

A Compositional Theory for STM Haskell

Johannes Borgström Karthikeyan Bhargavan Andrew D. Gordon

Microsoft Research, Cambridge, UK
{jborg,karthb,adg}@microsoft.com

Abstract

We address the problem of reasoning about Haskell programs that use Software Transactional Memory (STM). As a motivating example, we consider Haskell code for a concurrent non-deterministic tree rewriting algorithm implementing the operational semantics of the ambient calculus. The core of our theory is a uniform model, in the spirit of process calculi, of the run-time state of multi-threaded STM Haskell programs. The model was designed to simplify both local and compositional reasoning about STM programs. A single reduction relation captures both pure functional computations and also effectful computations in the STM and I/O monads. We state and prove liveness, soundness, completeness, safety, and termination properties relating source processes and their Haskell implementation. Our proof exploits various ideas from concurrency theory, such as the bisimulation technique, but in the setting of a widely used programming language rather than an abstract process calculus. Additionally, we develop an equational theory for reasoning about STM Haskell programs, and establish for the first time equations conjectured by the designers of STM Haskell. We conclude that using a pure functional language extended with STM facilitates reasoning about concurrent implementation code.

Categories and Subject Descriptors D.2.4 [Software/Program Verification]: Correctness Proofs; D.3.1 [Formal Definitions and Theory]: Syntax and Semantics; D.3.3 [Language Constructs and Features]: Concurrent Programming Structures—Software Transactional Memory.

General Terms Theory, verification

Keywords Transactional memory, compositional reasoning, ambient calculus

1. Introduction

Software Transactional Memory (STM), introduced by Shavit and Touitou [31], is a promising programming abstraction for shared-variable concurrency. Shared variables may only be accessed within transactions. Ingenious software techniques allow transactions to run in parallel while their semantics is as if they run in series. There is a good deal of promising research on efficient implementations, in the context of various languages [9, 10, 27]. Moreover, several formal techniques have been applied to verifying the underlying algorithms [30] and their implementations [22, 8, 20, 1].

In this paper, we explore the prospects for reasoning about software written using the STM abstraction. Transactional semantics undoubtedly simplifies the reasoning task compared to say lock-based concurrency [18, 23], but is no panacea.

We pursue the idea that theories of concurrency developed in abstract process calculi can fruitfully be recast in the concrete setting of transactional programming languages. We consider programs written in STM Haskell [10, 11], an embedding of transactional concurrency within the pure functional language Haskell.

As a concrete programming task, we investigate programming, specifying, and reasoning about an STM Haskell implementation of ambients. The ambient calculus [5] is a formalism for expressing and reasoning about mobile computation, more recently applied to biological systems [26]. An ambient process has a hierarchical structure that mutates over time, as ambients within the structure move inside and outside other ambients. Hence, an implementation of the ambient calculus amounts to a concurrent tree rewriting algorithm. The first concurrent implementation, by Cardelli in Java [4], was lock-based; here, we give a lock-free implementation of a programming API for the ambient calculus in STM Haskell.

As a basis for reasoning about this code, we present a core calculus for STM Haskell: a concurrent non-strict lambda calculus with transactional variables and atomic blocks. The syntax of lambda calculus expressions provides a compositional and uniform formalism for both source programs and their run-time states.

The original presentation of STM Haskell separated the transactional heap, its latest checkpoint and the currently running code. In contrast, our syntax uniformly represents all of these as expressions, hence facilitating compositional reasoning by allowing multiple threads with associated pieces of the heap to be composed using a parallel operator and rearranged using structural congruence.

We develop a formal semantics and type system for the calculus. Our semantics is based on a reduction relation \rightarrow on expressions, which range over both effect-free functional computations and effectful concurrent computations. It facilitates local reasoning by having transactions run against a part of the heap, making it easy to exhibit the smallest heap allowing a given transaction to make progress. On the other hand, we can show a precise correspondence with the original (but syntactically more complex) operational semantics of STM Haskell.

We import into STM Haskell behavioural equivalences and proof techniques that originate in process calculi. Notably, in our proofs, the notion of bisimulation up to a relation [19, 28] permits significant reductions in the size of the state space that needs to be considered when reasoning equationally about a program. This state space reduction comes beyond the already significant reduction afforded by the atomicity and serializability guarantees of the STM abstraction.

Using these proof techniques, our first result, Theorem 1, directly relates the operational semantics of ambient processes with expressions representing the run time states of their Haskell implementation. More specifically, Theorem 1 establishes liveness,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Haskell'09, August 30, 2009, Edinburgh, Scotland, UK.
Copyright © 2009 ACM 978-1-60558-508-6/09/09...\$5.00

soundness, completeness, safety, and termination properties. Building on the first theorem, our second result, Theorem 2, is a concise statement of full correctness of our implementation in terms of a bisimulation between the source process and its Haskell translation. The use of transactions makes the Haskell code rather simpler than Cardelli’s original Java program; still, it is non-trivial but feasible to establish the intended correspondences between the implementation code and the formal specification of ambients.

Finally, we adopt the standard definition of Morris-style contextual equivalence of expressions, and develop a sound equational theory. We show the monad laws of the STM monad hold, and also establish some equations proposed by the designers of STM Haskell [10].

Our main contributions are the following:

- We use notions of behavioural equivalence and proof techniques from process calculi to specify and prove STM Haskell code for a complex concurrent tree rewriting algorithm (an implementation of the ambient calculus).
- We develop uniform syntax and reduction semantics, in the style of process calculi, for STM Haskell. The uniformity of the syntax facilitates compositional reasoning over multiple threads with associated pieces of the heap.
- We prove soundness of an equational theory for STM Haskell, including monad laws and properties of operators for transactional control flow.

Outline of the Paper Section 2 introduces the formal syntax of core STM Haskell, and reviews the programming model. (Since our motivating example does not use exceptions, we omit them from the core calculus. The extended version of the paper shows how to extend our core calculus with exceptions.) Section 3 describes our code and data structures for programming ambients in Haskell. Section 4 completes our formalization of STM Haskell; we define the operational semantics and type system, and compare with the original semantics. Section 5 recalls the semantics of the ambient calculus, and states and proves the expected correctness properties; the proof depends on a detailed correspondence between ambient processes and expressions representing the corresponding states of our implementations. Section 6 develops an equational theory based on our operational semantics. Section 7 describes related work, and Section 8 concludes, and discusses potential future work.

An extended version of this paper, with additional explanations and proofs, and code listings, is available [3].

2. A Core Calculus for STM Haskell

The GHC implementation of STM [10] uses a monad to encapsulate all accesses to shared transactional variables (TVars). To execute an STM transaction we use the function `atomically`, which has type $\text{STM } t \rightarrow \text{IO } t$. If the execution of a transaction `returns` a result, the run-time system guarantees that it executed independently of all other transactions in the system. An STM expression can also `retry`, which may cause the transaction to be rolled back and run at some later time.

The original definition of STM Haskell made use of an implicit functional core language. For verification purposes, and in order to make this paper self-contained, we formalize STM Haskell as a concurrent non-strict lambda calculus, with memory cells (TVars) and atomic blocks.

2.1 Syntax

Our syntax treats source programs, heaps made up of transactional variables (TVars), and concurrent threads uniformly as expressions in the calculus, similarly to Concurrent Haskell [24].

We assume denumerable distinct sets \mathbf{X} of (lambda calculus) variables and \mathbf{N} of (TVar) addresses; we let x, y, z range over \mathbf{X} , and a, b range over \mathbf{N} . We let f range over ADT constructors.

Expressions of Core STM Haskell:

$M, N \in \mathbf{M} ::=$	expression
x	$(x \in \mathbf{X})$ variable (value)
a	$(a \in \mathbf{N})$ address (value)
$\lambda x.M$	lambda abstraction (value)
$f \bar{M}$	construction (value)
$M N$	application
$\text{case } M \text{ of } \overline{f \bar{x}} \rightarrow N$	case expression
$Y M$	fixpoint
$\text{equal } M N$	address equality
$\text{readTVar } M$	STM read variable
$\text{writeTVar } M N$	STM write variable
$\text{return}_{\text{STM}} M$	STM return
retry	STM retry transaction
$M \gg_{\text{STM}} N$	STM bind
$\text{orElse } M N$	STM prioritized choice
$\text{or } M N$	STM erratic choice
$\text{atomically } M$	IO execute transaction
$\text{return}_{\text{IO}} M$	IO return
$M \gg_{\text{IO}} N$	IO bind
$a \mapsto M$	transactional variable (TVar)
$(\nu a)M$	restriction (a bound in M)
$M N$	parallel composition
emp	empty heap

Our syntax is close to STM Haskell, but with minor differences. The STM Haskell functions `newTVar` and `fork` are not primitive expressions in our calculus, but are derived forms, explained below. In actual source programs, the monadic bind and return operators do not have the subscripts IO and STM; instead, the monads are inferred by the typechecker. The syntax $a \mapsto M$, $(\nu a)M$, $M | N$, `emp`, and addresses a exist as expressions only to represent run-time state, and cannot be written directly in source programs. Actual Haskell includes various abbreviations such as recursive definitions, pattern matching, and do-notation for monads, which can be reduced in standard ways to our core syntax.

We introduce some notational conventions. We write \bar{M} as shorthand for a possibly empty sequence $M_1 \cdots M_n$ (and similarly for \bar{x} , \bar{t} , etc.) and $\overline{f \bar{x}} \rightarrow N$ for a non-empty sequence $f_1 \bar{x}_1 \rightarrow N_1 \mid \cdots \mid f_m \bar{x}_m \rightarrow N_m$ where the f_i as well as the x_{ij} are assumed to be pairwise different (and similarly for $\overline{f \bar{t}}$). We write the empty sequence as \circ and denote concatenation of sequences using a comma. The length of a sequence \bar{x} is written $|\bar{x}|$. If ϕ is a phrase of syntax (such as an expression), we let $\text{fv}(\phi)$ and $\text{fn}(\phi)$ be the sets of variables and names occurring free in ϕ . We write $\phi \{M/x\}$ for the outcome of the capture-avoiding substitution of M for each free occurrence of x in ϕ .

We let $g \bar{M}$ range over the *builtin applications*, which are the expressions listed above starting with $Y M$ and ending with $M \gg_{\text{IO}} N$. (This notation is particularly useful in the typing rule (T BUILTIN) in Section 4.2.) If $g \bar{M}$ is a builtin application, we say g is a builtin function; Y and \gg_{IO} are examples. To simplify our reduction semantics, builtin functions are not themselves expressions, and only occur fully applied to their arguments. As usual, we can use lambda abstraction to represent unapplied or partially applied builtin functions as expressions.

2.2 Informal Semantics

Our uniform syntax for expressions covers single-threaded functional computations, as well as heaps, imperative transactions, and

concurrent computations. We describe each of these in turn, together with their semantics.

Functional Computations The core of our expression language is a call-by-name lambda calculus with algebraic data types. The simplest expressions are *values*, which return at once. As well as variables x , there are three kinds of value: names a representing the address of a mutable variable in the heap, lambda abstractions $\lambda x.M$, and constructions $f \bar{M}$ representing data tagged with the algebraic type constructor f . For example, `True`, `False`, `Nil`, and `Cons M1 M2` are constructions. Since we describe a non-strict language, constructor arguments \bar{M} may be general expressions and not just values.

