

Handling Local State with Global State

Koen Pauwels¹, Tom Schrijvers¹, and Shin-Cheng Mu²

¹ Department of Computer Science, KU Leuven, Belgium,
`first.last@cs.kuleuven.be`

² Institute of Information Science, Academia Sinica, Taiwan,
`scm@iis.sinica.edu.tw`

Abstract. Equational reasoning is one of the most important tools of functional programming. To facilitate its application to monadic programs, Gibbons and Hinze have proposed a simple axiomatic approach using laws that characterise the computational effects without exposing their implementation details. At the same time Plotkin and Pretnar have proposed algebraic effects and handlers, a mechanism of layered abstractions by which effects can be implemented in terms of other effects.

This paper performs a case study that connects these two strands of research. We consider two ways in which the nondeterminism and state effects can interact: the high-level semantics where every nondeterministic branch has a local copy of the state, and the low-level semantics where a single sequentially threaded state is global to all branches.

We give a monadic account of the folklore technique of handling local state in terms of global state, provide a novel axiomatic characterisation of global state and prove that the handler satisfies Gibbons and Hinze’s local state axioms by means of a novel combination of free monads and contextual equivalence. We also provide a model for global state that is necessarily non-monadic.

Keywords: monads · effect handlers · equational reasoning · nondeterminism · state · contextual equivalence

1 Introduction

Monads have been introduced to functional programming to support side effects in a rigorous, mathematically manageable manner [11,14]. Over time they have become the main framework in which effects are modelled. Various monads were developed for different effects, from general ones such as IO, state, nondeterminism, exception, continuation, environment passing, to specific purposes such as parsing. Much research was also devoted to producing practical monadic programs.

Equational reasoning about pure functional programs is particularly simple and powerful. Yet, Hutton and Fulger [7] noted that a lot less attention has been paid to reasoning about monadic programs in that style. Gibbons and Hinze [4] argue that equational reasoning about monadic programs becomes particularly

convenient and elegant when one respects the abstraction boundaries of the monad. This is possible by reasoning in terms of axioms or laws that characterise the monad’s behavior without fixing its implementation.

This paper is a case study of equational reasoning with monadic programs. Following the approach of algebraic effects and handlers [12], we consider how one monad can be implemented in terms of another—or, in other words, how one can be simulated in by another using a careful discipline. Our core contribution is a novel approach for proving the correctness of such a simulation. The proof approach is a convenient hybrid between equational reasoning based on axioms and inductive reasoning on the structure of programs. To capture the simulation we apply the algebraic effects technique of *handling* a free monad representation [15]. The latter provides a syntax tree on which to perform induction. To capture the careful discipline of the simulation we use contextual equivalence and perform inductive reasoning about program contexts. This allows us to deal with a heterogeneous axiom set where different axioms may make use of different notions of equality for programs.

We apply this proof technique to a situation where each “monad” (both the simulating monad and the simulated monad) is in fact a combination of two monads, with differing laws on how these effects interact: non-determinism and state.

In the monad we want to implement, each non-deterministic branch has its own ‘local’ copy of the state. This is a convenient effect interaction which is provided by many systems that solve search problems, including Prolog. A characterisation of this ‘local state’ monad was given by Gibbons and Hinze [4].

We realise this local state semantics in terms of a more primitive monad where a single state is sequentially threaded through the non-deterministic branches. Because this state is shared among the branches, we call this the ‘global state’ semantics. The appearance of local state is obtained by following a discipline of undoing changes to the state when backtracking to the next branch. This folklore backtracking technique is implemented by most sequential search systems because of its relative efficiency: remembering what to undo often requires less memory than creating multiple copies of the state, and undoing changes often takes less time than recomputing the state from scratch. To the best of our knowledge, our axiomatic characterisation of the global state monad is novel.

In brief, our contributions can be summarized as follows:

- We provide an axiomatic characterisation for the interaction between the monadic effects of non-determinism and state where the state is persistent (i.e., does not backtrack), together with a model that satisfies this characterisation.
- We prove that—with a careful discipline—our characterisation of persistent state can correctly simulate Gibbons and Hinze’s monadic characterisation of backtrackable state [4]. We use our novel proof approach (the core contribution of this paper) to do so.
- Our proof also comes with a mechanization in Coq.³

³ The proof can be found at <https://github.com/KoenP/LocalAsGlobal>.

