# Rolling Your Own Mutable ADT <br> - A Connection between Linear Types and Monads - 

Chih-Ping Chen and Paul Hudak

Yale University<br>Department of Computer Science

New Haven, CT 06520
\{chen-chih-ping, hudak\}@cs.yale.edu


#### Abstract

A methodology is described whereby a linear ADT may be rigorously encapsulated within a state monad. A CPS-like translation from the original ADT axioms into monadic ones is also described and proven correct, so that reasoning can be accomplished at the monadic level without exposing the state. The ADT axioms are suitably constrained by a linear type system to make this translation possible. This constraint also allows the state to be "updated in place," a notion made precise via a graph-rewrite operational semantics.


## 1 Introduction

In recent years, numerous proposals for I/O, destructive updates to data structures, mutable variables, nondeterminism, and concurrency have been put forth, all using monads to structure programs in such a way that details of the computation are effectively hidden and encapsulated [18, 20]. One of the most important uses of monads is in dealing with state, resulting in a style of programming referred to appropriately by Peyton Jones and Wadler as imperative functional programming [16]. In this style, a state monad is used to encapsulate the state in an abstract datatype, thus avoiding the "plumbing" usually associated with pure functional programs that manipulate state. Access to the hidden state is achieved by abstract operations that are sequenced via a small set of monadic combinators, typically called bind and unit. This style of programming has been adopted wholeheartedly by the Haskell [8] community, and its utility has been verified through many non-trivial ap-

[^0]plications.
Surprisingly, however, little work has been done in establishing a formal connection between monads and state. In particular:

- There are few formal connections between monads and single-threaded state, where the property of being single-threaded is precisely what is needed to guarantee safe, "in-place" update. The severity of this problem is easily demonstrated by an example: For any state-transformer monad with type $M a=\mathcal{S} \rightarrow(\mathcal{S}, a)$, there is a trivial operation getState :: $M \mathcal{S}$ that will instantly destroy any hope for single-threadedness:

$$
\text { getState } s=(s, s)
$$

- Nor are there many formal connections between monads and the state-manipulating operations themselves. That is, the standard monad operations tell us how the state-manipulating operations may be sequenced, but say nothing about what they do, or how they are derived.
- Finally, although laws for reasoning about several specific monads exist [18, 15], there is no general way to derive these laws.
To solve these problems, we outline a methodology that allows one to begin with a conventional axiomatization of a type of interest, rigorously encapsulate it within a monad, and reason about it abstractly in a monadic fashion. This is only possible if the axiomatization satisfies a certain linearity condition, established using linear types $[6,9,4,19,5]$. This same condition allows us to prove that the type of interest (i.e. the state) may be updated in place.

Our combination of linear types and monads unveils a connection between two very different methods for
dealing with state in a functional language. It is interesting that a CPS-like conversion, with strong ties to notions of sequentiality, formalizes this connection.

From a pragmatic perspective, the methodology allows one to use conventional ADT techniques to design a type of interest and operations on it. The use of the linear type system is confined just to the $A D T$ axioms. Once the axioms type-check, the linear type system never needs to be dealt with again. In other words, the monad encapsulates not just the ADT, but its linearity as well.

The methodology is also easily implementable (although we have not yet done so). The graph-rewrite semantics is quite similar to conventional graph reduction, the basis of many functional language implementations. Moreover, the monadic axioms can be derived automatically, yielding a source of program transformation rules for compile-time optimization. These rules may also be used by semantics-directed compilers and interpreters that are based on monads [12].

An Example. Consider a simple integer list ADT, with operations:

$$
\begin{array}{lll}
\text { nil } & :: & \text { IntList } \\
\text { cons } & :: & \text { Int } \rightarrow \text { IntList } \rightarrow \text { IntList } \\
\text { nth-select } & :: & \text { Int } \rightarrow \text { IntList } \rightarrow \text { Int } \\
\text { nth-update } & :: & (\text { Int, Int }) \rightarrow \text { IntList } \rightarrow \text { IntList }
\end{array}
$$

These operations are axiomatized by:

$$
\begin{array}{lll}
\text { nth-select } 0(\text { cons } x x s) & \Rightarrow & x \\
\text { nth-select }(i+1)(\text { cons } x x s) & \Rightarrow & \text { nth-select } i x s \\
\text { nth-update }(0, v)(\text { cons } x x s) & \Rightarrow & \text { cons } v x s \\
\text { nth-update }(i+1, v)(\text { cons } x x s) & \\
\quad \Rightarrow \text { cons } x(\text { nth-update }(i, v) x s)
\end{array}
$$

This axiomatization is well typed with respect to the linear type system described in Section 3. Consequently, a mutable version of this ADT can be derived automatically, with operations encapulated in a monad having types:

$$
\begin{array}{lll}
\text { nilM } & :: & \text { M Int } \rightarrow \text { Int } \\
\text { consM } & :: & \text { Int } \rightarrow M() \\
\text { nth-selectM } & :: & \text { Int } \rightarrow M \text { Int } \\
\text { nth-updateM } & :: & (\text { Int }, \text { Int }) \rightarrow M()
\end{array}
$$

The following small program which manipulates this mutable ADT:

$$
\begin{array}{ll}
\text { nilM } & \\
(\text { consM } 2 & \gg \\
\text { consM 1 } & \gg \\
\text { nth-updateM }(1,5) & \gg \\
\text { nth-selectM 1 } &
\end{array}
$$

returns the value 5. Furthermore, using the graphrewrite semantics of Section 5, nth-updateM will update the list in-place (in contrast with the original ADT, which must recreate the first $n$ elements). In this example the benefits are only constant-factor amortized improvements in time and space, but in other examples the improvements are more dramatic (see Appendix A for the design of an array ADT, where a linear-factor improvement is realized). Even constant-factor improvements may be important in certain situations, for example if the improvement can reduce load on the garbage collector (we will show later that nth-updateM induces only a constant number of heap or stack allocations).

Finally, a new set of axioms for the ADT can be automatically obtained by applying the translation scheme in Section 4.3 to the original axioms above, yielding:

```
consM x > nth-selectM 0}>>\lambdaz->
    consMx>m[x/z]
consMx>nth-selectM (i+1) >> \z->m
    nth-selectM i}>>\lambdaz->\mathrm{ consM }x>
consM x > nth-updateM (0,v)
    consMv
consMx> nth-updateM (i+1,v)
    # nth-updateM (i,v) > consM x
consM x > unit y }=>\mathrm{ unit y
nilM (unit y) }=>
```

$(\geq$ and $\gg$ are the monadic bind and seq operators, respectively.) Note that these axioms allow reasoning about integer lists at the level of monads, without exposing the underlying representation.