The other expressions of the lambda calculus core are as follows. An application expression $M N$ evaluates to $M' \{N/x\}$, if evaluation of M returns a function $\lambda x.M'$. A case expression `case M of $\bar{f} \bar{x} \rightarrow \bar{N}$` attempts to match one of the clauses $\bar{f} \bar{x} \rightarrow \bar{N}$ against M . If $f_j \bar{x}_j \rightarrow N_j$ is the first of these clauses such that the value of M is $f_j \bar{M}$ for some \bar{M} , the whole expression behaves as $N_j \{ \bar{M} / \bar{x}_j \}$; if there is no such clause, the whole expression is stuck. A builtin application $Y M$ evaluates to $M (Y M)$. (We use Y to explain recursive definitions in the standard way.) A builtin application `equal M N` evaluates both M and N , and returns `True` if they evaluate to the same address, and `False` otherwise.

Heaps Next, we describe *heap-expressions*, which consist of a possibly empty composition of transaction variables, or TVars, for short. An expression $a \mapsto M$ denotes a heap consisting of a single TVar, with address a and current content M . The expression `emp` represents the empty heap, and the composition $M | N$ represents the concatenation of heaps M and N .

More generally, parallel composition $M | N$ represents M and N running in parallel, where M and N may be a mixture of heaps, STM-expressions, and IO-expressions. Unusually, the $|$ operator is not fully commutative; the result of $M | N$, if any, is the result of N . To evaluate a restriction $(\nu a)M$ one creates a fresh address a and then evaluates M ; as in process calculi, a restriction may also be read declaratively as meaning that the address a is known only within its scope M .

Imperative Transactions An STM-expression typically has the form $H | M$, where H is a heap-expression and M is a running (single-threaded) *transaction*. Transactions are composed from reads (`readTVar a`) and writes (`writeTVar a M`) to TVars. In source programs, we create TVars by the following abbreviation:

`newTVar M := (νa)(a ↦ M | returnSTM a)` where $a \notin \text{fn}(M)$

A transaction may `returnSTM` a result M (to commit), or `retry` (to roll back any updates and restart). There are two different kinds of choice. Prioritized choice `M orElse N` behaves as M unless M does a `retry` (in which case it behaves as N). Erratic choice `or M N` behaves as M or as N nondeterministically.

Concurrency Finally, an IO-expression typically has the form $H | M_1 | \dots | M_n$, where H is a heap-expression and M_1, \dots, M_n are *threads* running in parallel. A thread M_i may run a transaction to have side-effects on the shared heap H . Each thread may eventually terminate by returning a result. In our source programs, we create parallel threads using the abbreviation:

`fork M := (M | returnIO ())`

A thread `atomically M` behaves as if the transaction M executes in isolation against the heap, in one big atomic step. There are two possibilities to consider. If M returns a result M' , the updates to the heap are committed, and the whole expression returns M' . If M retries, any updates to the heap are rolled back, and the whole expression snaps back to `atomically M`, ready to try again.

Both IO and STM computations can be sequenced using the bind operator $M \gg_{=IO} N$, which behaves as $N M'$ if M returns M' , and otherwise as M .

For an example, consider a function `swap` that swaps the values of two TVars:

`swap := λxa.λxb.readTVar xa >>=STM λy.readTVar xb >>=STM λz.writeTVar xa z >>=STM λw.writeTVar xb y`

Here, `swap` takes two TVar arguments x_a and x_b and calls the builtin function `readTVar` to read their values. As discussed (infix) function $\gg_{=STM}$ binds the return value of the expression to its left to the function on its right. Hence, y and z get the values of x_a and x_b , respectively; `swap` then calls the function `writeTVar` to write x_a and x_b with z and y , respectively, thus completing the swap. To call the function with two TVars a, b , we can write `swap a b`. However, two parallel executions of `swap` could yield an inconsistent state. Instead, we call the function and have it execute atomically, by writing `atomically (swap a b)`.

3. Application: An API for Managing Ambients

As a verification challenge, we consider a sizeable program written in STM Haskell. Our target application is an API that implements the mobility primitives of the ambient calculus [5]. Programs that use this API spawn multiple threads that concurrently modify a shared data structure. We consider such programs to be exemplary applications of STM Haskell.


Our implementation is inspired by Cardelli's earlier implementation in Java using locks [4], which, although carefully specified, was never proved correct. A more recent programming comparison between a lock-based and STM-based implementation of ambients finds the "STM implementation to be much easier to reason about and much faster to implement" [32]. We seek to evaluate whether an STM implementation is also easy to verify, by developing the first proof for an implementation of the ambient calculus.

A Tree of TVars. The underlying data structure is a tree of named ambients, implemented as follows. The figure on the right depicts an example tree with a node named a that has two children named b and c .

```

type Ambient = TVar AmbData
data AmbData =
  AD (Name,
     Maybe Handle,
     [Handle],
     [(Name, Ambient)])
type Name = TVar String
type Handle = TVar Ambient

```



An Ambient (depicted as a named rectangle in the figure) is a TVar containing an `AmbData`, which consists of four items: a `Name`, a `Handle` pointing to its parent node (if it has one), a list of all incoming handles pointing to the current node, and an association list of child nodes, mapping node names to `Ambients`. A `Name` is an identifier (depicted as an italicized variable), implemented as a transactional variable (TVar) containing a string; the same `Name` may appear on multiple nodes. A `Handle` (depicted as a circle) is a pointer to an Ambient; in the figure, such pointers are depicted as dashed arrows. The additional level of indirection using `Handles` is used when merging a child node with its parent; all handles pointing to the child can then simply be pointed over to its parent.

We illustrate this data structure through a few simple functions that are used in our subsequent development. These functions can only be called from within STM atomic blocks.

```
readAmb :: Handle -> STM AmbData
readAmb h = do { a <- readTVar h;
                ad <- readTVar a;
                return ad}
```

The function `readAmb` takes a `Handle h` and returns the `AmbData` it points to (through two levels of indirection). Since `readAmb` reads transactional variables using `readTVar`, the result is an STM action.

```
writeAmb :: Handle -> AmbData -> STM ()
writeAmb h ad = do { a <- readTVar h;
                   writeTVar a ad }
```

The function `writeAmb` both reads and writes transactional variables; it writes an `AmbData` to the location pointed to by a handle.

```
parentOf :: Handle -> STM Handle
parentOf a = do { AD (_,p,_,_) <- readAmb a;
                case p of
                  Nothing -> retry;
                  Just ph -> return ph}
```

The function `parentOf` takes a `Handle h` pointing to a node and returns a handle to the parent of the node; if there is no parent (that is, `h` points to the root node), then it calls the STM `retry` function to rollback the transaction and restart.


Mobile Agents. An agent represents an IO thread that is initialized with a handle to an ambient. We say that the agent executes within that ambient.

```
data AGENT a = Agent (Handle -> IO a)
type Agent = AGENT ()
nil :: Agent
root :: Agent -> IO ()
new :: String -> AGENT Name
amb :: Name -> Agent -> Agent
```

The simplest agent is `nil`, which denotes an inactive agent that has finished its computation. The function `root` creates a fresh (unnamed) root node and attaches an agent `A` to this node. The agent `A` can then create subtrees by calling `new` to generate fresh node names, and calling `amb` to create new child nodes.

Using these functions, we can now create our example tree and attach agents to each node:

```
ex = root $ do {
  a <- new "a";
  b <- new "b";
  c <- new "c";
  amb a (do
    amb b (into c);
    amb c nil)}
```




Here, the agent attached to `b` uses a new function `into`:

```
into :: Name -> Agent
```

The figure treats the more general case when the agent performs other actions `e` after calling `into`. However, for simplicity, it does not depict `nil` processes or back pointers to handles.

When the agent `into c` is executed, it has the effect of moving the (subtree rooted at) node `b` to become a child of node `c`. The resulting tree is depicted on the right. If there is no neighbouring tree labelled `c`, the operation blocks until one becomes available.



As usual, when concurrent threads modify the tree in this way, there is a risk of the tree ending up in an inconsistent state. Our implementation of `into` in STM Haskell below uses the STM construct `atomically` to avoid inconsistency:

```
into c = Agent $ \bHandle -> atomically $ do {
  bAmbient <- readTVar bHandle;
  AD (b,bp,bh,bc) <- readAmb bHandle;
  aHandle <- parentOf bHandle;
  AD (a,ap,ah,ac) <- readAmb aHandle;
  let bSiblings = delete (b,bAmbient) ac in do
  cAmbient <- lookup' c bSiblings;
  AD (_,cp,ch,cc) <- readTVar cAmbient;
  let cHandle = head ch in do
  writeAmb cHandle (AD (c, cp, ch,
                       (b,bAmbient):cc));
  writeAmb aHandle (AD (a, ap, ah, bSiblings));
  writeAmb bHandle (AD (b,Just cHandle,bh,bc))}
```

The function `into` takes as argument the name `c` of the target ambient and creates an agent parameterized by a handle `bHandle` to the source ambient named `b`. The function proceeds in two phases. First, it reads the values of three ambient nodes: `b`, `b`'s parent `a`, and some sibling of `b` named `c`. Then, it writes updated values to all three nodes.

The function begins by reading the ambient at `b`, `bAmbient`, and calls `readAmb` to read its contents. It calls `parentOf` to find the parent ambient at `a` and reads its contents, including the list of its children `ac`. It computes the siblings of `b` by deleting `(b,bAmbient)` from the association list `ac`. It then finds a target ambient `cAmbient` by calling `lookup'` which non-deterministically chooses some sibling ambient with name `c`. (This non-determinism is motivated by the desired correctness properties of Section 5.) It reads the contents of `cAmbient`, including its children `cc` and one of its handles, `cHandle`.


Finally, the function updates the ambient at `c` by adding a child `(b,bAmbient)` to `ch`, it updates the ambient at `a` by deleting the child `b` from `ac`, and it updates the ambient at `b` by changing its parent to `c`.

Note that `into` is a local operation that only modifies three nodes of the tree; Agents manipulating other parts of the tree can be scheduled to run in parallel without causing any conflicts.

The Full API. The full Ambient API consists of several other functions:

```
out :: Name -> Agent
open :: Name -> Agent
fork :: Agent -> Agent
```

The agent `out c` is the inverse of `into c`; it moves an ambient out of its parent (if the parent is named `c`). The agent `open c` deletes a child node named `c` and swings all handles of `c` to point to its parent. This has the effect of causing all of `c`'s children to become children of the parent; all agents running on `c` are similarly affected. The figure below depicts its effect on an example graph.



The agent `fork A` forks off a new thread running the agent `A` within the same ambient.

Programmatically, agents form a Reader monad, where the value `read` is a handle to the location in which the agent is running¹. The $\gg=$ and \gg operators have their standard definition: $(\text{Reader } f) \gg= g$ running at a location with handle h evaluates $f h$ to some v , evaluates $g v$ to some $\text{Reader } g'$, and then evaluates $g' h$. Similarly, $(\text{Reader } f) \gg (\text{Reader } g)$ reading handle h evaluates $f h$, discards the result, and then evaluates $g h$.

```
instance Monad AGENT where
  return a = Agent $ \s -> return a
  a >>= g = Agent $ \s -> case a of Agent f -> f s
    >>= \v -> case (g v) of Agent ff -> (ff s)
  a >> b = Agent $ \s -> case a of Agent f -> f s
    >>= \v -> case b of Agent ff -> (ff s)
```

When verifying the ambient API, we are interested in establishing full functional correctness, not only the preservation of certain invariants of the location tree. To do this, we need to give a formal account of the semantics of our core calculus for STM Haskell.

