exists p. i \Rightarrow a * w \Rightarrow w * v \Rightarrow - * x \Rightarrow - \times a \mapsto p * \langle\langle l \cdot w \rangle\rangle^{(p, nil)} \right\} \\
\quad \quad \quad \text{(see Fig. 12)} \\
\quad \quad \left\{ \exists p. i \Rightarrow a * w \Rightarrow w * v \Rightarrow w * x \Rightarrow \mathbf{nil} \times a \mapsto p * \langle\langle l \cdot w \rangle\rangle^{(p, nil)} \right\} \\
\quad \quad \text{if } x = \mathbf{nil} \text{ then} \\
\quad \quad \quad v := x \\
\quad \quad \text{else} \\
\quad \quad \quad v := [x.\text{value}] \\
\quad \quad \quad \left\{ \exists p. i \Rightarrow a * w \Rightarrow w * v \Rightarrow \mathbf{nil} * x \Rightarrow \mathbf{nil} \times a \mapsto p * \langle\langle l \cdot w \rangle\rangle^{(p, nil)} \right\} \\
\quad \quad \quad \left\{ \exists p. i \Rightarrow a * w \Rightarrow w * v \Rightarrow \mathbf{nil} \times a \mapsto p * \langle\langle l \cdot w \rangle\rangle^{(p, nil)} \right\} \\
\quad \quad \left\{ \exists p. i \Rightarrow a * w \Rightarrow w * v \Rightarrow \mathbf{nil} \times a \mapsto p * \langle\langle l \cdot w \rangle\rangle^{(p, nil)} * \langle\langle ls \rangle\rangle \right\} \\
\quad \quad \left\{ \llbracket i \Rightarrow a * w \Rightarrow w * v \Rightarrow \mathbf{nil} \times a \Rightarrow [l \cdot w] * ls \rrbracket \right\} \\
\quad \}
\end{array}$$

Fig. 14. Proof outline for getNext implementation (failure case)

Note that the transformation from locality-breaking to locality-preserving translations is not a perfect inverse of the transformation in the other direction. This is because locality-breaking transformations do not embody the notion of representing a substructure, and therefore it is necessary to define the representation functions by providing the *entire* enveloping context as the interface.

7.2. On Abstract Predicates

One way of viewing the translation functions is as abstract predicates: that is, $\llbracket P \rrbracket$ is an abstract predicate parametrised by P . However, viewing this as a completely abstract entity does not confer abstract local reasoning. Exposing such axioms as $\llbracket P \rrbracket \vee \llbracket Q \rrbracket \leftrightarrow \llbracket P \vee Q \rrbracket$ is a start, it allows the low-level disjunction rule to implement its high-level counterpart, and is possible with some formulations of abstract predicates (Dinsdale-Young et al., 2010a). However, abstract predicates do not currently provide a mechanism for exporting meta-theorems such as the soundness of the abstract frame rule. That is to say, there is no way to expose the fact that if $\{\{\llbracket P \rrbracket\}\} \mathbb{C} \{\{\llbracket Q \rrbracket\}\}$ holds then so does $\{\{\llbracket K \bullet P \rrbracket\}\} \mathbb{C} \{\{\llbracket K \bullet Q \rrbracket\}\}$. This work suggests that including such a mechanism could be a valuable addition to the abstract predicate methodology.

7.3. On Concurrency

Extending the results presented here to a concurrent setting is not entirely trivial. A separation logic-style parallel rule is only appropriate when the model has a commutative $*$ -like operator; while the heap and list modules have such an operator, the tree model does not. (The segment logic of Gardner and Wheelhouse (2009) is a promising approach to introducing a commutative $*$ to structured models such as trees.)

Assume that we have a sound module translation from an abstract module to a concrete one, and that both models include a commutative $*$ connective, such that, for every state d there is some context c such that, for all states d' , $d * d' = c \bullet d'$. If the implementations of the module operations behave as atomic operations, then the separation logic parallel rule should be sound at the abstract level. That is, if

$$\begin{aligned} & \{[P_1]\} [C_1] \{[Q_1]\} \\ & \{[P_2]\} [C_2] \{[Q_2]\} \end{aligned}$$

then

$$\{[P_1 * P_2]\} [C_1 \parallel C_2] \{[Q_1 * Q_2]\}.$$

Intuitively, this is because the operations of the threads are effectively interleaved, and so the state of the other thread at each point can be viewed as a frame.