## 2 Preliminaries

Define the language $\mathcal{L}$ as the lambda calculus extended with constants, which include primitive datatypes and operations on them, such as the unit type (), integers, booleans, and pairs. Its syntax is given in Figure 1. In addition, for convenience we will often use a nonrecursive let expression defined by:

$$
\text { let } x=e_{1} \text { in } e_{2} \equiv\left(\lambda x \rightarrow e_{2}\right) e_{1}
$$

The only reduction rules that interest us at this stage are the conventional $\beta$ rule:

$$
(\lambda x \rightarrow e) y \Rightarrow_{\beta} \quad e[y / x]
$$

and the $\delta$ rules that govern the behavior of conditional expressions, primitive datatypes, and recursion. For the conditional and pairing, we have:

$$
\begin{gathered}
\text { if True } \quad e_{1} \quad e_{2} \\
\text { if False } \quad e_{\delta}
\end{gathered} e_{2} \quad e_{1}
$$



Figure 1: Syntax of $\mathcal{L}$

Note that the above rules imply a non-strict semantics.
For recursion, we assume $f i x \in C o n$ with $\delta$ rule:

$$
\text { fix } f \Rightarrow_{\delta} \quad f(\text { fix } f)
$$

Define the relation $\Rightarrow$ as the union of $\Rightarrow_{\beta}$ and $\Rightarrow_{\delta}$. We write $e_{1} \Downarrow_{C B N} e_{2}$ to denote the reduction of $e_{1}$ to $e_{2}$ using $\Rightarrow$ under a call-by-name (leftmost, outermost) reduction strategy. Evaluation of an expression $e$ is defined as reduction of $e$ to a value $v$, defined as: ${ }^{1}$

$$
v::=k \mid\left(v_{1}, v_{2}\right)
$$

Evaluation is a partial function.

### 2.1 Abstract Datatypes (ADTs)

Definition 1 Given a type of interest $T$, a simple $A D T$ is one in which each operation can be classified as either a generator of type $X_{1} \rightarrow T$, a modifier of type $X_{2} \rightarrow$ $T \rightarrow T$, or a selector of type $X_{3} \rightarrow T \rightarrow X_{4}$, where the $X_{i}$ are arbitrary auxiliary datatypes.

This is a standard classification for ADTs, and most conventional ADTs such as arrays, lists, stacks, and queues can be expressed in this way. However, it does not include tree-shaped ADT's, since a modifier for such an ADT would need a type such as $X \rightarrow T \rightarrow T \rightarrow T$. This limitation is discussed further in Section 6.

We allow $\mathcal{L}$ to be extended with simple ADTs by extending the base types as necessary and using the ADT axioms as $\delta$ rules in the calculus. We assume that adding new axioms does not destroy confluence (which can also be achieved by choosing a determinate order on the delta rules, as is done in most modern functional languages). For an $\operatorname{ADT} \mathcal{A}$ we refer to the extended language as $\mathcal{L}^{\mathcal{A}}$.

Definition 2 Given a simple ADT with a set of generators $G$, modifiers $M$, and selectors $S$, an axiomatization distinguishes a set $D=S \cup U$, where $U \subseteq M$. Each operator $d \in D$ has one or more defining axioms whose syntax is given in Figure 2.

[^1]This definition subdivides modifiers into two categories: $U$ is the set of mutators, each having defining axioms; and the rest $(M \backslash U)$ are called constructors, whose behavior is defined implicitly by the axioms. For example, for the integer list ADT in the introduction, nil is a generator, cons is a constructor, nth-update is a mutator, and nth-select is a selector.

Note that the syntax in Figure 2 restricts the axioms to be first-order.

## 3 The Linearity Condition

Not every ADT can be converted to monadic form; it must be "linear," or "single-threaded," in some sense. In this section we describe a linear type system which constrains the ADT sufficiently to allow the conversion. It is inspired by the linear type system with read-only access in [19], which in turn was influenced by the singlethreadedness condition in [17].

A type in our system can be a basic type, a pair of auxiliary types, a function type, or the type of interest $T$. In addition, it could be a linear type of interest, denoted $; T$. Formally: ${ }^{2}$

$$
\begin{array}{lcc|c|c}
u, v & := & \text { Int } & \text { Bool } & \ldots \\
& \mid & (u, v) & (u \rightarrow v) & \\
& T & i T &
\end{array}
$$

Facilitated with this finer-grained type system, we are able to type each class of ADT operators to better reflect the intended use of $T$ :

$$
\begin{array}{ccc}
g & : & X_{1} \rightarrow \mathfrak{¡} T \\
m & : & X_{2} \rightarrow \mathfrak{i} T \rightarrow \mathfrak{i} T \\
s & : & X_{3} \rightarrow T \rightarrow X_{4}
\end{array}
$$

Generators and modifiers manipulate $; T$ to ensure that at any time they either create or have the sole pointer to the type of interest. On the other hand, selectors take a non-linear $T$ as an argument because they only read the type of interest and, as a result, the linearity condition can be relaxed.

[^2]
$g \in G \quad$ (generators)
lhs $\quad::=m x$ pat $\mid$ s x pat
$s \in S \quad$ (selectors)
pat $::=t|g x| c x$ pat

Figure 2: Axiom Syntax Scheme

An assumption list associates variables and constants with types:

$$
A::=x_{1}: u_{1}, \ldots, x_{n}: u_{n}
$$

where $x_{1}, \ldots, x_{2}$ are distinct variables or constants. The base assumption list $B_{\mathcal{A}}$ for an $\mathrm{ADT} \mathcal{A}$ associates all the operators of $\mathcal{A}$, primitives, and constants with their types. We write $A, B$ to denote the concatenation of two assumption lists; an implicit side condition of a concatenation is that $A$ and $B$ contain distinct variables.
$R(A)$ is an assumption list identical to $A$ except that all the associations of mutators are elided, and those of generators, constructors, and variables of type $¡ T$ are "non-linearized." That is, the i's in their type signatures are removed. Intuitively, $R$ prepares a "read-only" assumption list for an expression which may do multiple reads on the type of interest.

The typing rules are listed in Figure 3. The predicate $N L(u)$ ("non-linear $u$ ") checks that $u$ does not have a linear type of interest as a component. It is defined by:

$$
\begin{aligned}
N L(u) \Longleftrightarrow & u \neq \mathrm{i} T \text { and } \\
& u=(v, w) \Longrightarrow N L(v) \text { and } N L(w)
\end{aligned}
$$

The predicate $H Y(u)$ ("hygienic $u$ ") checks that $u$ is not a function type and does not have any type of interest, either linear or nonlinear, as a component. It is defined by:

$$
\begin{aligned}
H Y(u) \Longleftrightarrow & u \neq v \rightarrow w \text { and } \\
& u \neq i T \text { and } \\
& u \neq T \text { and } \\
& u=(v, w) \Longrightarrow H Y(v) \text { and } H Y(w)
\end{aligned}
$$

Note that our system allows (Weakening) on any variable because the "no discarding" property of a linear type of interest is not of our concern. This design in turn enables the use of only one (Identity) rule. On the other hand, (Contraction) can only be applied to a variable of non-linear type because the "no duplicate" property of a linear type of interest is the theme of the whole type system. There are three application rules: (¡T-Application) is identical to the application
rules in other linear type system except that here we constrain the resulting type of the application to be $¡ T$ ( T -Application) allows read-only access in the argument term whose type is $T$, and the hygienic condition ensures that no type of interest can be smuggled out (a side-effect of this constraint is that the resulting type can not be a function type); (Application) handles applications whose operands are of hygienic type. Also note that read-only access is used when typing the predicate of a conditional expression.