4. The Core Calculus, Concluded

This section concludes the definition of our core calculus, begun in Section 2. We define the operational semantics and type system, and make a comparison with the original semantics. In the next section, we apply the calculus to specifying and verifying the Haskell code from Section 3.

4.1 Operational Semantics

We define a *reduction relation*, $M \rightarrow M'$, which specifies the run time behaviour of STM programs. A single reduction relation captures pure functional computation, imperative transactions, and concurrency. We rely on some auxiliary notions to define reduction. First, we define three kinds of evaluation contexts.

Contexts: Pure (\mathcal{R}_β), **parallel** (\mathcal{R}_\parallel), and **transactional** (\mathcal{R}_{\mapsto})

$$\begin{aligned} \mathcal{R}_\beta &:: [\cdot] \mid \mathcal{R}_\beta M \mid \text{case } \mathcal{R}_\beta \text{ of } f \bar{x} \rightarrow \bar{N} \mid \text{equal } \mathcal{R}_\beta M \\ &\quad \mid \text{equal } a \mathcal{R}_\beta \mid \text{readTVar } \mathcal{R}_\beta \mid \text{writeTVar } \mathcal{R}_\beta M \\ \mathcal{R}_\parallel &:: [\cdot] \mid (va)\mathcal{R}_\parallel \mid (\mathcal{R}_\parallel \mid M) \mid (M \mid \mathcal{R}_\parallel) \mid \\ &\quad (\mathcal{R}_\parallel \gg_{\text{IO}} M) \mid (\mathcal{R}_\parallel \gg_{\text{STM}} M) \\ \mathcal{R}_{\mapsto} &:: [\cdot] \mid (va)\mathcal{R}_{\mapsto} \mid (a \mapsto M \mid \mathcal{R}_{\mapsto}) \end{aligned}$$

The second auxiliary notion is *structural equivalence*, $M \equiv M'$. The purpose of this relation is to re-arrange the structure of an expression—for example, by pulling restrictions to the top, or by moving TVars beside reads or writes—so as to enable reduction steps. Structural equivalence is the least equivalence relation closed under the following rules. Let $\text{bn}(\mathcal{R}_\parallel)$ be the names bound by the context \mathcal{R}_\parallel , and let $\text{n}(\mathcal{R}_\parallel) = \text{bn}(\mathcal{R}_\parallel) \cup \text{fn}(\mathcal{R}_\parallel)$.

Structural Equivalence: $M \equiv N$

$$\begin{aligned} M &\equiv \text{emp} \mid M && (\text{STRUCT EMP}) \\ M \mid \mathcal{R}_\parallel[N] &\equiv \mathcal{R}_\parallel[M \mid N] \text{ if } \text{bn}(\mathcal{R}_\parallel) \cap \text{fn}(M) = \emptyset && (\text{STRUCT FLOAT}) \\ \mathcal{R}_\parallel[(va)M] &\equiv (va)\mathcal{R}_\parallel[M] \text{ if } a \notin \text{n}(\mathcal{R}_\parallel) && (\text{STRUCT RES CTX}) \\ M \equiv N &\Rightarrow \mathcal{R}_\parallel[M] \equiv \mathcal{R}_\parallel[N] && (\text{STRUCT CTX}) \end{aligned}$$

Let *reduction*, $M \rightarrow M'$, be the least relation closed under the rules in groups (R1), (R2), and (R3) displayed below. The first group consists of standard rules for functional and concurrent computation.

(R1) Reductions without Side-Effects: $M \rightarrow M'$

$$\begin{aligned} &(\lambda x.M) N \rightarrow M\{N/x\} && (\text{BETA}) \\ &\text{case } f_j(\bar{M}) \text{ of } \overline{f \bar{x}} \rightarrow \bar{N} \rightarrow N_j\{\bar{M}/\bar{x}_j\} && (\text{CASE}) \\ &Y M \rightarrow M (Y M) && (\text{FIX}) \\ &\text{equal } a a \rightarrow \text{True} && (\text{EQUAL TRUE}) \\ &\text{equal } a b \rightarrow \text{False} \text{ if } a \neq b && (\text{EQUAL FALSE}) \\ &(\text{return}_{\text{IO}} M \gg_{\text{IO}} N) \rightarrow N M && (\text{IO BIND RETURN}) \\ & && (\text{STRUCT}) \\ & \frac{}{M \equiv N} && (\text{PURE CTX}) \\ & \frac{}{M \equiv N} && (\text{RED CTX}) \\ & \frac{M \rightarrow M'}{\mathcal{R}_\beta[M] \rightarrow \mathcal{R}_\beta[M']} && \frac{M \rightarrow M'}{\mathcal{R}_\parallel[M] \rightarrow \mathcal{R}_\parallel[M']} && \frac{M \equiv N \quad N \rightarrow N' \quad N' \equiv M'}{M \rightarrow M'} \end{aligned}$$

The second group of reduction rules concerns the core behaviour of STM-expressions. A heap-expression H is a parallel composition of transactional variables

$$\Pi_i(a_i \mapsto M_i) := a_1 \mapsto M_1 \mid \dots \mid a_n \mapsto M_n \mid \text{emp}$$

where the a_i are pair-wise distinct. We write \rightarrow^* for the transitive closure of \rightarrow .

(R2) Core Reductions for STM Transactions: $M \rightarrow M'$

$$\begin{aligned} &(\text{STM READ TVar}) \\ &(a \mapsto M) \mid \text{readTVar } a \rightarrow (a \mapsto M) \mid \text{return}_{\text{STM}} M \\ &(\text{STM WRITE TVar}) \\ &(a \mapsto M) \mid \text{writeTVar } a M' \rightarrow (a \mapsto M') \mid \text{return}_{\text{STM}} () \\ &(\text{return}_{\text{STM}} M \gg_{\text{STM}} N) \rightarrow N M && (\text{STM BIND RETURN}) \\ &(\text{retry} \gg_{\text{STM}} N) \rightarrow \text{retry} && (\text{STM BIND RETRY}) \\ &(\text{ATOMIC RETURN}) \\ & \frac{H \mid M \rightarrow^* \mathcal{R}_{\mapsto}[\text{return}_{\text{STM}} N]}{H \mid \text{atomically } M \rightarrow \mathcal{R}_{\mapsto}[\text{return}_{\text{IO}} N]} \end{aligned}$$

(STM READ TVar) and (STM WRITE TVar) allow transactional variable to be read and written within a transaction.

(STM BIND RETURN) allows return values to propagate through the STM bind operator, much as through the IO bind operator, while (STM BIND RETRY) allows retry to propagate directly through the bind operator, much like an exception.

The rule (ATOMIC RETURN) turns a successful many-step transaction of an STM-expression $H \mid M$ into a single-step computation of the IO-expression $H \mid \text{atomically } M$. If the transaction yields `retry` then (ATOMIC RETURN) is not applicable, so there is no transition in this case. In the STM Haskell implementation, a transaction that `retrys` is aborted by the run-time system and queued for later execution.

The final group of rules concerns choices within transactions.

(R3) Reductions for OrElse and Or: $M \rightarrow M'$

$$\begin{aligned} &(\text{STM ORElse RETURN}) \\ & \frac{H \mid N_1 \rightarrow^* \mathcal{R}_{\mapsto}[\text{return}_{\text{STM}} N'_1]}{H \mid (N_1 \text{ orElse } N_2) \rightarrow \mathcal{R}_{\mapsto}[\text{return}_{\text{STM}} N'_1]} \\ &(\text{STM ORElse RETRY}) \\ & \frac{H \mid N_1 \rightarrow^* \mathcal{R}_{\mapsto}[\text{retry}]}{H \mid (N_1 \text{ orElse } N_2) \rightarrow H \mid N_2} \\ &M \text{ or } N \rightarrow M && (\text{STM OR LEFT}) \\ &M \text{ or } N \rightarrow N && (\text{STM OR RIGHT}) \end{aligned}$$

Rules (STM OrElse Return) and (STM OrElse Retry) formalize the idea that $N_1 \text{ orElse } N_2$ behaves as N_1 if N_1 terminates with `returnSTM N1'`. If N_1 terminates with `retry` then its effects are discarded, and we instead run N_2 on the original heap H .

¹The Haskell programmer familiar with monad transformers will notice that it is effectively a `ReaderT Handle IO a`.

Rules (STM Or Left) and (STM Or Right) define M or N as making a nondeterministic choice within a transaction. Such choices may be derived at the level of the IO monad, but this operator introduces nondeterminism into transactions (which otherwise are deterministic). Nondeterminism is used in our programming example only to ensure completeness with respect to its specification; without nondeterminism we would still have soundness.

4.2 Type System

We complete our formalization of STM Haskell by defining a simple type system that prevents ill-formed expressions, such as the inappropriate mixing of pure, STM and I/O expressions.

The type system only permits the reading and writing of transactional variables inside transactions, which a fortiori enforces *static separation* [1] and permits us to reason about transactions as if they occur in a single step.

Let the *domain*, $\text{dom}(M)$, of an expression M be the set of (free) addresses of the transactional variables at top level in the expression. We have $\text{dom}(a \mapsto M) = \{a\}$, $\text{dom}(M \gg_{\text{IO}} N) = \text{dom}(M)$, $\text{dom}(M \gg_{\text{STM}} N) = \text{dom}(M)$, $\text{dom}(M \mid N) = \text{dom}(M) \cup \text{dom}(N)$ and $\text{dom}((\text{va})M) = \text{dom}(M) \setminus \{a\}$. Otherwise, $\text{dom}(M) = \emptyset$. In particular, expressions that are not in a top-level evaluation context should have no free transactional variables, so the type system enforces that their domain is empty.

Here is the syntax of types. For the sake of simplicity, we formalize only a monomorphic type system. We make the standard assumption that uses of Hindley-Milner style polymorphism may be represented by monomorphising via code duplication.

Types:

$u ::= t \mid T$	type
$t ::= t \rightarrow t \mid X \mid \text{TVar } t \mid \text{IO } \emptyset t \mid \text{STM } \emptyset t$	expression type
$T ::= \text{IO } \bar{a} t \mid \text{STM } \bar{a} t \mid \text{heap } \bar{a} \mid \text{proc } \bar{a}$	Configuration type

An *expression type* t describes the eventual value of a pure functional computation. They are either function types ($t \rightarrow t$), algebraic datatypes (X), TVar reference types ($\text{TVar } t$), IO computation types ($\text{IO } \emptyset t$) or STM transaction types ($\text{STM } \emptyset t$). We usually write $\text{IO } t$ for $\text{IO } \emptyset t$, and $\text{STM } t$ for $\text{STM } \emptyset t$.

A *configuration type* T describes the structure, heap and potential return value (if any) of imperative and concurrent expressions. Heap-expressions with domain \bar{a} have type $\text{heap } \bar{a}$. Both running transactions and STM-expressions with domain \bar{a} have type $\text{STM } \bar{a} t$ for some t . Both threads and IO-expressions with domain \bar{a} have type $\text{IO } \bar{a} t$ for some t .

Finally, the type $\text{proc } \bar{a}$ consists of concurrent expressions with domain \bar{a} that are executed in the background for their effects, but whose results will be discarded. Given T , we write $\text{dom}(T)$ for its domain.

We assume that all polymorphic algebraic datatypes X and their constructors f have been monomorphized by instantiating each of their occurrences. For instance, the type $\text{Maybe } a$ is instantiated at the unit type $()$ as $\text{data Maybe } () = \text{Nothing } () \mid \text{Just } ()$. We assume a set of predefined algebraic types $()$, Error , List_t , Bool , and Maybe_t , with constructors $()$, Nil_t , Cons_t , True , False , Nothing_t , and Just_t .