An alternative approach is that of Dinsdale-Young, Dodds, Gardner, Parkinson, and Vafeiadis (2010a) on *concurrent abstract predicates (CAP)*. CAP presents a fiction of disjointness for concurrent program modules by building abstractions on top of a powerful permission system. This makes it possible to express sophisticated notions of abstract resource, akin to those we consider here. Ultimately, we hope to reach a unified theory of abstraction and refinement that supports fictions of locality and disjointness in both the sequential and concurrent settings.

References

- Cristiano Calcagno, Philippa Gardner, and Uri Zarfaty. Context Logic and tree update. In *POPL '05: Proceedings of the 32nd ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 271–282, New York, NY, USA, 2005. ACM Press. ISBN 1-58113-830-X. .
- Cristiano Calcagno, Philippa Gardner, and Uri Zarfaty. Local reasoning about data update. *Electron. Notes Theor. Comput. Sci.*, 172:133–175, 2007a. ISSN 1571-0661. .
- Cristiano Calcagno, Peter W. O'Hearn, and Hongseok Yang. Local action and abstract separation logic. In *LICS '07*, pages 366–378, Washington, DC, USA, 2007b. IEEE Computer Society.
- Willem-Paul de Roever and Kai Engelhardt. *Data Refinement: Model-Oriented Proof Methods and Their Comparison*. Cambridge University Press, Cambridge, UK, 1999. ISBN 0521641705.
- Thomas Dinsdale-Young, Mike Dodds, Philippa Gardner, Matthew J. Parkinson, and Viktor Vafeiadis. Concurrent abstract predicates. In Theo D'Hondt, editor, *ECOOP*,

- volume 6183 of *Lecture Notes in Computer Science*, pages 504–528. Springer, 2010a. ISBN 978-3-642-14106-5.
- Thomas Dinsdale-Young, Philippa Gardner, and Mark Wheelhouse. Abstraction and refinement for local reasoning. Technical report, Imperial College London, 2010b. URL <http://www.doc.ic.ac.uk/~td202/papers/alrfull.pdf>.
- Philippa Gardner and Mark J. Wheelhouse. Small specifications for tree update. In Cosimo Laneve and Jianwen Su, editors, *WS-FM*, volume 6194 of *Lecture Notes in Computer Science*, pages 178–195. Springer, 2009. ISBN 978-3-642-14457-8.
- Philippa A. Gardner, Gareth D. Smith, Mark J. Wheelhouse, and Uri D. Zarfaty. Local hoare reasoning about dom. In *PODS '08*, pages 261–270, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-108-8. .
- Maurice P. Herlihy and Jeannette M. Wing. Linearizability: a correctness condition for concurrent objects. *TOPLAS*, 12(3):463–492, 1990. ISSN 0164-0925. .
- C. A. R. Hoare. Proof of correctness of data representations. *Acta Inf.*, 1(4):271–281, 1972.
- Samin S. Ishtiaq and Peter W. O’Hearn. BI as an assertion language for mutable data structures. In *POPL 2001*, New York, 2001. ACM Press.
- Ivana Mijajlović, Noah Torp-Smith, and Peter W. O’Hearn. Refinement and separation contexts. In *FSTTCS '04*, pages 421–433, Berlin/Heidelberg, Germany, 2004. Springer.
- Peter W. O’Hearn, John Reynolds, and Hongseok Yang. Local reasoning about programs that alter data structures. In *CSL*, Lecture Notes in Computer Science. Springer, 2001.
- M. J. Parkinson and G. M. Bierman. Separation logic and abstraction. In *POPL*, pages 247–258, 2005.
- Mohammad Raza and Philippa Gardner. Footprints in local reasoning. *Logical Methods in Computer Science*, 5(2), 2009.
- John C. Reynolds. Separation Logic: A logic for shared mutable data structures. In *LICS '02: Proceedings of the 17th Annual IEEE Symposium on Logic in Computer Science*, pages 55–74, Washington, DC, USA, 2002. IEEE Computer Society. ISBN 0-7695-1483-9.