Definition 3 An ADT $\mathcal{A}$ is linear if each axiom $l h s \Rightarrow r h s$ has the following properties:

- There exists a type $u$ and assumption list $A$ such that:

$$
\begin{aligned}
& -B_{\mathcal{A}}, A \vdash \text { lhs }: u, \text { and } \\
& - \text { If } u=\mathrm{i} T \text {, then } B_{\mathcal{A}}, A \vdash r h s: u ; \\
& \quad \text { otherwise, } R\left(B_{\mathcal{A}}, A\right) \vdash r h s: u .
\end{aligned}
$$

- Each free variable has exactly one occurrence in lhs.

As an example of a well-typed linear ADT, consider the integer list ADT given in the introduction. As an example of an ADT that would not be well-typed, consider adding the following (contrived) operator to the integer list ADT:

$$
\begin{aligned}
& \text { silly } l \Rightarrow \text { let } l^{\prime}=\text { nth-update }(1,2) l \\
& \text { in nth-select } 1 l^{\prime}
\end{aligned}
$$

This is not linear because silly is a selector, but there is an update to the type of interest on the right-hand side, which our type system does not allow.

## 4 Monadic ADTs

To say that an ADT is linear is to say something about its axioms. But if we naively implement a linear ADT in a functional language based on a Hindley-Milner type system, the linear axioms won't buy us much: we still will not be able to do in-place updates on the type of

$$
\begin{aligned}
& \overline{x: u \vdash x: u} \text { (Identity) } \frac{A, y: u, z: u, B \vdash e: v \quad N L(u)}{A, x: u, B \vdash e[x / y, x / z]: v} \text { (Contraction) } \\
& \frac{A \vdash e: v}{A, x: u \vdash e: v} \text { (Weakening) } \quad \frac{A, x: u, y: v, B \vdash e: t}{A, y: v, x: u, B \vdash e: t} \text { (Exchange) } \\
& \frac{R(A) \vdash p: \text { Bool } \quad A \vdash e_{1}: u \quad A \vdash e_{2}: u}{A \vdash \text { if } p \text { then } e_{1} \text { else } e_{2}: u} \text { (If) } \\
& \frac{R(A) \vdash e_{1}: u \quad R(A) \vdash e_{2}: v \quad H Y(u) \quad H Y(v)}{A \vdash\left(e_{1}, e_{2}\right):(u, v)} \text { (Pair) } \\
& \frac{A \vdash e_{1}: \mathfrak{i} T \rightarrow \mathfrak{i} T \quad B \vdash e: \mathrm{i} T}{A, B \vdash e_{1} e: ¡ T}(\mathrm{iT} \text {-Application) } \\
& \frac{A \vdash e_{1}: T \rightarrow v \quad R(B) \vdash e_{2}: T \quad H Y(v)}{A, B \vdash e_{1} e_{2}: v} \text { (T-Application) } \\
& \frac{A \vdash e_{1}: u \rightarrow v \quad B \vdash e_{2}: u \quad H Y(u)}{A, B \vdash e_{1} e_{2}: v} \text { (Application) }
\end{aligned}
$$

Figure 3: Linear Typing Rules
interest because the type system is not strong enough to guarantee that the argument to an ADT mutator is unshared. For example, the program:

$$
\text { let } \begin{aligned}
& l_{1}=\text { cons } 1(\text { cons } 2 \text { nil }) \\
& \quad l_{2}=\text { nth-update }(1,5) l_{1} \\
& \text { in nth-select } 1 l_{2}+\text { nth-select } 1 l_{1}
\end{aligned}
$$

should return 7. But if the update were done destructively, the result might be 10 , which is clearly wrong.

We could, of course, enrich the type system with linear types throughout the language, in which case the above program would be ill-typed (at least one functional language, Clean, uses this approach [2]). We choose instead a different solution: we will encapsulate the linear ADT in a monad, and use the monadic structure to enforce linearity of the type of interest throughout the language.

### 4.1 Monad Basics

Definition 4 A monad is a triple $(M, \gg$, unit) where $M$ is a type constructor, and $\gg$ and unit are functions which satisfy the following three laws:

$$
\begin{aligned}
\text { unit } a \ggg b \rightarrow n= & n[a / b] \\
n \gg \text { unit }= & n \\
n_{1} \geqq\left(\lambda a \rightarrow n_{2} \geqq \lambda b \rightarrow n_{3}\right)= & \left(n_{1} \gg \lambda a \rightarrow n_{2}\right) \\
& \geqq \lambda b \rightarrow n_{3}
\end{aligned}
$$

In the last law, $a$ does not appear free in $n_{3}$.
Operationally speaking, $M a$ is a computation which produces an answer of type $a$ when executed; $\gg$ combines two computations into a larger one which performs the two in sequence; and unit injects a single value into a computation. Of particular interest to us is the state monad which abstracts over a type of interest to be treated as mutable state:

$$
\begin{array}{ll}
\text { type } M a & =T \rightarrow(a, T) \\
\gg & :: M a \rightarrow(a \rightarrow M b) \rightarrow M b \\
\text { unit } & :: a \rightarrow M a \\
(n \gg k) t & \Rightarrow \text { let } p=n t \quad \text { in } \quad k(\text { fst } p)(\text { snd } p) \\
(\text { unit } a) t & \Rightarrow(a, t)
\end{array}
$$

It can be easily verified that these definitions satisfy the three monad laws.

For convenience, we specialize a version of $\gg$ which we write $>$ :

$$
\left(n_{1} \gg n_{2}\right) t \Rightarrow n_{2}\left(\operatorname{snd}\left(m_{1} t\right)\right)
$$

$\geqq$ and $\gg$ are like Haskell's $\gg=$ and $\gg$, respectively, and are sometimes called bind and seq.