The return type of an expression is the type of its rightmost thread. The typing rule for parallel composition guarantees that an expression consists of some transactional variables together with either several IO threads or a single rightmost STM thread (currently running a transaction). Moreover, it ensures that there is at most one transactional variable at each location a . It uses the partial non-commutative operation $T \otimes T'$, defined as follows, where $\bar{a} \uplus \bar{b}$ is $\bar{a} \cup \bar{b}$ if \bar{a} and \bar{b} are disjoint.

$$\begin{aligned}
\text{heap } \bar{a} \otimes \text{heap } \bar{b} &:: \text{heap } \bar{a} \uplus \bar{b} \\
\text{proc } \bar{a} \otimes \text{heap } \bar{b} &:: \text{proc } \bar{a} \uplus \bar{b} \\
\text{heap } \bar{a} \otimes \text{STM } \bar{b} t &:: \text{STM } \bar{a} \uplus \bar{b} t \\
\text{IO } \bar{a} t \otimes \text{heap } \bar{b} &:: \text{proc } \bar{a} \uplus \bar{b} \\
T \otimes \text{proc } \bar{a} &:: \text{proc } \text{dom}(T) \uplus \bar{a} \text{ if } T \neq \text{STM } \bar{b} t' \\
T \otimes \text{IO } \bar{a} t &:: \text{IO } \text{dom}(T) \uplus \bar{a} t \text{ if } T \neq \text{STM } \bar{b} t'
\end{aligned}$$

In particular, note that $\text{STM } \bar{a} t \otimes \text{STM } \bar{b} t'$ is undefined, and hence the type system does not allow two transactions to run at once.

Lemma 1. $(T_1 \otimes T_2) \otimes T_3 = T_1 \otimes (T_2 \otimes T_3) = T_2 \otimes (T_1 \otimes T_3) = (T_2 \otimes T_1) \otimes T_3$.

A typing environment $E \in \mathbf{E}$ is a finite mapping from $\mathbf{X} \cup \mathbf{N}$ to types. Each individual map is written as $a :: \text{TVar } t$ or $x :: t$. We write $\bar{x} :: \bar{t}$ for the environment $x_1 :: t_1, \dots, x_n :: t_n$ where n is the length of \bar{x} and \bar{t} . We write E, E' for the union of E and E' when E and E' have disjoint domains. The full typing rules are given in Figure 1 on page 7.

The rule (T BUILTIN) appeals to a relation $g :: u_1 \rightarrow \dots \rightarrow u_n \rightarrow u'$, defined as follows, which gives a type for each application of a builtin function g . In the following, all types t, t' and domains \bar{a} are universally quantified, and $\bar{u} \rightarrow u'$ stands for u' when $|\bar{u}| = 0$, and otherwise for $u_1 \rightarrow \dots \rightarrow u_n \rightarrow u'$.

Types for Builtin Functions: $g :: \bar{u} \rightarrow u'$

Y	$:: (t \rightarrow t) \rightarrow t$
equal	$:: \text{TVar } t' \rightarrow \text{TVar } t' \rightarrow \text{Bool}$
readTVar	$:: \text{TVar } t \rightarrow \text{STM } t$
writeTVar	$:: \text{TVar } t \rightarrow t \rightarrow \text{STM } ()$
$\text{return}_{\text{STM}}$	$:: t \rightarrow \text{STM } t$
retry	$:: \text{STM } t$
\gg_{STM}	$:: \text{STM } \bar{a} t' \rightarrow (t' \rightarrow \text{STM } t) \rightarrow \text{STM } \bar{a} t$
orElse	$:: \text{STM } t \rightarrow \text{STM } t \rightarrow \text{STM } t$
or	$:: \text{STM } t \rightarrow \text{STM } t \rightarrow \text{STM } t$
atomically	$:: \text{STM } t \rightarrow \text{IO } t$
$\text{return}_{\text{IO}}$	$:: t \rightarrow \text{IO } t$
\gg_{IO}	$:: \text{IO } \bar{a} t' \rightarrow (t' \rightarrow \text{IO } t) \rightarrow \text{IO } \bar{a} t$

For example, the function swap has type $\text{TVar } t \rightarrow \text{TVar } t \rightarrow \text{STM } ()$ for each t . Hence, the expression $a \mapsto M \mid b \mapsto N \mid \text{swap } a b$ is well-typed, by (T PAR), (T CELL), and (T APP). But the expression $a \mapsto M \mid b \mapsto N \mid \text{swap } a b \mid \text{swap } a b$ is not well-typed, since it has two STM threads and $\text{STM } t \otimes \text{STM } t'$ is undefined. As a second example, the expression $\lambda x. (x \mid x) (a \mapsto ())$ is not well-typed since the transactional variable $a \mapsto ()$ has type $\text{heap } \bar{a}$; $\text{heap } \bar{a}$ is not an expression type, so we cannot derive any valid function type $t \rightarrow t'$ for the lambda-expression. Indeed, this expression would yield $a \mapsto () \mid a \mapsto ()$, which has two transactional variables with the same location. Such ill-formed expressions are untypable, due to the disjointness conditions of \otimes (see (T PAR)). Similarly, the expression $\lambda x. (x \mid x) (a \mapsto ()) \mid \text{return}_{\text{IO}} ()$ is not well-typed since x must have an expression type, which always has empty domain. However, $\lambda x. (x \mid x)$ has type $\text{IO } t \rightarrow \text{IO } t$ for each t , by (T PAR) and (T LAMBDA). Thus, the expression $\lambda x. (x \mid x) (\text{va})(a \mapsto ()) \mid \text{return}_{\text{IO}} ()$ is well-typed.

For example, for a well-typed application of swap , we have the expected result,

$$a \mapsto M \mid b \mapsto N \mid \text{swap } a b \rightarrow^* a \mapsto N \mid b \mapsto M \mid \text{return}_{\text{STM}} ()$$

but an ill-typed application may have an undesirable outcome.

$$\begin{aligned}
a \mapsto M \mid b \mapsto N \mid \text{swap } a b \mid \text{swap } a b &\rightarrow^* \\
a \mapsto N \mid b \mapsto N \mid \text{return}_{\text{STM}} () \mid \text{return}_{\text{STM}} () &
\end{aligned}$$

<p>(T VAR)</p> $\frac{}{E, x :: t \vdash x :: t}$	<p>(T ADDR)</p> $\frac{}{E, a :: \text{TVar } t \vdash a :: \text{TVar } t}$	<p>(T EMP)</p> $\frac{}{E \vdash \text{emp} :: \text{heap } \emptyset}$
<p>(T LAMBDA)</p> $\frac{}{E, x :: t \vdash M :: t'}$ $\frac{}{E \vdash \lambda x. M :: (t \rightarrow t')}$	<p>(T APP)</p> $\frac{}{E \vdash M :: t \rightarrow t' \quad E \vdash N :: t}$ $\frac{}{E \vdash M N :: t'}$	<p>(T BUILTIN) $(g :: \bar{u} \rightarrow u')$</p> $\frac{}{E \vdash M_1 :: u_1 \quad \dots \quad E \vdash M_n :: u_n}$ $\frac{}{E \vdash g M_1 \dots M_n :: u'}$
<p>(T ADT) $(\text{data } X = f_1 \bar{t}_1 \mid \dots \mid f_m \bar{t}_m, \bar{t}_i = \bar{M}_i)$</p> $\frac{}{E \vdash M_1 :: t_1^1 \quad \dots \quad E \vdash M_m :: t_1^m}$ $\frac{}{E \vdash f_i \bar{M} :: X}$	<p>(T CASE) $(\text{data } X = f_1 \bar{t}_1 \mid \dots \mid f_m \bar{t}_m)$</p> $\frac{}{E \vdash M :: X \quad E, \bar{x}_1 :: \bar{t}_1 \vdash N_1 :: t' \quad \dots \quad E, \bar{x}_m :: \bar{t}_m \vdash N_m :: t'}$ $\frac{}{E \vdash \text{case } M \text{ of } f \bar{x} \rightarrow \bar{N} :: t'}$	
<p>(T CELL)</p> $\frac{}{E, a :: \text{TVar } t \vdash N :: t}$ $\frac{}{E, a :: \text{TVar } t \vdash a \mapsto N :: \text{heap } a}$	<p>(T PAR)</p> $\frac{}{E \vdash M :: T_M \quad E \vdash N :: T_N}$ $\frac{}{E \vdash M \mid N :: T_M \otimes T_N}$	<p>(T RES)</p> $\frac{}{E, b :: \text{TVar } t \vdash M :: \text{heap } b \otimes T}$ $\frac{}{E \vdash (vb)M :: T}$

Figure 1. Type system

Lemma 2 (Subject Reduction).

If $E \vdash M :: u$ and $M \rightarrow M'$ then $E \vdash M' :: u$.

From this point, we only consider well-typed processes (that is, such that there is a typing environment under which they have a type). This is motivated by Lemma 2. Moreover, due to the structural definition of the type system, every subexpression of a well-typed process is well-typed. In order to reason compositionally about multi-step reductions, we develop some simple conditions for when two reductions are independent. We use these conditions in our correctness proofs, where we often consider only transactions and reason up to β -equivalence. We begin by dividing reductions into pure \rightarrow_β and impure \rightarrow_{STM} . (This distinction is different from the one in [10], where the transition relation is stratified and there is only one kind of top-level transition.)

Definition 3. We write $M \rightarrow_\beta N$ if $M \rightarrow N$ can be derived using only the rules in group (R1). We write \rightarrow_{STM} for $(\rightarrow \setminus \rightarrow_\beta)$ and \rightarrow^* for $\rightarrow_\beta^* \rightarrow_{STM}$ (the composition of \rightarrow_β^* and \rightarrow_{STM}). We let $=_\beta$ be the smallest equivalence relation containing \rightarrow_β and \equiv .

Using Lemma 2, we can show that the pure reductions of a single thread are deterministic, and that they commute with reductions in other threads. β -reduction thus enjoys the diamond property.

Lemma 4. If $M \rightarrow M_1$ and $M \rightarrow_\beta M_2$ with $M_1 \not\equiv M_2$ then $M_1 \rightarrow_\beta M'$ and $M_2 \rightarrow M'$ for some M' .

4.3 Comparison with the Original Semantics

The original STM Haskell semantics [10] is based on three different transition relations: I/O transitions, administrative transitions, and STM transitions. These are defined on structures built from expressions, heaps, and multiple threads. In contrast, our semantics of STM Haskell is in the style of a process calculus (like the semantics of Concurrent Haskell [24], for example) and consists of a single reduction relation defined on expressions, whose syntax subsumes heaps and concurrent threads.

The difference in styles, though, is essentially syntactic. We can show that our reduction relation is equivalent to the original semantics. In the extended version of this paper we show a straightforward translation between our syntax and the original run-time syntax, which yields a strong operational correspondence.

Having finished the development of our theory, we suspect it would be quite possible to recast it directly on top of the original semantics.

Still, we contend that our use of a uniform syntax of expressions is better suited to the development of theories for reasoning about STM Haskell programs. One reason is because it allows us to define contextual equivalence (in Section 6) in the standard way, and to import ideas from process calculus, such as bisimulation, directly. Another reason is that our STM reduction rules (in groups (R2) and (R3)) operate on the adjacent piece H of the heap, as opposed to the full heap; this facilitates reasoning about the part of the heap that is actually used by a transaction. Moreover, we can easily represent parts of the run-time state, such as a thread together with a small piece of the heap. The syntax also allows multiple threads with local state to be composed using the parallel operator.