The rest of the paper is structured as follows. First, Section 2 gives an overview of the main concepts used in the paper and defines our terminology. Then, Section 3 informally explores the differences between local and global state semantics. Next, Section 4 explains how to handle local state in terms of global state. Section 5 formalizes this approach and proves it correct. Finally, Sections 6 and 7 respectively discuss related work and conclude.

2 Background

This section briefly summarises the main concepts we need for equational reasoning about effects. For a more extensive treatment we refer to the work of Gibbons and Hinze [4].

2.1 Monads, Nondeterminism and State

Monads A monad consists of a type constructor $M :: * \rightarrow *$ and two operators $return :: a \rightarrow M a$ and “bind” ($\gg=$) $:: M a \rightarrow (a \rightarrow M b) \rightarrow M b$ that satisfy the following *monad laws*:

$$return\ x \gg= f = f\ x \quad , \quad (1)$$

$$m \gg= return = m \quad , \quad (2)$$

$$(m \gg= f) \gg= g = m \gg= (\lambda x \rightarrow f\ x \gg= g) \quad . \quad (3)$$

Nondeterminism The first effect we introduce is nondeterminism. Following the trail of Hutton and Fulger [7] and Gibbons and Hinze, we introduce effects based on an axiomatic characterisation rather than a specific implementation. We define a type class to capture this interface as follows:

class Monad $m \Rightarrow$ **MNonDet** m **where**

$\emptyset :: m\ a$

$(\square) :: m\ a \rightarrow m\ a \rightarrow m\ a \quad .$

In this interface, \emptyset denotes failure, while $m \square n$ denotes that the computation may yield either m or n . Precisely what laws these operators should satisfy, however, can be a tricky issue. As discussed by Kiselyov [8], it eventually comes down to what we use the monad for.

It is usually expected that (\square) and \emptyset form a monoid. That is, (\square) is associative, with \emptyset as its zero:

$$(m \square n) \square k = m \square (n \square k) \quad , \quad (4)$$

$$\emptyset \square m = m = m \square \emptyset \quad . \quad (5)$$

It is also assumed that monadic bind distributes into (\square) from the end, while \emptyset is a left zero for $\gg=$:

$$\text{left-distributivity} : (m_1 \square m_2) \gg= f = (m_1 \gg= f) \square (m_2 \gg= f) \quad , \quad (6)$$

$$\text{left-zero :} \quad \emptyset \gg f = \emptyset . \quad (7)$$

We will refer to the laws (4), (5), (6), (7) collectively as the *nondeterminism laws*.

One might intuitively expect some additional laws from a set of nondeterminism operators, such as idempotence ($p \parallel p = p$) or commutativity ($p \parallel q = q \parallel p$). However, our primary interest lies in the study of combinations of effects and – as we shall see very soon – in particular the combination of nondeterminism with *state*. One of our characterisations of this interaction would be incompatible with both idempotence and commutativity, at least if they are stated as strongly as we have done here. We will eventually introduce a weaker notion of commutativity, but it would not be instructive to do so here (as its properties would be difficult to motivate at this point).

State The state effect provides two operators:

```
class Monad m => MState s m | m -> s where
  get :: m s
  put :: s -> m () .
```

The *get* operator retrieves the state, while *put* overwrites the state by the given value. They satisfy the *state laws*:

$$\text{put-put :} \quad \text{put } st \gg \text{put } st' = \text{put } st' , \quad (8)$$

$$\text{put-get :} \quad \text{put } st \gg \text{get} = \text{put } st \gg \text{return } st , \quad (9)$$

$$\text{get-put :} \quad \text{get} \gg \text{put} = \text{return } () , \quad (10)$$

$$\text{get-get :} \quad \text{get} \gg (\lambda st \rightarrow \text{get} \gg k \ st) = \text{get} \gg (\lambda st \rightarrow k \ st \ st) , \quad (11)$$

where $m_1 \gg m_2 = m_1 \gg \lambda_ \rightarrow m_2$, which has type $(\gg) :: m \ a \rightarrow m \ b \rightarrow m \ b$.