### 4.2 Monadic Encapsulation of ADT's

For each modifier $m:: X_{2} \rightarrow T \rightarrow T$ and selector $s::$ $X_{3} \rightarrow T \rightarrow X_{4}$ in an ADT $\mathcal{A}$, we introduce monadic counterparts $m M$ and $s M$, respectively, defined in terms of $m$ and $s$ :

$$
\begin{array}{lll}
m M & :: & X_{2} \rightarrow M() \\
s M & :: & X_{3} \rightarrow M X_{4} \\
m M x t & \Rightarrow & ((), m x t) \\
s M x t & \Rightarrow & (s x t, t)
\end{array}
$$

The typing $M$ () which results from applying $m M$ implies that the modifier only effects the type of interest and does not return anything interesting.

In addition, for each generator $g:: X_{1} \rightarrow T$ in an ADT $\mathcal{A}$ we define a monadic generator $g M$ by:

$$
\begin{array}{ll}
g M & :: X_{1} \rightarrow M a \rightarrow a \\
g M x m & \Rightarrow f_{s t}(m(g x))
\end{array}
$$

This definition reflects the use of $g M$ as the operator which invokes a computation: it applies the monadic computation to the type of interest generated by the original ADT generator and returns the first element of the resulting pair of the computation while discarding the second element, the (possibly altered) type of interest.

### 4.3 Deriving Monadic Axioms

The encapsulation just described leads to a monadic implementation of a linear ADT, and we will see in a later section that it allows in-place update of the type of interest. But it is not an especially good mechanism for reasoning about a monadic ADT, since it reveals the details of the encapsulation, and, once the original ADT is exposed, reverts to reasoning in terms of that ADT. It would be better if we had monadic axioms, corresponding to the original linear ones, that did not expose the underlying ADT representation. In this section we describe a method for deriving such axioms from those in the original linear ADT.

Our translation scheme resembles, perhaps not surprisingly, a CPS conversion, and is related to call-byvalue monad transformation $[13,18]$. But it is also different from these in several ways, the two most important being: (1) the conversion is performed on axioms rather than just terms, and (2) the conversion is done with respect to a particular value - the type of interest and in fact is not even valid for non-linear ADTs.

In the following, we use $\lceil\cdots\rceil$ and $\langle\cdots\rangle$ to quote and un-quote, respectively, the text being translated. First we define a translation function $\mathcal{D}$ for expressions that, given an expression $e$ and a continuation $k$, returns an
equivalent expression using the monadic ADT:

$$
\begin{aligned}
& \mathcal{D}\lceil s\langle x\rangle\langle t\rangle\rceil k=\mathcal{D} t(\mathcal{D} x\lceil\lambda y \rightarrow s M y \geqq\langle k\rangle\rceil) \\
& \mathcal{D}\lceil m\langle x\rangle\langle t\rangle\rceil k=\mathcal{D} t(\mathcal{D} x\lceil\lambda y \rightarrow m M y \gg\langle k\rangle\rceil) \\
& \mathcal{D}\lceil g\langle x\rangle\rceil k=\lceil g M\langle x\rangle\langle k\rangle\rceil \\
& \mathcal{D}\lceil\text { if }\langle p\rangle \text { then }\langle c\rangle \text { else }\langle a\rangle\rceil k \\
& =\mathcal{D} p\lceil\lambda x \rightarrow \text { if } x \text { then }\langle\mathcal{D} \text { c } k\rangle \text { else }\langle\mathcal{D} \text { a } k\rangle\rceil \\
& \mathcal{D}\left\lceil\left(\left\langle e_{1}\right\rangle,\left\langle e_{2}\right\rangle\right)\right\rceil k \\
& =\mathcal{D} e_{1}\left\lceil\lambda y \rightarrow\left\langle\mathcal{D} e_{2}\lceil\lambda z \rightarrow(\text { unit }(y, z)) \geqq\langle k\rangle\rceil\right\rangle\right\rceil \\
& \mathcal{D}\left\lceil f\left\langle e_{1}\right\rangle\left\langle e_{2}\right\rangle\right\rceil k \\
& =\mathcal{D} e_{1}\left\lceil\lambda y \rightarrow\left\langle\mathcal{D} e_{2}\lceil\lambda z \rightarrow(\text { unit }(f y z)) \geqq\langle k\rangle\rceil\right\rangle\right\rceil \\
& \mathcal{D}\lceil t\rceil k=k \\
& \mathcal{D}\lceil x\rceil(\lambda y \rightarrow n)=n[x / y] \\
& \mathcal{D}\lceil c\rceil(\lambda y \rightarrow n)=n[c / y]
\end{aligned}
$$

Here $s, m$, and $g$ represent an arbitrary selector, modifier, and generator, respectively; $s M, m M$, and $g M$ are their monadic counterparts; $t$ is an identifier whose type is the type of interest; $x$ is an identifier having auxiliary type; $f$ is a primitive function; and $c$ is a constant having auxiliary type. We assume suitable $\alpha$-renaming to avoid name clashes with $y, z$, and $k$.

Given $\mathcal{D}$, we can now define translations $\mathcal{S}$ and $M$ for axioms that define selectors and mutators, respectively.

$$
\begin{aligned}
\mathcal{S} & \left\lceil s\left\langle x_{1}\right\rangle\langle t\rangle \Rightarrow\left\langle x_{2}\right\rangle\right\rceil \\
& =\left\lceil\left\langle\mathcal{D}\left\lceil s\left\langle x_{1}\right\rangle\langle t\rangle\right\rceil\lceil(\lambda y \rightarrow n)\rceil\right\rangle\right. \\
& \left.\Rightarrow\left\langle\mathcal{D} x_{2}\lceil\lambda y \rightarrow\langle\mathcal{D} t\lceil n\rceil\rangle\rceil\right\rangle\right\rceil \\
M & \left\lceil m\langle x\rangle\left\langle t_{1}\right\rangle \Rightarrow\left\langle t_{2}\right\rangle\right\rceil \\
& =\left\lceil\left\langle\mathcal{D}\left\lceil m\langle x\rangle\left\langle t_{1}\right\rangle\right\rceil\lceil k\rceil\right\rangle \Rightarrow\left\langle\mathcal{D} t_{2}\lceil k\rceil\right\rangle\right\rceil
\end{aligned}
$$

Here also we assume suitable $\alpha$-renaming to avoid name clashes with $y, z$, and $k$. Note that $\mathcal{S}$ translates the right-hand side in the context of the nested type of interest on the left-hand side, whereas $M$ translates both sides independently. Also note that if an axiom is defined in terms of type of interest identifier $t$, then the type of interest disappears during the translation (thus completing our mission of hiding the state!).

In addition to the axioms derived from the original ADT axioms, we also need to axiomatize:

- the interactions between generators and unit. For each generator $g M$, add the rule:

$$
g M(\text { unit } y) \Rightarrow y
$$

- the interactions between constructors and $\gg$. For each constructor $c M$, add the rule:

$$
c M x \gg \text { unit } y \Rightarrow \text { unit } y
$$

An example of the overall translation is the monadic axioms for the integer list ADT given in the introduction. It is interesting to note in this translation process
that mutators and selectors trade roles! This is an artifact of the CPS-like nature of the translation, which uncovers a "duality" between constructors and selectors, much like the duality noted by Filinski in [3].