On the other hand, although our expression syntax is uniform, we need to introduce configuration types, as well as conventional types, to rule out certain ill-formed expressions. This is certainly a cost we must pay for the uniform syntax, but we have not found it so onerous; we need a type system anyway, and the additional rules are not hard to work with.

5. Verifying the Ambient API

We are now in a position to specify the expected behaviour of the Haskell code for the ambient API in Section 3, and to verify it. We do so by showing that the API is a fully abstract implementation of the *ambient calculus*, a small calculus of tree-manipulating processes. Theorem 1, below, shows soundness and completeness of the API, while Theorem 2 shows that ambient processes and their Haskell implementations are in fact bisimilar.

Although the high-level statement of correctness is fairly intuitive, the definitions of correspondence between the run time states of our Haskell code and the ambient calculus are rather detailed and technical. The proofs themselves, in the long version of this paper, are also rather complicated. Still, the theorems and their proofs show the viability of our theory for reasoning about STM Haskell code. To the best of our knowledge, ours is the first theory for equational reasoning about concurrent Haskell programs (as opposed to say the correctness of implementations).

5.1 An (Imperative) Ambient Calculus

Our Haskell API is intended to implement the primitives of an ambient calculus, defined as follows. calculus [5]. Readers familiar with the ambient calculus will notice that every syntactic form of the original calculus also exists as an imperative operation in iAmb.

Syntax of the Ambient Calculus:

$\pi ::=$	simple capability
into a	enter a
out a	leave a
open a	open a
amb $a C$	create ambient $a[C]$
fork C	fork thread C
new(a) C	a fresh in C
$C ::= \pi \mid \text{nil} \mid C.C$	capabilities
$P ::=$	Process
$\mathbf{0}$	inactivity
$ a[P]$	ambient
$ C.P$	prefixed thread
$ (va)P$	restriction
$ P \mid P$	parallel
$\mathcal{R} ::= [\cdot] \mid a[\mathcal{R}] \mid (va)\mathcal{R} \mid \mathcal{R} \mid P \mid P \mid \mathcal{R}$	Reduction context

We often omit the $\mathbf{0}$ in $C.\mathbf{0}$ and $a[\mathbf{0}]$. Free and bound names of capabilities and processes are defined as expected. The scope of the bound name a extends to P in $(va).P$ and to C in $\text{new}(a).C$.

The reduction semantics of the ambient calculus are defined as follows. Structural equivalence \equiv is the least congruence on processes, with respect to the reduction (\mathcal{R}) contexts, that satisfies commutative monoid laws for \mid with $\mathbf{0}$ as unit and the rules below.

Structural Equivalence for Ambient Processes: $P \equiv Q$

$\text{nil}.P \equiv P$	(A EPS)
$(C_1.C_2).P \equiv C_1.(C_2.P)$	(A ASSOC)
$\mathcal{R}[(va)P] \equiv (va)\mathcal{R}[P]$ if $n \notin \text{n}(\mathcal{R})$	(A RES)

Reduction \rightarrow of processes is the least relation satisfying the following rules.

Reduction for Ambient Processes: $P \rightarrow Q$

$b[\text{into } a.P \mid Q] \mid a[R] \rightarrow a[b[P \mid Q] \mid R]$	(A IN)
$a[b[\text{out } a.P \mid Q] \mid R] \rightarrow b[P \mid Q] \mid a[R]$	(A OUT)
$\text{open } a.P \mid a[Q] \rightarrow P \mid Q$	(A OPEN)
$(\text{new}(a) C).P \rightarrow (va)C.P$ if $a \notin \text{fn}(P)$	(A NEW)
$\text{amb } a C.P \rightarrow a[C.\mathbf{0}] \mid P$	(A AMB)
$\text{fork } C.P \rightarrow C.\mathbf{0} \mid P$	(A FORK)
$P \rightarrow P' \implies \mathcal{R}[P] \rightarrow \mathcal{R}[P']$	(A R CTX)
$P \equiv \implies P' \implies P \rightarrow P'$	(A STRUCT)

The first three rules specify how the tree structure can be modified. If $\text{into } a$ is executed inside a location b that has a sibling a , then b is moved inside a . Conversely, if $\text{out } a$ is executed inside a location b that is a child of a , then b is moved outside a . Finally, $\text{open } a$ opens a single child named a of the ambient it is running in.

As a simple example, we take the ambient tree $a[p[\text{out } a.\text{into } b]] \mid b[]$, where the ambient p represents a packet that intends to move from a to b : $a[p[\text{out } a.\text{into } b]] \mid b[] \rightarrow a[] \mid p[\text{into } b] \mid b[] \rightarrow a[] \mid b[p[]]$. We define the delay operator τ as $\tau.P := a_\tau[] \mid \text{open } a_\tau.P$ for some distinguished a_τ .

In this setting, processes such as $C.a[P]$ are ill-formed, since they have no direct correspondent in the API. We instead use $C.\text{amb } a P$. Formally, we treat only the following subcalculus; processes that result from the execution of a closed process $C.\mathbf{0}$.

Normal form for a subcalculus of iAmb

$P_N ::= a[P_N] \mid (va)P_N \mid (P_N \mid P_N) \mid C.\mathbf{0} \mid \mathbf{0}$

We write \mathbf{P}_N for the set of all P_N . As an example, $(\text{out } a.\text{into } b).\mathbf{0} \in \mathbf{P}_N$, but $\text{out } a.(\text{into } b.\mathbf{0}) \notin \mathbf{P}_N$. Note that \mathbf{P}_N is not closed under structural equivalence, although it is closed (modulo structural equivalence) under reduction. We write \rightarrow_N for \rightarrow restricted to

$\mathbf{P}_N \times \mathbf{P}_N$. In the remainder of the paper, we only consider processes $P \in \mathbf{P}_N$. Continuing the running example:

$$\begin{aligned} & \text{amb } a (\text{amb } p (\text{out } a.\text{into } b)).\text{amb } b \text{nil}.\mathbf{0} \\ & \rightarrow_N a[\text{amb } p (\text{out } a.\text{into } b).\mathbf{0}] \mid \text{amb } b \text{nil}.\mathbf{0} \\ & \rightarrow_N a[p[\text{out } a.\text{into } b.\mathbf{0}]] \mid \text{amb } b \text{nil}.\mathbf{0} \\ & \rightarrow_N a[p[\text{out } a.\text{into } b.\mathbf{0}]] \mid b[] \end{aligned}$$

5.2 Statement of Correctness

Cardelli [4] defined a notion of correctness for implementations of the ambient calculus, which we quote here:

The problem. We want to find a (nondeterministic) implementation of the reduction relation \rightarrow^* , such that each P_i in an ambient is executed by a concurrent thread (and so on recursively in the subambients $m_j[\dots]$).

Desirable properties of the implementation are:

- *Liveness*: If $P \rightarrow Q$ then the implementation must reduce P .
- *Soundness*: If the implementation reduces P to Q , then we must have $P \rightarrow^* Q$.
- *Completeness*: If $P \rightarrow^* Q$, then the implementation must be able (however unlikely) to reduce P to some $Q' \equiv Q$.

Additional Properties. In addition to the three properties proposed by Cardelli, we formalize the following two, and establish all five as Theorem 1.

- *Safety*: If the implementation reduces P to M then M can reduce further to some Q .
- *Termination*: If the implementation of P has an infinite reduction, then P also does.

Compared to [4], we additionally treat the open capability (and in an extended version of this paper, communication of both names and capabilities).

The proof of Theorem 1 proceeds as follows: We begin by giving a simple correspondence between ambient capabilities and their Haskell implementation. In Definition 5, we define how an ambient process is implemented as a Haskell expression, including heap and running capabilities. Definition 6 bridges the gap between this intensional specification and the expressions that arise when executing the expressions; the main difference is due to the lack of garbage collection in our semantics. Then, Lemma 7 guarantees that the correspondence does not confuse unrelated ambient processes.

With the static correspondence in place, we can then show how it is preserved by execution. Lemma 8 details how the execution of the implementation of a prefix corresponds to its semantics in the ambient calculus. Finally, in the proof of Theorem 1 we close the result of Lemma 8 under contexts, yielding a strong operational correspondence.

5.3 Correspondence between Haskell Code and Ambients

The encoding $\llbracket C \rrbracket$ into Haskell of imperative ambient capabilities is homomorphic, except for two cases:

$$\begin{aligned} \llbracket \text{new}(a) C \rrbracket & := (\text{new } []) \gg \lambda a \rightarrow \llbracket C \rrbracket \\ \llbracket C'.C \rrbracket & := \llbracket C' \rrbracket \gg \llbracket C \rrbracket \end{aligned}$$

Continuing the running example, we have:

$$\begin{aligned} & \llbracket \text{amb } a (\text{amb } p (\text{out } a.\text{into } b)).\text{amb } b \text{nil} \rrbracket \\ & = \text{amb } a (\text{amb } p (\text{out } a \gg \text{into } b)) \gg \text{amb } b \text{nil} \end{aligned}$$

We can then give a compositional definition of what it means for the run-time state of a Haskell program to correspond to (the structure

of) a given iAmb process. This definition encapsulates both the heap shape invariant preserved by the functions of the API, and how a given ambient calculus process is represented in the heap. The definition has two levels. At the inner level (Definition 5), we inductively match the structure of an ambient process against a structured decomposition of a process term. At the outer level (Definition 6), we perform sanity checks, open restrictions, discard unused heap items and identify the root ambient.

Definition 5. We identify association lists with the corresponding binary relations, that must be injective. We identify other lists with multisets. We then say that $(D_n, D_p, D_h, D_c) \in (D_n, D_p, D_h, D'_c) \oplus (D_n, D_p, D_h, D'_c)$ if $D_c \in D'_c \cup D'_c$. We write D for an $\text{AD}(D_n, D_p, D_h, D_c)$. An agent C at location h is $\llbracket C.\mathbf{0} \rrbracket_h := \text{case } \llbracket C \rrbracket \text{ of Agent } x \rightarrow x h$.

Informally, we write $(a \mapsto D, H_h, H, M) \in \mathbf{M}(P)$ if $a \mapsto D$ is the current ambient, H_h its handles, H the data and handles of all its subambients and M the running capabilities in P . $\mathbf{M}(P)$ is inductively defined as follows:

(Completed agent)

$$(a \mapsto (D_n, D_p, D_h, \square), \Pi_{h \in D_h} h \mapsto a, \text{emp}, \text{return}_{\text{IO}} ()) \in \mathbf{M}(P) \text{ if } P \equiv \mathbf{0}.$$

(Agent running in the current ambient)

$$(a \mapsto (D_n, D_p, D_h, \square), \Pi_{h \in D_h} h \mapsto a, \text{emp}, \llbracket C \rrbracket_h) \in \mathbf{M}(P) \text{ if } P \equiv C.\mathbf{0} \text{ and } h \in D_h$$

(Child of the current ambient)

$$(a \mapsto (D_n, D_p, D_h, [(b, c)]), H_h, H, M) \in \mathbf{M}(P) \text{ if } P \equiv b[Q] \text{ and } H \equiv c \mapsto D' \mid \Pi_{h \in D'_h} h \mapsto c \mid H' \text{ where } (c \mapsto D', \Pi_{h \in D'_h} h \mapsto c, H', M) \in \mathbf{M}(Q), D'_n = b \text{ and } D'_p = \text{Some } h' \text{ with } h' \in D_h$$

(Parallel decomposition)

$$(a \mapsto D, H_h, H, M) \in \mathbf{M}(P) \text{ if } P \equiv Q_1 \mid Q_2, H \equiv H_1 \mid H_2, M \equiv M_1 \mid M_2, D \in D_1 \uplus D_2 \text{ with } (a \mapsto D_1, H_h, H_1, M_1) \in \mathbf{M}(Q_1) \text{ and } (a \mapsto D_2, H_h, H_2, M_2) \in \mathbf{M}(Q_2).$$

We can then define what it means for M to be a run-time state corresponding to an ambient process P_0 .