2.2 Combining Effects

As Gibbons and Hinze already noted, an advantage of defining our effects axiomatically, rather than by providing some concrete implementation, is that it is straightforward to reason about combinations of effects. In this paper, we are interested in the interaction between nondeterminism and state, specifically.

```
class (MState s m, MNondet m) => MStateNondet s m | m -> s .
```

The type class `MStateNondet s m` simply inherits the operators of its superclasses `MState s m` and `MNondet m` without adding new operators, and implementations of this class should comply with all laws of both superclasses.

However, this is not the entire story. Without additional ‘interaction laws’, the design space for implementations of this type class is left wide-open with respect to questions about how these effects interact. In particular, it seems hard to imagine that one could write nontrivial programs which are agnostic towards the question of what happens to the state of the program when the program backtracks. We discuss two possible approaches.

Local State Semantics One is what Gibbons and Hinze call “backtrackable state”, that is, when a branch of the nondeterministic computation runs into a dead end and the continuation of the computation is picked up at the most recent branching point, any alterations made to the state by our terminated branch are invisible to the continuation. Because in this scheme state is local to a branch, we will refer to these semantics as *local state semantics*. We characterise local state semantics with the following laws:

$$\text{right-zero :} \quad m \gg \emptyset = \emptyset \quad , \quad (12)$$

$$\text{right-distributivity :} \quad m \gg (\lambda x \rightarrow f_1 x \parallel f_2 x) = (m \gg f_1) \parallel (m \gg f_2). \quad (13)$$

With some implementations of the monad, it is likely that in the lefthand side of (13), the effect of m happens once, while in the righthand side it happens twice. In (12), the m on the lefthand side may incur some effects that do not happen in the righthand side.

Having (12) and (13) leads to profound consequences on the semantics and implementation of monadic programs. To begin with, (13) implies that for (\parallel) we have some limited notion of commutativity. For instance, both the left and right distributivity rules can be applied to the term $(\text{return } x \parallel \text{return } y) \gg \lambda z \rightarrow \text{return } z \parallel \text{return } z$. It is then easy to show that this term must be equal to both $\text{return } x \parallel \text{return } x \parallel \text{return } y \parallel \text{return } y$ and $\text{return } x \parallel \text{return } y \parallel \text{return } x \parallel \text{return } y$.⁴

In fact, having (12) and (13) gives us very strong and useful commutative properties. To be clear what we mean, we give a formal definition:

Definition 1. *Let m and n be two monadic programs such that x does not occur free in m , and y does not occur free in n . We say m and n commute if*

$$\begin{aligned} m \gg \lambda x \rightarrow n \gg \lambda y \rightarrow f x y &= \\ n \gg \lambda y \rightarrow m \gg \lambda x \rightarrow f x y &. \end{aligned} \quad (14)$$

We say that effects ϵ and δ commute if any m and n commute as long as their only effects are respectively ϵ and δ .

One important result is that, in local state semantics, non-determinism commutes with any effect :

Theorem 1. *If right-zero (12) and right-distributivity (13) hold in addition to the other laws, non-determinism commutes with any effect.*

Implementation-wise, (12) and (13) imply that each nondeterministic branch has its own copy of the state. To see that, let $m = \text{put } 1$, $f_1 () = \text{put } 2$, and $f_2 () = \text{get}$ in (13) — the state we *get* in the second branch does not change, despite the *put* 2 in the first branch. One implementation satisfying the laws is

⁴ Gibbons and Hinze [4] were mistaken in their claim that the type $s \rightarrow [(a, s)]$ constitutes a model of their backtrackable state laws; it is not a model because its (\parallel) does not commute with itself. One could consider a relaxed semantics that admits $s \rightarrow [(a, s)]$, but that is not the focus of this paper.