Calling the above translation scheme $\mathcal{T}$, its correctness is captured by its relationship to the monadic encapsulation previously defined:

Theorem 1 The monadic encapsulation of a linear ADT $\mathcal{A}$ satisfies the monadic axiomatization derived from $\mathcal{A}$ using translation $\mathcal{T}$.

Proof: By applying the translation scheme to a generic axiom of each kind, then using equational reasoning to unfold the monadic encapsulation on each side and reduce them to equivalent terms. A proof for one of the three axioms is given in Appendix B.

## 5 GRS Operational Semantics

In [7] we informally defined:
Definition 5 A mutable ADT, or MADT, is any ADT whose operational semantics permits in-place update (i.e. destructive reuse) of the type of interest, while still retaining confluence.

We will now make this definition more precise. To do so, we need an operational semantics that captures sharing and identifies all relevant costs of execution. The evaluation function $\Downarrow_{C B N}$ is not suitable for this purpose since it is essentially a term-rewriting system. Instead, we choose a graph-rewriting system (GRS) to capture operational semantics. Although other approaches to operational semantics abound - such as recent efforts at modeling lazy evaluation [10, 1, 14]-we find a GRS as particularly clear (a picture is worth a thousand words), abstract (no heap, for example), expressive (capturing notions of sharing, updating, and sequencing), and intuitive (resembling conventional graph reduction).

Figure 4 defines a function lift which maps terms of $\mathcal{L}$ into their corresponding graphs (called $\mathcal{L}$-graphs). For each CBN $\delta$ rule, a corresponding GRS rule can be obtained by lifting both sides of the rule. For example, Figure 5 shows the GRS rules for $f s t$ and $s n d$. The GRS rule corresponding to the CBN $\beta$ rule is also shown in Figure 5, where the notation $e\{y / x\}$ denotes the result of redirecting the edges which point to the bound variable $x$ in the subgraph $e$ to the subgraph $y$; thus $y$ is shared (also, as in [21], we assume that a fresh copy of the function body $e$ is created before the subsitution is made).

We write $e_{1} \Downarrow_{G R S} e_{2}$ to denote the reduction of $e_{1}$ to $e_{2}$ using the GRS-lifted reduction $\Rightarrow$ under a call-by-name (leftmost, outermost) reduction strategy. The
evaluation function $\Downarrow_{G R S}$ can be seen as the call-by-need equivalent of the call-by-name evaluator $\Downarrow_{C B N}$.
Theorem 2 For every term e in $\mathcal{L}$,

$$
e \Downarrow_{\mathrm{CBN}} v \quad \Longleftrightarrow \quad \operatorname{lift}(e) \Downarrow_{\mathrm{GRS}} v
$$

where $v$ is a value.
Proof: See [21].

### 5.1 Lifting of ADT's

The lifting operation can also be extended to an ADT by lifting each side of each axiom. For example, the GRS rules for the two nth-update axioms given in the introduction are shown in Figure 6.

When implementing GRS rules-for example using conventional graph reduction techniques - all of the reductions are assumed to be "pure." That is, other than the top vertex on each of the left- and right-hand sides, no vertices are assumed to be shared. For example, in the first $n$ th-update rule of Figure 6, the top vertex on the left is assumed to be the same as the top vertex on the right, but the vertices that comprise the cons operation on the right are assumed to be "fresh" vertices, distinct from those on the left. This is necessary, of course, to preserve confluence, as we have discussed earlier. But this is also exactly where we would like to do better: if some vertices on the left are discarded as a result of applying the rule, there is no reason why they can't be immediately re-used on the right, thus achieving "in-place update." In those cases where we abandon the "pure" reduction strategy in a rule, we will annotate the vertices with names to indicate which are being reused, and how.

### 5.2 GRS Rules for Encapsulated ADT

The monad laws can also be lifted into the GRS, as shown in Figure 7, where $t$ denotes the subgraphs representing state. Note in the rule for unit that the state $t$ is single-threaded; i.e. $t$ is reachable from the root of each side of the rule in exactly one way. On the other hand, the rule for $\gg$ is not single-threaded, at least not in this simplistic way: in reality it depends on the behavior of $k$ and $n$. This observation hints that if:

- a linear ADT operator is only used "monadically," i.e. having type $M a$ and only combined with other monadic computations via $\geqq$, and
- a fresh (unshared) copy of the type of interest is passed into a monadic computation when it is invoked,

$$
\begin{aligned}
& \operatorname{lift}(x)=x \quad \operatorname{lift}(k)=k \quad \operatorname{lift}(\lambda x->e)=\left.\right|_{\operatorname{lift}(e)} ^{\lambda x} \\
& \operatorname{lift}(e 1 \text { e2) }=\overbrace{\text { lift }(e 1)} \\
& \operatorname{lift}((e 1, e 2))=
\end{aligned}
$$

Figure 4: Lifting Terms to Graphs


Figure 5: GRS Rules for Application and Pair Selections


Figure 6: GRS Rules for nth-update


Figure 7: GRS Rules for Monadic Combinators


Figure 8: GRS Rules for Monadic ADT Operators


Figure 9: GRS Rules for nth-update Reusing Vertices
then the single-threadedness of the type of interest can be guaranteed throughout the entire computation.

We can achieve this effect by first lifting the encapsulation strategy of Section 4.2 into the GRS, as shown in Figure 8. However, the GRS rule for $m M$ will have one important twist: the mutator $m$ that it exposes will be allowed to reuse the type of interest; i.e. in-place update is allowed. To see this concretely, Figure 9 shows the rules for $n$ th-update, using the convention of vertex reuse described earlier. It is interesting to note that the second rule implies that no vertices are needed to realize the recursive call structure of nth-update; that is, it is fully tail-recursive, or iterative. Thus, in our operational semantics, an application of nth-updateM requires allocation of exactly two vertices, regardless of the length of the list.

There is one problem, however: the right-hand side of the rule for $s M$ introduces sharing in $t$, but using CBN reduction there is no guarantee that the lifted graph of ( $\left.\begin{array}{lll}s & x & t\end{array}\right)$ on the right-hand side is reduced before some subsequent mutation to $t$. We can fix this problem by making fst and snd hyper-strict in the first component of their arguments. This is the only sensible way to ensure that there are no dangling references to the type of interest, and is similar to the use of a hyper-strict let expression in [19]. Thus we define new versions of the pair selectors, $f s t A$ and $s n d A$ :

$$
\begin{array}{llll}
f s t A & (v, e) & \Rightarrow_{\delta} & v \\
\operatorname{sndA} & (v, e) & \Rightarrow_{\delta} & e
\end{array}
$$

where $v$ is a value. We then use these to redefine $\geq$ and $\gg$ :

$$
\begin{array}{ll}
(m \gg k) t & \Rightarrow \operatorname{let} p=m t \text { in } k(\text { fstA } p)(\text { sndA } p) \\
(m \gg n) t & \Rightarrow n(\operatorname{sndA}(m t))
\end{array}
$$

The GRS versions of these rules are exactly the same as those in Figure 7 except that $f s t$ is replaced by $f s t A$ and snd by sndA.