Definition 6. $M \in \mathcal{M}(P_0)$ iff

1. There are P, \bar{e} such that $P_0 \equiv (\nu \bar{e})P$ and P is not a $\mathcal{R}[(\nu a)Q]$ (the top-level restrictions of P_0 are \bar{e});
2. $\text{fn}(P_0) \subseteq \text{dom}(M)$ and $E \vdash M :: \text{IO } \bar{a} ()$ for $E := \{a_i :: \text{TVar } [\text{Char}] \mid a_i \in \text{dom}(M)\}$ (M has the free names of P_0 in its domain, and is well-typed);
3. $M \equiv (\nu a \bar{b} \bar{c} \bar{e})(a \mapsto \square \mid b \mapsto (a, \text{None}, D_h, D_c) \mid H_0 \mid H_1 \mid H_2 \mid H_3 \mid M')$ (we can split M into the root ambient, some heaps and some running code);
4. $H_0 = \Pi_i d_i \mapsto N_i$ with $\bar{d} \cap \text{fn}(D_h \mid D_c \mid H_1 \mid H_2 \mid H_3 \mid M') = \emptyset$. Moreover, if $N_i = D'$ then $D'_p \neq \text{None}$ (H_0 is unreachable garbage not containing a root ambient);
5. $H_1 = \Pi_{n \in \text{fn}(P)} n \mapsto s_n$ with $\emptyset \vdash s_n :: \text{String}$ (H_1 is the free names of P , and is well-typed);
6. $H_2 = \Pi_{h \in D_h} h \mapsto b$ (H_2 is the handles of the root ambient);
7. There are no \mathcal{R}_1, a, M'' such that $H_3 \mid M' \equiv \mathcal{R}_1[(\nu a)M'']$ (there are no further restricted heap cells at the top level); and
8. $(a \mapsto D, H_2, H_3, M') \in \mathbf{M}(P)$.

Both \mathbf{M} and \mathcal{M} characterize \mathbf{P}_N modulo structural equivalence.

Lemma 7. If $P \equiv Q$ then $\mathcal{M}(P) = \mathcal{M}(Q)$ and $\mathbf{M}(P) = \mathbf{M}(Q)$. Conversely, if $\mathbf{M}(P) \cap \mathbf{M}(Q) \neq \emptyset$ or $\mathcal{M}(P) \cap \mathcal{M}(Q) \neq \emptyset$ then $P \equiv Q$.

5.4 Operational Semantics of the Implementation

The transactions of the implementations of prefixes exactly correspond to the axioms of the ambient calculus operational semantics,

lifted to Haskell using the \mathbf{M} function. We show the case of the into prefix.

Lemma 8. If $C.\mathbf{0} \equiv \text{into } a.P$ and $(d \mapsto D, H_2, H_3, M) \in \mathbf{M}(a[Q] \mid b[C.\mathbf{0} \mid R_1] \mid R_2)$, $M = \mathcal{R}_1[\llbracket C.\mathbf{0} \rrbracket_{h_3}]$, $\{(a, d_2), (b, d_3)\} \in D_c$ with $d_2 \neq d_3$, $H_3 \equiv d_2 \mapsto D_2 \mid h_3 \mapsto d_3 \mid d_3 \mapsto D_3 \mid H'_3$ with $D_3_p = \text{just } h$ and $H_2 \equiv h \mapsto d \mid H'_2$, then $d \mapsto D \mid H_2 \mid H_3 \mid M \Rightarrow_{\beta} d \mapsto D' \mid H_2 \mid d_2 \mapsto D_2' \mid h_3 \mapsto d_3 \mid d_3 \mapsto D_3' \mid H'_3 \mid \mathcal{R}_1[\llbracket C'.\mathbf{0} \rrbracket_{h_3}]$ where $C'.\mathbf{0} \equiv P$ and $(d \mapsto D', H_2, d_2 \mapsto D_2' \mid h_3 \mapsto d_3 \mid d_3 \mapsto D_3' \mid H'_3, \mathcal{R}_1[\llbracket C'.\mathbf{0} \rrbracket_{h_3}]) \in \mathbf{M}(a[Q] \mid C'.\mathbf{0} \mid R_1] \mid R_2)$.

5.5 Main Results About the Haskell Code

Our first correctness result establishes direct correspondences between ambient processes and the states of the Haskell implementation; the different properties in this theorem generalize the properties sought by Cardelli [4]. Recall the definition of $\rightarrow := \rightarrow_{\beta}^* \rightarrow_{STM}$, intuitively “performing a transaction”.

Theorem 1.

- Liveness, Completeness:
If $P \rightarrow_N Q$ and $M \in \mathcal{M}(P)$ then $M \Rightarrow_{\beta} M' \in \mathcal{M}(Q)$.
- Safety, Soundness:
If $M \in \mathcal{M}(P)$ and $M \rightarrow M'$ then $P \rightarrow_N Q$ with $M' =_{\beta} M' \in \mathcal{M}(Q)$.
- Termination:
If $M \in \mathcal{M}(P)$ and M has an infinite reduction then P has an infinite reduction.

Proof sketch.

1. Assume that $M \rightarrow M'$ and that $M \in \mathcal{M}(P)$ where $P \equiv (\nu \bar{e})P_0$ such that P_0 does not have any top-level restrictions. By assumption, $M \equiv (\nu a \bar{b} \bar{c} \bar{e})(a \mapsto \square \mid b \mapsto (a, \text{None}, D_h, D_c) \mid H_0 \mid H_1 \mid H_2 \mid H_3 \mid N)$ such that $H_1 \mid H_2 \mid H_3 \mid N \rightarrow H'_1 \mid H'_2 \mid H'_3 \mid N'$ and $A := (b \mapsto (a, \text{None}, D_h, D_c), H_2, H_3, N) \in \mathbf{M}(P_0)$. By induction on the derivation of $A \in \mathbf{M}(P_0)$, $N = \Pi_i N_i$ is a parallel composition of several $N_i = \llbracket C_i \rrbracket_{h_i}$. Then there is j such that $H_1 \mid H_2 \mid H_3 \mid \llbracket C_j \rrbracket_{h_j} \rightarrow H'_1 \mid H'_2 \mid H'_3 \mid N'_j$ with $N' =_{\beta} N'_j \mid \Pi_{i \neq j} N_i$. As shown in Lemma 8 for the in prefix, and in the extended version for the other prefixes, we then have $H_1 \mid H_2 \mid H_3 \equiv H_R \mid d \mapsto D \mid H_h \mid H_S$ such that $P_0 \equiv \mathcal{R}[\mathcal{R}_2[C_j.Q]]$, $(d \mapsto D, H_h, H_S) \in \mathbf{M}(\mathcal{R}_2[C_j])$ and $H'_1 \mid H'_2 \mid H'_3 \equiv H_R \mid d \mapsto D' \mid H'_h \mid H'_S$ such that $(d \mapsto D', H'_h, H'_S) \in \mathbf{M}(\mathcal{R}_2'[Q])$ where $C_j.\mathbf{0} \equiv C'_j.Q'$ and $\mathcal{R}_2[C_j.Q'] \rightarrow \mathcal{R}_2'[Q']$ is an axiom. By induction on the derivation of $A \in \mathbf{M}(P_0)$, $M' =_{\beta} (\nu a \bar{b} \bar{c} \bar{e})(a \mapsto \square \mid b \mapsto (a, \text{None}, D_h, D_c) \mid H_0 \mid H'_1 \mid H'_2 \mid H'_3 \mid N'_j \mid \Pi_{i \neq j} N_i)$. $M'_{\beta} \in \mathcal{M}(\mathcal{R}[\mathcal{R}_2'[Q]])$ follows by Lemma 7.

2. Assume that $P \rightarrow P'$. Let \bar{e} be the top-level restrictions of P . If the reduction occurred inside an ambient, then there are a, Q, R and contexts $\mathcal{R}_1, \mathcal{R}_2$ where $P \equiv (\nu \bar{e})\mathcal{R}_1[a[\mathcal{R}_2[\pi.Q] \mid R]]$, $\mathcal{R}_2[\pi.Q] \rightarrow \mathcal{R}'_2[Q]$ is an instance of an axiom and $P' \equiv (\nu \bar{e})\mathcal{R}[a[\mathcal{R}'_2[Q] \mid R]]$.

By assumption $M \in \mathcal{M}(P)$, so $N \equiv \mathcal{R}_1[d \mapsto D \mid H_h \mid H \mid N]$ such that $(d \mapsto D, H_h, H, N) \in \mathbf{M}(a[\mathcal{R}_2[\pi.Q] \mid R])$. Thus, $H \equiv c \mapsto D' \mid H_1 \mid H_2 \mid \Pi_{h \in D'_h} h \mapsto c$ and $N \equiv N_1 \mid N_2$ with $D'_n = b$, $D'_p = \text{Some } h'$, $h' \in D_h$ and $D \in D'_1 \uplus D'_2$ with $A := (c \mapsto D', \Pi_{h \in D'_h} h_i \mapsto c, H_1, N_1) \in \mathbf{M}(\mathcal{R}_2[\pi.Q])$ and $(c \mapsto D'_2, \Pi_{h \in D'_h} h \mapsto c, H_2, N_2) \in \mathbf{M}(R)$.

By induction on the derivation of $A \in \mathbf{M}(\mathcal{R}_2[\pi.Q])$, we have $N_1 \equiv \llbracket C' \rrbracket_{h_1} \mid N'_1$ with $C'.\mathbf{0} \equiv \pi.Q$. We treat the case where π is not $\text{new}(a)C$. As shown in Lemma 8 for the into prefix, and in the extended version for the other prefixes, $c \mapsto D'_1 \mid$

$\prod_{h_i \in D'_i} h_i \mapsto c \mid H_1 \mid \llbracket C \rrbracket_{h_i} \rightarrow c \mapsto D'_1 \mid H'_h \mid H'_l \mid \llbracket C_Q \rrbracket_{h_i}$ with $C_Q.0 \equiv Q$ and $(c \mapsto D'_1, H'_h, H'_l, \llbracket C_Q \rrbracket_{h_i}) \in \mathbf{M}(\mathcal{R}'_2[C_Q.0])$.

If the reduction occurs at top level, we have $P \equiv (\nu \bar{e})(Q \mid R)$, and $N \equiv \mathcal{R}[d \mapsto D \mid H_h \mid H \mid N]$ such that $(d \mapsto D, H_h, H, N) \in \mathbf{M}(Q \mid R)$. The rest of the proof proceeds analogously.