$\mathbf{M} s a = s \rightarrow \mathbf{Bag} (a, s)$, where $\mathbf{Bag} a$ is an implementation of a multiset or “bag” data structure. If we ignore the unordered nature of the \mathbf{Bag} type, this implementation is similar to $\mathbf{StateT} s$ ($\mathbf{ListT Identity}$) in the Monad Transformer Library [5]. With effect handling [15,9], the monad behaves similarly (except for the limited commutativity implied by law (13)) if we run the handler for state before that for list.

Global State Semantics Alternatively, we can choose a semantics where state reigns over nondeterminism. In this case of non-backtrackable state, alterations to the state persist over backtracks. Because only a single state is shared over all the branches of the nondeterministic computation, we call this semantics *global state semantics*. We will return later to the question of how to define laws that capture our intuition for this kind of semantics, because (to the best of our knowledge) this constitutes a novel contribution.

Even just figuring out an implementation of a global state monad that matches our intuition is already tricky. One might believe that $\mathbf{M} s a = s \rightarrow ([a], s)$ is a natural implementation of such a monad. The usual, naive implementation of (\gg) using this representation, however, does not satisfy left-distributivity (6), violates monad laws, and is therefore not even a monad. The type $\mathbf{ListT} (\mathbf{State} s)$ generated using the Monad Transformer Library [5] expands to essentially the same implementation, and is flawed in the same way. More careful implementations of \mathbf{ListT} , which do satisfy (6) and the monad laws, have been proposed [3,13]. Effect handlers (e.g. Wu [15] and Kiselyov and Ishii [9]) do produce implementations which match our intuition of a non-backtracking computation if we run the handler for non-determinism before that of state.

We provide a direct implementation to aid the intuition of the reader. Essentially the same implementation is obtained by using the type $\mathbf{ListT} (\mathbf{State} s)$ where \mathbf{ListT} is implemented as a correct monad transformer. This implementation has a non-commutative (\parallel).

$$\mathbf{M} s a = s \rightarrow (\mathbf{Maybe} (a, \mathbf{M} s a), s) \text{ .}$$

The \mathbf{Maybe} in this type indicates that a computation might fail to produce a result. But note that the s is outside of the \mathbf{Maybe} : even if the computation fails to produce any result, a modified state may be returned (this is different from local state semantics). \emptyset , of course, returns an empty continuation ($\mathbf{Nothing}$) and an unmodified state. (\parallel) first exhausts the left branch (always collecting any state modifications it performs), before switching to the right branch.

$$\begin{aligned} \emptyset &= \lambda s \rightarrow (\mathbf{Nothing}, s) \text{ ,} \\ p \parallel q &= \lambda s \rightarrow \mathbf{case} \ p \ s \ \mathbf{of} \ (\mathbf{Nothing}, t) \rightarrow q \ t \\ &\quad (\mathbf{Just} (x, r), t) \rightarrow (\mathbf{Just} (x, r \parallel q), t) \text{ .} \end{aligned}$$

The state operators are implemented in a straightforward manner.

$$\begin{aligned} \mathit{get} &= \lambda s \rightarrow (\mathbf{Just} (s, \emptyset), s) \text{ ,} \\ \mathit{put} \ s &= \lambda t \rightarrow (\mathbf{Just} ((), \emptyset), s) \text{ .} \end{aligned}$$

And this implementation is also a monad. The implementation of $p \gg k$ extends every branch within p with k , threading the state through this entire process.

$$\begin{aligned} \text{return } x = \lambda s &\rightarrow (\text{Just } (x, \emptyset), s) \text{ ,} \\ p \gg k &= \lambda s \rightarrow \text{case } p \text{ s of } (\text{Nothing}, t) \rightarrow (\text{Nothing}, t) \\ &\quad (\text{Just } (x, q), t) \rightarrow (k \ x \ \parallel \ (q \gg k)) \ t \text{ .} \end{aligned}$$

3 Motivation

In the previous section we discussed two possible semantics for the interaction of state and nondeterminism: global and local state semantics. In this section, we will further explore the differences between these two interpretations. Using the classic n -queens puzzle as an example, we show that sometimes we end up in a situation where we want to write our program according to local state semantics (which is generally speaking easier to reason about), but desire the space usage characteristics of global state semantics.