We can see that using hyper-strict pair selectors restores single-threadedness by reasoning as follows:

- If $f s t A(s M x t)$ is reduced before $\operatorname{sndA}(s M x t)$, then $t$ becomes unshared because $f s t A$ forces $(s x t)$ to be fully evaluated.
- Similarly, if $s n d A(s M x t)$ is reduced first, $s n d A$ likewise forces ( $s x t$ ) to be fully evaluated, so $t$ becomes unshared before it is further used.

The GRS system was necessarily made more strict in order to achieve single-threadedness. Thus it is possible that some computations which do not contribute to the evaluation of the final answer are evaluated (and may
not terminate). For example:

$$
\begin{array}{lll}
g M c_{1} & & \\
\left(s M c_{2}\right. & \gg & \lambda x \rightarrow \\
s M c_{3} & \geqq & \lambda y \rightarrow \\
\text { unit } y & & )
\end{array}
$$

where $x$ does not appear in $c_{3}$. The GRS would evaluate the subgraph representing $\left(s c_{2}\left(g c_{1}\right)\right)$, whose value is obviously irrelevant to the final answer. Nevertheless, GRS still offers the beauty of lazy evaluation in many cases. For example:

$$
\begin{array}{lll}
g M c_{1} \\
\left(m M c_{2}\right. & \geqq & \lambda x \rightarrow \\
m M c_{3} & \geqq & \lambda y \rightarrow \\
\text { unit } 1 & & )
\end{array}
$$

Here the GRS returns 1 without creating even the initial copy of the type of interest, let alone performing the two mutations.

### 5.3 Correctness of In-place Update

Using our somewhat stricter notion of GRS evaluation, we can prove a correctness result for in-place update. Define $\mathcal{L}^{\mathcal{M}}$ as the language $\mathcal{L}$ extended with any monadically encapulated, linear ADT $\mathcal{M}$.

Theorem 3 For every term e in $\mathcal{L}^{\mathcal{M}}$,

$$
e \Downarrow_{\mathrm{CBN}} v \quad \Longleftrightarrow \quad \operatorname{lift}(e) \Downarrow_{\mathrm{GRS}} v
$$

where $v$ is a value.
Proof: See Apendix C.

## 6 Discussion

In this section we discuss limitations and possible extensions of our methodology.

Although we believe that the added strictness introduced in Section 5.2 is minimal, and that some added strictness is inherent in any single-threaded reduction system, we are investigating ways to improve this, such as dynamically tagging closures having no references to the type of interest. But it is not clear that such a small improvement is worthwhile.

Polymorphism should be easy to handle, at the expense of carrying around more type information. For example, the integer list ADT could be turned into a polymorphic list, but the monad type would have to be extended to something like $M a b$, where $a$ is the type returned from the computation, and $b$ is the type of the list elements.

The linear type system proposed in Section 3 is limited to first-order values, so that, for example, a selector cannot return a function as its result. We believe, however, that the type system can be extended to handle abstractions, and we have worked out the preliminary ideas of such an extension.

A more serious limitation of the current system is that it does not handle more than one mutable ADT. One possible solution to this is to introduce references, so that individual objects can be named and manipulated, and use the extension to the Hindley-Milner type system proposed in [11] to prevent the escape of references from the monadic computation.

Related to this issue is the inability to handle treelike structures, ruled out by our definition of a simple ADT. Rather than using references, we have been working on a different solution to this problem, based on the following idea.

Suppose we are designing a database-like ADT, whose axiomatization is given by:

$$
\begin{array}{lll}
\text { new }:: D B \\
\text { insert } & :: \text { Int } \rightarrow D B \rightarrow D B \\
\text { remove }:: \text { Int } \rightarrow D B \rightarrow D B
\end{array} \quad \begin{aligned}
& \\
& \text { remove } a \text { new } \text { new } \\
& \text { remove } a(\text { insert } b x)= \text { if } a=b \text { then } x \\
&\text { else insert } b \text { (remove } a x)
\end{aligned}
$$

This is clearly linear, so our methodology can be applied, yielding a mutable ADT with monadic axioms, etc. However, it is not very efficient. So, as in a traditional ADT design, we may wish to implement the ADT using binary search trees. The implementation details should be hidden, of course, and we do so below using Haskell module syntax:

```
module DB (DB, new, insert, find, remove) where
data DB = Leaf | Branch Int Tree Tree
new = Leaf
insert a Leaf = Branch a Leaf Leaf
insert a (Branch b lt rt) =
    if a<=b then Branch b (insert a lt) rt
        else Branch b lt (insert a rt)
remove a Leaf = Leaf
remove a (Branch b lt rt) =
    if a==b
    then if lt==Leaf then rt else
        if rt==Leaf then lt else
        let pred = largest lt
            lt' = remove pred lt
        in Branch pred lt' rt
    else if a<=b then Branch b (remove a lt) rt
            else Branch b lt (remove a rt)
    where largest (Branch x l r) =
                        if r==Leaf then x else largest r
```

The interesting thing about this solution is that the tree is used linearly in the implementation viewed as
a set of axioms, only our type system is currently not strong enough to determine this. If it could be suitably strengthened, then this encapsulation of the axioms should allow in-place update of the tree.

## 7 Conclusions

Using state monads to achieve safe destructive reuse of the state is by no means an innovation; several earlier efforts have already demonstrated its suitability. What have been missing in these previous efforts are a systematic way to recognize when monads work and when they don't, an operational semantics to formally reason about the correctness of in-place updates, a method for designing efficient state monads, and a way to reason about them abstractly. Our contributions take positive steps toward resolving these deficiencies.

## 8 Acknowledgements

Thanks to Phil Wadler and the anonymous referees for their helpful comments on previous versions of this paper. This research was supported in part by NSF under Grant CCR-9404786.