- This follows from the completeness above and the fact that $\mathcal{M}(P)$ is \rightarrow_β -convergent (modulo \equiv). \square

The proof of this theorem uses Lemma 8 to prove that an agent can progress whenever the corresponding ambient process does and to get the shape of the result of the transition. The proof also uses the compositionality of the calculus; specifically in order to separate an agent (running as part of an expression in the IO monad) and the heap it needs to progress.

Next, we define a notion of bisimulation between ambient processes and STM Haskell expressions.

Definition 9. $\mathbf{R} \subseteq \mathbf{M} \times \mathbf{P}_N$ is a bisimulation iff for all $(M, P) \in \mathbf{R}$

- If $M \rightarrow M'$ then $P \rightarrow_N P'$ with $(M', P') \in \mathbf{R}$; and
- If $P \rightarrow_N P'$ then $M \rightarrow M'$ with $(M', P') \in \mathbf{R}$.

The expression M is bisimilar to the process P if there is some bisimulation \mathbf{R} with $\mathbf{M} \mathbf{R} P$.

Theorem 2. $H_C \mid \text{root} \llbracket C \rrbracket$ is bisimilar to $\tau.C.0$, where $H_c := \prod_{a_i \in \text{fn}(C)} a_i \mapsto \text{''}$.

Bisimulation between the expressions of our calculus and processes of the ambient calculus allows a succinct statement of the theorem. The proof relies on the soundness of bisimulation up to \rightarrow_β . We could probably replicate this definition using the original semantics of STM Haskell, but it would require many cases; our reformulated semantics allows a simple and direct definition.

6. Equational Reasoning

One of the nice things about functional programming is that we can hope for two expressions to be equivalent, in the sense that they can be substituted for each other in any context. In this section, we develop a proof technique for a Morris-style contextual equivalence. In particular, we prove a number of equations asserted in [10].

6.1 Contextual Equivalence

We begin by defining a notion of a typed relation, stating that two terms are related at a given type under a typing environment.

Definition 10 (Typed Relation). $\mathbf{R} \subseteq \mathbf{E} \times \mathbf{M} \times \mathbf{M} \times \mathbf{T}$ is a typed relation if whenever $(E, M_1, M_2, u) \in \mathbf{R}$ we have $E \vdash M_1 :: u$ and $E \vdash M_2 :: u$. We write $E \vdash M_1 \mathbf{R} M_2 :: u$ for $(E, M_1, M_2, u) \in \mathbf{R}$.

An expression M has terminated, written $M \downarrow$, if its rightmost thread returns. Termination is our only top-level observation.

Termination

(TERM RETURN)	(TERM RES)	(TERM PAR)
$\frac{}{\text{return}_{\text{IO}} M \downarrow}$	$\frac{M \downarrow}{(va)M \downarrow}$	$\frac{M \downarrow}{N \mid M \downarrow}$

An expression M terminates, written $M \downarrow$, if $M \rightarrow^* N$ such that $N \downarrow$.

Definition 11. Contextual equivalence, written \simeq , is the typed relation such that $E \vdash M_1 \simeq M_2 :: u$ if and only if for all contexts \mathcal{C} such that $\circ \vdash \mathcal{C}[M_1] :: \text{IO } \bar{a}()$ and $\circ \vdash \mathcal{C}[M_2] :: \text{IO } \bar{a}()$ we have $\mathcal{C}[M_1] \downarrow$ if and only if $\mathcal{C}[M_2] \downarrow$.

6.2 STM Expressions as Heap Relations

Because of the isolation between different transactions provided by the run-time systems, STM expressions are completely defined by

their effect on the transactional heap. For simplicity (cf. [16, 17, 29]), we work with a pure heap, where the types of elements in the heap do not mention the STM or IO monads.

Definition 12. A type t is pure if it is either $t_1 \rightarrow t_2$ where t_1 and t_2 are pure, if it is $\text{TVar } t'$ where t' is pure, or if it is X such that $\text{data } X = f_1 \bar{t}_1 \mid \dots \mid f_m \bar{t}_m$ where all t_m^i are pure. An environment E is a pure store environment if E is of the form $\cup_i b_i :: \text{TVar } t_i$ where all t_i are pure.

A derivation $E \vdash M :: u$ is pure, written $E \vdash_p M :: u$, if E is a pure store environment and t is pure in all occurrences of $\text{TVar } t$ in the derivation. We then say that M uses only pure heap.

Two STM threads that only use pure heap are equivalent if they modify the heap in the same way and return the same result.

Definition 13. Heap transformer equivalence, written $=_{HT}$, is defined by $E \vdash M =_{HT} N :: u$ if and only if $u = \text{STM } t$, $E \vdash_p M :: u$, $E \vdash_p N :: u$, M and N are β -threads, and for all STM contexts $\mathcal{R}_{\rightarrow}, \mathcal{R}'_{\rightarrow}$, and heaps H such that $E \vdash H :: \text{heap } \bar{a}$ we have $H \mid M \rightarrow^* \mathcal{R}_{\rightarrow}[\text{return}_{\text{STM}} M']$ iff $H \mid N \rightarrow^* \mathcal{R}'_{\rightarrow}[\text{return}_{\text{STM}} M']$; and $H \mid M \rightarrow^* \mathcal{R}_{\rightarrow}[\text{retry}]$ iff $H \mid N \rightarrow^* \mathcal{R}'_{\rightarrow}[\text{retry}]$.

Theorem 3. The relation $=_{HT}$ is sound, that is, $=_{HT} \subseteq \simeq$.

Proof. We let $\overset{\mathcal{C}}{=}_{HT}$ be the smallest typed congruence containing $=_{HT}$. We prove that $\overset{\mathcal{C}}{=}_{HT} \subseteq \simeq$. The proof has three parts:

- If $E \vdash_p M :: t$ and $E \vdash H :: \text{heap } \bar{a}$ then reductions of $H \mid M$ only depend on the pure cells in H .
- Let $\overset{\mathcal{C}}{\simeq}_{HT}$ be the smallest typed congruence such that $E \vdash M \overset{\mathcal{C}}{=}_{HT} N :: t$ with t pure and M, N closed implies $E \vdash M \overset{\mathcal{C}}{\simeq}_{HT} N :: t$.
If $E \vdash_p M :: t$, and G and H are pure heaps related by $\overset{\mathcal{C}}{\simeq}_{HT}$, then derivatives of $G \mid M$ and $H \mid M$ are related by $\overset{\mathcal{C}}{\simeq}_{HT}$.
- We can then derive that $\overset{\mathcal{C}}{=}_{HT}$ is a barbed bisimulation, so it is contained in \simeq . The interesting case is as follows:

Assume that $E \vdash M =_{HT} N :: \text{STM } t$, $E \vdash H \overset{\mathcal{C}}{=}_{HT} G :: \text{heap } \bar{c}$ and $H \mid M \rightarrow^* \mathcal{R}_{\rightarrow}[B]$. To prove that $G \mid N \rightarrow^* \mathcal{R}'_{\rightarrow}[B']$ such that $E \vdash \mathcal{R}_{\rightarrow}[B] \overset{\mathcal{C}}{=}_{HT} \mathcal{R}'_{\rightarrow}[B'] :: \text{STM } \bar{c} t$ we first use 1. and 2. to prove that $G \mid M \rightarrow^* \mathcal{R}''_{\rightarrow}[B'']$ such that $E \vdash \mathcal{R}_{\rightarrow}[B] \overset{\mathcal{C}}{=}_{HT} \mathcal{R}''_{\rightarrow}[B''] :: \text{STM } \bar{c} t$.

Then $G \mid N \rightarrow^* \mathcal{R}'_{\rightarrow}[B']$ such that $E \vdash \mathcal{R}''_{\rightarrow}[B''] \overset{\mathcal{C}}{=}_{HT} \mathcal{R}'_{\rightarrow}[B'] :: \text{STM } \bar{c} t$ by the definition of $=_{HT}$. By transitivity, $E \vdash \mathcal{R}_{\rightarrow}[B] \overset{\mathcal{C}}{=}_{HT} \mathcal{R}'_{\rightarrow}[B'] :: \text{STM } \bar{c} t$. \square

We write $M \leftrightarrow N$ if for all pure store environments E and types t such that $E \vdash_p M :: \text{STM } t$ and $E \vdash_p N :: \text{STM } t$ we have $E \vdash M =_{HT} N :: \text{STM } t$. We can now use Theorem 3 to prove classic equations between expressions.

6.3 Proving the Monad Laws

To be a proper monad, the $\text{return}_{\text{STM}}$ and \gg_{STM} functions must work together according to three laws:

Lemma 14.

- $((\text{return}_{\text{STM}} M) \gg_{\text{STM}} N) \leftrightarrow NM$.
- $(M \gg_{\text{STM}} \lambda x. \text{return}_{\text{STM}} x) \leftrightarrow M$
- $((M \gg_{\text{STM}} f) \gg_{\text{STM}} g) \leftrightarrow (M \gg_{\text{STM}} (\lambda x. fx \gg_{\text{STM}} g))$

Proof.

- The only transition of $H \mid (\text{return}_{\text{STM}} M) \gg_{\text{STM}} N$ is $H \mid (\text{return}_{\text{STM}} M) \gg_{\text{STM}} N \rightarrow \equiv H \mid NM$
- Take $M' \in \{\text{retry}, \text{return}_{\text{STM}} M''\}$. We then have $H \mid M \rightarrow^* \mathcal{R}_{\rightarrow}[M']$ iff

$$M \gg_{\text{STM}} \rightarrow^* \mathcal{R}_{\rightarrow}[M'] \gg_{\text{STM}} \lambda x. \text{return}_{\text{STM}} x. \\ \equiv \mathcal{R}_{\rightarrow}[M'] \gg_{\text{STM}} \lambda x. \text{return}_{\text{STM}} x]$$

We proceed by case analysis on M' .

- $M' = \text{retry}$ iff, using (STM BIND RETRY),
 $\mathcal{R}_{\rightarrow}[M'] \gg_{\text{STM}} \lambda x. \text{return}_{\text{STM}} x \rightarrow \mathcal{R}_{\rightarrow}[\text{retry}]$.
- $M' = \text{return}_{\text{STM}} M''$ iff $\mathcal{R}_{\rightarrow}[M'] \gg_{\text{STM}} \lambda x. \text{return}_{\text{STM}} x \rightarrow \mathcal{R}_{\rightarrow}[\text{return}_{\text{STM}} M'']$, using (STM BIND RETURN) and (BETA).

3. as 2. \square

6.4 Proving Other Equations

We prove classical single-threaded imperative equivalences, such as the commutativity of accesses to independent memory cells.

Lemma 15.

- $(\text{readTVar } a \gg_{\text{STM}} \lambda x. \text{writeTVar } a x) \leftrightarrow \text{return}_{\text{STM}} ()$.
- $(\text{writeTVar } a M \gg_{\text{STM}} \text{writeTVar } b N) \leftrightarrow (\text{writeTVar } b N \gg_{\text{STM}} \text{writeTVar } a M)$ if $a \neq b$.
- $(\text{readTVar } a \gg_{\text{STM}} \lambda x. \text{writeTVar } b M \gg_{\text{STM}} \text{return}_{\text{STM}} x) \leftrightarrow (\text{writeTVar } b M \gg_{\text{STM}} \text{readTVar } a)$ if $a \neq b$

We also prove absorption and associativity laws for `orElse`, as proposed in [10], and associativity and commutativity laws for `or`.

Lemma 16.