## References

[1] Z. Ariola, M. Felleisen, J. Maraist, M. Odersky, and P. Wadler. A call-by-need lambda calculus. In Proc. 22nd ACM Symposium on Principles of Programming Languages, pages 233-241, New York, January 1995. ACM Press.
[2] E. Barendsen and S. Smetsers. Conventional and uniqueness typing in graph rewrite systems (extended abstract). In Proc. 13th Conference on the Foundations of Software Technology $\mathcal{E}^{3}$ Theoretical Computer Science, Bombay, India, 1993.
[3] A. Filinski. Declarative continuations: An investigation of duality in programming language semantics. In D. H. Pitt, editor, Category Theory and Computer Science, pages 224-249. Springer-Verlag, Berlin, 1989. (Lect. Notes in Comp. Science Vol. 389).
[4] J. Guzmán and P. Hudak. Single-threaded polymorphic lambda calculus. In Proc. Symposium on Logic in Computer Science, pages 333-343. IEEE, June 1990.
[5] J.C. Guzman. On Expressiong the Mutation of State in a Functional Programming Language. PhD thesis, Yale University, 1993.
[6] S. Holmstrom. A linear functional language. In Proc. the Workshop on the Implementation of Lazy Functional Languages, PMG Report 53, pages 1332, 1988.
[7] P. Hudak. Mutable abstract datatypes - or - how to have your state and munge it too. Research Report YALEU/DCS/RR-914, Yale University, Department of Computer Science, December 1992.
[8] P. Hudak, S. Peyton Jones, and P. Wadler (editors). Report on the Programming Language Haskell, A Non-strict Purely Functional Language (Version 1.2). ACM SIGPLAN Notices, 27(5), May 1992.
[9] Y. Lafont. The linear abstract machine. Theoretical Computer Science, 59:157-180, 1988.
[10] J. Launchbury. A natural semantics for lazy evaluation. In Proc. 20th ACM Symposium on Principles of Programming Languages, pages 144-154, New York, January 1993. ACM Press.
[11] J. Launchbury and S. Peyton Jones. Lazy functional state threads. In Proc. Symposium on Programming Language Design and Implementation '94, Orlando, June 1994. ACM.
[12] S. Liang and P. Hudak. Modular denotational semantics for compiler construction. In European Symposium on Programming '96, Linkoping, Sweden, April 1996.
[13] E. Moggi. Computational lambda-calculus and monads. In Proceedings of Symposium on Logic in Computer Science. IEEE, June 1989.
[14] G. Morriset, M. Felleisen, and R. Harper. Abstract models of memory management. In Proc. Conference on Functional Programming Languages and Computer Architecture, pages 66-77, New York, June 1995. ACM/IFIP, ACM Press.
[15] M. Odersky, D. Rabin, and P. Hudak. Call-byname, assignment, and the lambda-calculus. In Proc. 20th ACM Symposium on Principles of Programming Languages, Charleston, South Carolina, January 1993.
[16] S. Peyton Jones and P. Wadler. Imperative functional programming. In Proc. 20th Symposium on Principles of Programming Languages. ACM, January 1993. (to appear).
[17] D. A. Schmidt. Detecting global variables in denotational specification. ACM Transactions on Programming Languages and Systems, 7(2):299-310, 1985.
[18] P. Wadler. Comprehending monads. In Proc. Symposium on Lisp and Functional Programming, pages $61-78$, Nice, France, June 1990. ACM.
[19] P. Wadler. Linear types can change the world! In IFIP TC 2 Working Conference on Programming Concepts and Methods, Sea of Galilee, Isreal, April 1990. IFIP.
[20] P. Wadler. The essence of functional programming. In Proc. 19th Symposium on Principles of Programming Languages, pages 1-14. ACM, January 1992.
[21] C.P. Wadsworth. Semantics and Pragmatics of the Lambda Calculus. PhD thesis, Oxford University, 1971.

## A A Complete Example: Array ADT

Consider a very simple, fixed-size integer array ADT, with operations:

$$
\begin{array}{lll}
\text { newArr } & ::(\text { Int }, \ldots, \text { Int }) \rightarrow \text { Array } \\
\text { update } & ::(\text { Ix, Int }) \rightarrow \text { Array } \rightarrow \text { Array } \\
\text { select } & :: & \text { Ix } \rightarrow \text { Array } \rightarrow \text { Int }
\end{array}
$$

whose array axiomatization is given by:

$$
\begin{aligned}
& \text { select } i\left(\text { newArr }\left(x_{1}, \ldots, x_{i}, \ldots, x_{n}\right)\right) \Rightarrow x_{i} \\
& \text { update }(i, y)\left(\text { newArr }\left(x_{1}, \ldots, x_{i}, \ldots, x_{n}\right)\right) \\
& \quad \Rightarrow \operatorname{newArr}\left(x_{1}, \ldots, y, \ldots, x_{n}\right)
\end{aligned}
$$

Intuitively, newArr $x s$, where $x s$ is an $n$-tuple, creates an array of size $n$, each of whose elements is initialized to its corresponding component in $x s$; update $(i, v) a$ returns an array identical to $a$ except that the value of the $i$ th element is $v$; and select $i$ a returns the value of the $i$ th element of $a$.

It can be verified that these axioms satisfy our linearity condition. The GRS rules for this ADT are given in Figure 10. Monadically encapsulating newArr, update, and select, we get:

$$
\begin{array}{lll}
\text { type } M a & =\text { Array } \rightarrow(a, \text { Array }) \\
\text { newArrM } & ::(\text { Int }, \ldots, \text { Int }) \rightarrow M a \rightarrow a \\
\text { updateM } & ::(\text { Ix, Int }) \rightarrow M() \\
\text { selectM } & :: \text { Ix } \rightarrow M \text { Int }
\end{array}
$$

Their GRS rules are shown in Figure 11; and a programming example in this MADT is given in Figure 12. The correctness of destrutively reusing the array follows from Theorem 3.


Figure 10: GRS Rules for Array ADT


Figure 11: GRS Rules for Monadic Array MADT

Finally, we apply the axiom translation scheme to obtain the following monadic MADT axiomatization:

$$
\begin{aligned}
& \text { newArrM }\left(\ldots x_{i} \ldots\right)(\text { select } M i>g) \\
& \quad \Rightarrow \text { newArrM }\left(\ldots x_{i} \ldots\right)\left(g x_{i}\right) \\
& \text { newArrM }\left(\ldots x_{i} \ldots\right)(\text { updateM }(i, y) \gg d) \\
& \quad \Rightarrow \text { newArrM }(\ldots y \ldots) d \\
& \text { newArrM }\left(\ldots x_{i} \ldots\right)(\text { unit } b) \Rightarrow b
\end{aligned}
$$