1. $\text{orElse } \text{retry } M \leftrightarrow M$
2. $\text{orElse } M \text{ retry} \leftrightarrow M$
3. $\text{orElse } M_1 (\text{orElse } M_2 M_3) \leftrightarrow \text{orElse} (\text{orElse } M_1 M_2) M_3$
4. $\text{or } M N \leftrightarrow \text{or } N M$
5. $\text{or } M_1 (\text{or } M_2 M_3) \leftrightarrow \text{or} (\text{or } M_1 M_2) M_3$

7. Related Work

Prior semantics for languages with STM, such as STM Haskell [10], were developed with an aim to specify and compare transaction models [33] and their implementations [30, 14, 1], or to study the interaction of transactions with other language features [20]. Hu and Hutton [13] show correctness for a compiler for a small transaction language, inspired by STM Haskell. In contrast, our semantics is designed to enable equational reasoning about source programs. In this respect, the development closest to ours is of an equational theory for a process algebra with STM [2]; this work is not about actual code, and includes no substantial example.

Proof techniques for STM programs have focused on checking invariants of shared transactional state, not on equational reasoning. An extension of STM Haskell with run time invariant checking [12] defines a semantics and implementation but does not attempt program verification. A program logic [21] allows invariants to be specified as pre- and post-conditions within a dependent type system and proofs are by typechecking; unlike STM Haskell, this system has no explicit transaction abort or `retry`.

Our main case study is a verification of a centralized shared-memory implementation of ambients. There are several distributed implementations of ambients described in the literature [6, 25, 7]. These have also been verified using techniques from process calculus, but the algorithms are based on message-passing rather than transactional memory. We recently learnt of an independent, but unverified, implementation of ambients within STM Haskell [32]. We intend to investigate whether our verification techniques also apply to this code.

8. Conclusions

It has been prominently argued that functional programming in pure languages like Haskell facilitates equational reasoning [15]

and that transactional memory enables compositional reasoning about concurrent programs [11]. Here we realize this promise in the context of STM Haskell and show how to verify equational properties of a sizeable STM program.

As future work, we want to extend our proof techniques to statically check invariants, and to investigate connections between our model of heaps and concurrency, spatial logics for process calculi, and separation logics for imperative programming languages. A possible further case study to exercise our theory would be to verify an STM implementation of the join calculus.

Acknowledgements Discussions with Cédric Fournet, Tim Harris, Simon Peyton Jones, and Claudio Russo were useful.

A. Source code

This appendix contains the remainder of the source code for the ambient API of Section 3.

Ambient Functions

```
nil = Agent $ \s -> return ()
```

```
new arg = Agent $ \s ->
  atomically $ newTVar arg
```

```
root agent = do
  rHandle <- (atomically $
    do rName <- newTVar "root";
       newAmb Nothing rName);
  case agent of Agent f -> f rHandle
```

```
amb a agent = Agent $ \bHandle -> do {
  aHandle <- atomically $ do {
    aHandle <- newAmb (Just bHandle) a;
    aAmbient <- readTVar bHandle;
    AD (n,p,h,c) <- readAmb bHandle;
    writeAmb bHandle (AD (n,p,h,(a,aAmbient):c));
    return aHandle};
  forkIO $ case agent of Agent f -> f aHandle
;return ()}
```

```
out c = Agent $ \bHandle -> atomically $ do {
  bAmbient <- readTVar bHandle;
  AD (bn,bp,bh,bc) <- readAmb bHandle;
  cHandle <- parentOf bHandle;
  AD (cn,cp,ch,cc) <- readAmb cHandle;
  aHandle <- if (cn == c)
    then parentOf cHandle
    else retry;
  AD (an,ap,ah,ac) <- readAmb aHandle;
  writeAmb aHandle (AD (an,ap,ah,
    (bn,bAmbient):ac));
  writeAmb cHandle (AD (cn,cp,ch,
    delete (bn,bAmbient) cc));
  writeAmb bHandle (AD (bn,Just aHandle,bh,bc))}
```

```
open c = Agent $ \aHandle -> atomically $ do {
  aAmbient <- readTVar aHandle;
  AD (an,ap,ah,ac) <- readAmb aHandle;
  cAmbient <- lookup' c ac;
  AD (cn,cp,ch,cc) <- readTVar cAmbient;
  rePoint aAmbient ch;
  writeAmb aHandle
    (AD (an, ap, ah++ch,
      (delete (cn,cAmbient) ac)++cc))}
```

```

fork agent = Agent $ \s -> do {
  atomically $ return ();
  ; forkIO $ case agent of
    Agent f -> f s
  ; return ()}

```

Helper Functions

```

newAmb :: (Maybe Handle) -> Name -> STM Handle
newAmb p n = do {
  me <- newTVar (AD (n, p, [], []));
  pMe <- newTVar me;
  writeTVar me (AD (n, p, [pMe], []));
  return pMe}

```

```

rePoint :: Ambient -> [Handle] -> STM ()
rePoint a [] = return ()
rePoint a (x:xs) = do writeTVar x a;
  rePoint a xs

```

Non-deterministic Lookup

```

choose :: [a] -> STM a
choose [] = retry
choose (x:[]) = return x
choose (x:xs) = or (return x) (choose xs)

```

```

assoc :: Name -> [(Name,Ambient)] -> [Ambient]
assoc f [] = []
assoc f ((a,x):xs) =
  if f==a then (x:assoc f xs)
  else assoc f xs
lookup' x l = choose (assoc x l)

```

References

- [1] ABADI, M., BIRRELL, A., HARRIS, T., AND ISARD, M. Semantics of transactional memory and automatic mutual exclusion. In *Proc. POPL'08* (2008), pp. 63–74.
- [2] ACCIAI, L., BOREALE, M., AND DAL-ZILIO, S. A concurrent calculus with atomic transactions. In *Proc. ESOP'07* (2007), R. D. Nicola, Ed., vol. 4421 of *LNCS*, Springer, pp. 48–63.
- [3] BORGSTRÖM, J., BHARGAVAN, K., AND GORDON, A. D. A compositional theory for STM Haskell. Tech. Rep. MSR-TR-2009-66, Microsoft Research, 2009.
- [4] CARDELLI, L. Mobile ambient synchronization. Technical Note 1997-013, Digital Equipment Corporation, Systems Research Center, 1997.
- [5] CARDELLI, L., AND GORDON, A. D. Mobile ambients. *Theoretical Computer Science* 240 (2000), 177–213.
- [6] FOURNET, C., LÉVY, J.-J., AND SCHMITT, A. An asynchronous, distributed implementation of mobile ambients. In *Proc. TCS'00* (2000), Springer, pp. 348–364.
- [7] GIANNINI, P., SANGIORGI, D., AND VALENTE, A. Safe ambients: Abstract machine and distributed implementation. *Science of Computer Programming* 59, 3 (2006), 209–249.
- [8] GUERRAOU, R., HENZINGER, T. A., AND SINGH, V. Completeness and nondeterminism in model checking transactional memories. In *Proc. CONCUR'08* (2008), F. van Breugel and M. Chechik, Eds., vol. 5201 of *LNCS*, Springer, pp. 21–35.
- [9] HARRIS, T., AND FRASER, K. Language support for lightweight transactions. In *Proc. OOPSLA'03* (2003), pp. 388–402.
- [10] HARRIS, T., MARLOW, S., PEYTON JONES, S., AND HERLIHY, M. Composable memory transactions. In *Proc. PPOPP'05* (2005), K. Pingali, K. A. Yelick, and A. S. Grimshaw, Eds., ACM, pp. 48–60.
- [11] HARRIS, T., MARLOW, S., PEYTON JONES, S., AND HERLIHY, M. Composable memory transactions. *Communications of ACM* 51, 8 (2008), 91–100.
- [12] HARRIS, T., AND PEYTON JONES, S. Transactional memory with data invariants. In *Proc. TRANSACT'06* (2006).
- [13] HU, L., AND HUTTON, G. Towards a verified implementation of software transactional memory. In *The Symposium on Trends in Functional Programming* (2008). To appear.
- [14] HUCH, F., AND KUPKE, F. A high-level implementation of composable memory transactions in Concurrent Haskell. In *Proc. Implementation and Application of Functional Languages* (2005), vol. 4015 of *LNCS*, Springer, pp. 124–141.
- [15] HUGHES, J. Why functional programming matters. *Computer Journal* 32, 2 (Apr. 1989), 98–107.
- [16] JEFFREY, A., AND RATHKE, J. A theory of bisimulation for a fragment of concurrent ML with local names. *Theoretical Computer Science* 323, 1–3 (2004), 1–48.
- [17] KOUTAVAS, V., AND WAND, M. Small bisimulations for reasoning about higher-order imperative programs. In *Proc. POPL'06* (2006), ACM, pp. 141–152.
- [18] LEE, E. The problem with threads. *COMPUTER* (2006), 33–42.
- [19] MILNER, R. *Communication and Concurrency*. Prentice Hall, 1989.
- [20] MOORE, K. F., AND GROSSMAN, D. High-level small-step operational semantics for transactions. In *Proc. POPL'08* (2008), pp. 51–62.
- [21] NANEVSKI, A., GOVEREAU, P., AND MORRISSETT, G. Type-theoretic semantics for transactional concurrency. Tech. Rep. TR-08-07, Harvard University, July 2007.
- [22] O'LEARY, J., SAHA, B., AND TUTTLE, M. R. Model checking transactional memory with Spin. In *Proc. PODC'08* (2008), R. A. Bazzi and B. Patt-Shamir, Eds., ACM, p. 424.
- [23] OUSTERHOUT, J. Why threads are a bad idea (for most purposes). In *Presentation given at the 1996 Usenix Annual Technical Conference, January* (1996).
- [24] PEYTON JONES, S., GORDON, A., AND FINNE, S. Concurrent Haskell. In *Proc. POPL'96* (1996), pp. 295–308.
- [25] PHILLIPS, A., YOSHIDA, N., AND EISENBACH, S. A distributed abstract machine for boxed ambient calculi. In *Proc. ESOP'04* (2004), D. A. Schmidt, Ed., vol. 2986 of *LNCS*, Springer, pp. 155–170.
- [26] REGEV, A., PANINA, E. M., SILVERMAN, W., CARDELLI, L., AND SHAPIRO, E. BioAmbients: An abstraction for biological compartments. *Theoretical Computer Science* 325, 1 (2004), 141–167.
- [27] RINGENBURG, M. F., AND GROSSMAN, D. AtomCaml: first-class atomicity via rollback. In *Proc. ICFP '05* (2005), ACM, pp. 92–104.
- [28] SANGIORGI, D. On the bisimulation proof method. *Mathematical Structures in Computer Science* 8 (1998), 447–479.
- [29] SANGIORGI, D., KOBAYASHI, N., AND SUMII, E. Environmental bisimulations for higher-order languages. In *Proc. LICS'07* (2007), IEEE Computer Society, pp. 293–302.
- [30] SCOTT, M. L. Sequential specification of transactional memory semantics. In *Proc. TRANSACT'06* (2006).
- [31] SHAVIT, N., AND TOUITOU, D. Software transactional memory. *Distributed Computing* 10, 2 (1997), 99–116.
- [32] SUNSHINE-HILL, B., AND ZARKO, L. STM versus locks, ambiently, May 2008. CIS 552 Final Project, University of Pennsylvania.
- [33] VITEK, J., JAGANNATHAN, S., WELC, A., AND HOSKING, A. L. A semantic framework for designer transactions. In *Proc. ESOP'04* (2004), D. A. Schmidt, Ed., vol. 2986 of *LNCS*, Springer, pp. 249–263.