## B The Proof of Theorem 1

The proof is a tedious proof-by-cases on the formation of both sides of a linear axiom. In some cases, an induction argument is necessary. We illustrate the main idea by going through one sample case. Consider a selector axiom $s x_{1} t=x_{2}$. To omit the uninteresting details, let us assume $x_{1}$ and $x_{2}$ are two identifiers of proper auxiliary types. By inspecting the typing rules in Figure 3 and the definition of simple ADTs, we know that $t$ can either be an identifier of type of interest, an application of some constructor (this case is proved inductively), or an application of some generator. If it is the last case, the axiom looks like

$$
s x_{1}\left(g x_{3}\right) \Rightarrow x_{2}
$$

Applying $\mathcal{S}$ to this axiom, we get

$$
g M x_{2}\left(s M x_{1} \geqq \lambda y \rightarrow n\right) \Rightarrow g M x_{2} n\left[x_{2} / y\right] .
$$

Unfolding $g M, s M$, and $\geqq$ on the left-hand side of the axiom:

$$
\begin{aligned}
& g M x_{3}\left(s M x_{1} \geqq \lambda y \rightarrow n\right) \\
& \Rightarrow_{g M} \quad \text { fst }\left(s M x_{1} \geqq \lambda y \rightarrow n\right)\left(g x_{3}\right) \\
& \Rightarrow \geq f s t\left(\left(\lambda p \rightarrow(\lambda y \rightarrow n)\left(f s t\left(s M x_{1} p\right)\right)\right.\right. \\
& \left.\left.\left(\operatorname{snd}\left(s M x_{1} p\right)\right)\right)\left(g x_{3}\right)\right) \\
& \Rightarrow_{\beta} \quad f s t\left((\lambda y \rightarrow n)\left(f s t\left(s M x_{1}\left(g x_{3}\right)\right)\right)\right. \\
& \left.\left(\operatorname{snd}\left(s M x_{1}\left(g x_{3}\right)\right)\right)\right) \\
& \Rightarrow_{\beta} \quad f s t\left(n\left[f s t\left(s M x_{1}\left(g x_{3}\right)\right) / y\right]\right. \\
& \text { (snd } \left.\left.\left(s M x_{1}\left(g x_{3}\right)\right)\right)\right) \\
& \Rightarrow_{s M} \quad f s t\left(n\left[f s t\left(s x_{1}\left(g x_{3}\right),\left(g x_{3}\right)\right) / y\right]\right. \\
& \left.\left(\operatorname{snd}\left(s M x_{1}\left(g x_{3}\right)\right)\right)\right) \\
& \Rightarrow{ }_{s} \quad \operatorname{fst}\left(n\left[x_{2} / y\right]\left(\operatorname{snd}\left(s M x_{1}\left(g x_{3}\right)\right)\right)\right) \\
& \Rightarrow_{s M} \quad \text { fst }\left(n\left[x_{2} / y\right]\left(\operatorname{snd}\left(s x_{1}\left(g x_{3}\right),\left(g x_{3}\right)\right)\right)\right) \\
& \Rightarrow_{s M} \quad f s t\left(n\left[x_{2} / y\right]\left(g x_{3}\right)\right)
\end{aligned}
$$

Unfolding $g M$ on the right-hand side:

$$
\begin{array}{rl} 
& g M x_{2} n\left[x_{2} / y\right] \\
\Rightarrow g M & f s t\left(n\left[x_{2} / y\right]\left(g x_{3}\right)\right)
\end{array}
$$

This term is identical to the unfolding of the left-hand side and we are done.

## C The Proof of Theorem 3

Following the development in [5], it can be shown that for any term $e$ in $\mathcal{L}^{\mathcal{M}}$ :

$$
e \Downarrow_{C B N} v \Longleftrightarrow \text { lift }(e) \Downarrow_{G R S} v
$$

where the monadic mutators in GRS do updates nondestructively.

So if we can prove that it is always safe to do inplace updates in GRS as an "optimization," it follows that for a $\mathcal{L}^{\mathcal{M}}$ term $e$, its CBN evaluation yields the same result as its in-place update GRS evaluation.

| newArrM ( $0, \ldots, 0$ ) |  | -- allocate new array and invoke a computation |
| :---: | :---: | :---: |
| ( updateM $(1,1)$ | $>$ | - set 1st element to 1 |
| selectM 1 | $\geq \lambda x \rightarrow$ | -- read first element and bind it to x |
| selectM 1 | $\geq \lambda y \rightarrow$ | -- read first element and bind it to y |
| unit $(x+y)$ | ) | -- return the sum of $x+y$ (the value 2 ) |

Figure 12: A Simple Program Using Array MADT

In order to prove that in-place update on a type of interest is safe, it is sufficient to show that a mutator can only update an unshared type of interest. More precisely, we want to prove that for a monadic term:

$$
g M e_{1}\left(n_{1} \geqq n_{2} \ldots n_{i} \ldots \geqq n_{k}\right)
$$

the type of interest after reduction of $\operatorname{sndA}(n t)$ at any step $i$, where $n$ is the composition of computations $n_{1}$ through $n_{i}$, is unshared.

Base case: $i=0$ : the type of interest is supplied by the rule for $g M$, which by definition is always unshared.

Induction hypothesis: the type of interest after reduction of $\operatorname{snd} A(n t)$ at step $i$ is unshared.

Induction step: prove that the type of interest is unshared after reduction of $s n d A(n t)$ at step $i+1$. We proceed by case analysis of the operation at step $i+1$ :

- unit: The rule for unit shows that it simply passes on the type of interest without adding any pointers. So by the induction hypothesis, it is still unshared.
- $m M$ : The rule for $m M$ shows that it applies $m$ to the unshared type of interest and passes the result on. But our linearity condition guarantees that $m$ handles the type of interest single-threadedly. So the type of interest is still unshared.
- $s M$ : The rule for $s M$ adds a pointer to $t$ in the first component of the pair; i.e. the reference to $t$ in $s x t$. But the reduction rule for $s n d A$ reduces this value to normal form, and since it is also linear, it cannot have any dangling reference to $t . s M$ also simply passes the type of interest on, but this is now the only reference, so the type of interest is still unshared.

Thus by induction, the type of interest is always unshared just after the reduction of $s n d A(n t)$. It follows that $m M$ is always handed an unshared copy of the type of interest, and $m M$ can safely perform an in-place update. This further implies that:

$$
e \Downarrow_{C B N} v \quad \Longleftrightarrow \quad \text { lift }(e) \Downarrow_{G R S} v
$$

where $v$ is a value and GRS employs rules doing in-place update on the type of interest.


[^0]:    To appear in the 24 th ACM Symposium on Principles of Programming Langauges, January 1997, Paris, France.

[^1]:    ${ }^{1}$ Evaluation can be generalized to include abstractions as values, but for simplicity we choose not to here.

[^2]:    ${ }^{2}$ In this work, the linearity of the type of interest is all that we care about, so it would be a waste of ink if we follow the convention of linear logic and tag each non-linear type with !. Instead we tag the linear type of interest with a ${ }_{j}$ as is done in [19]. We pronounce ${ }_{i} T$ as "squirt $T$ " (gentler than a bang).