3.1 Example: The n -Queens Problem




Fig. 1: (a) This placement can be represented by $[3, 5, 7, 1, 6, 0, 2, 4]$. (b) Up diagonals. (c) Down diagonals.

The n -queens puzzle presented in this section is adapted and simplified from that of Gibbons and Hinze [4]. The aim of the puzzle is to place n queens on a n by n chess board such that no two queens can attack each other. Given n , we number the rows and columns by $[0..n-1]$. Since all queens should be placed on distinct rows and distinct columns, a potential solution can be represented by a permutation xs of the list $[0..n-1]$, such that $xs !! i = j$ denotes that the queen on the i th column is placed on the j th row (see Figure 1(a)). The specification can be written as a non-deterministic program:

$$\begin{aligned} \text{queens} &:: \text{MNondet } m \Rightarrow \text{Int} \rightarrow m \text{ [Int]} \\ \text{queens } n &= \text{perm } [0..n-1] \gg \text{filt safe} \text{ ,} \end{aligned}$$

where *perm* non-deterministically computes a permutation of its input, and the pure function *safe* :: [Int] → Bool, to be defined later, determines whether a solution is valid. The function *filt p x* returns *x* if *p x* holds, and fails otherwise. It can be defined in terms of a standard monadic function *guard*:

```

filt :: MNonDet m ⇒ (a → Bool) → a → m a
filt p x = guard (p x) >> return x ,
guard :: MNonDet m ⇒ Bool → m ()
guard b = if b then return () else ∅ .

```

The function *perm* can be written either as a fold or an unfold. For this problem we choose the latter, using a function *select*, which non-deterministically splits a list into a pair containing one chosen element and the rest. For example, *select [1, 2, 3]* yields one of (1, [2, 3]), (2, [1, 3]) and (3, [1, 2]).

```

select :: MNonDet m ⇒ [a] → m (a, [a])
select [] = ∅
select (x : xs) = return (x, xs) [] ((id × (x:)) $) select xs ,
perm :: MNonDet m ⇒ [a] → m [a]
perm [] = return []
perm xs = select xs >>= λ(x, ys) → (x:) $) perm ys ,

```

where *f \$ m* = *m* >> (*return* · *f*) which applies a pure function to a monadic value, and (*f* × *g*) (*x, y*) = (*f x, g y*).

This specification of *queens* generates all the permutations, before checking them one by one, in two separate phases. We wish to fuse the two phases, which allows branches generates a non-safe placement to be pruned earlier, and thus produce a faster implementation.

A Backtracking Algorithm In our representation, queens cannot be put on the same row or column. Therefore, *safe* only needs to make sure that no two queens are put on the same diagonal. An 8 by 8 chess board has 15 *up diagonals* (those running between bottom-left and top-right). Let them be indexed by [0..14] (see Figure 1(b)). Similarly, there are 15 *down diagonals* (running between top-left and bottom right, indexed by [-7..7] in Figure 1(c)). We can show, by routine program calculation, that whether a placement is safe can be checked in one left-to-right traversal — define *safe xs* = *safeAcc* (0, [], []) *xs*, where

```

safeAcc :: (Int, [Int], [Int]) → [Int] → Bool
safeAcc (i, us, ds) [] = True
safeAcc (i, us, ds) (x : xs) = ok (i', us', ds') ∧ safeAcc (i', us', ds') xs ,
  where (i', us', ds') = (i + 1, (i + x : us), (i - x : ds)) ,
ok (i, (x : us), (y : ds)) = x ∉ us ∧ y ∉ ds .

```

Operationally, (*i, us, ds*) is a “state” kept by *safeAcc*, where *i* is the current column, while *us* and *ds* are respectively the up and down diagonals encountered

so far. Function *safeAcc* behaves like a fold-left. Indeed, it can be defined using *scanl* and *all* (where $all\ p = foldr\ (\wedge)\ True \cdot map\ p$):

$$safeAcc\ (i, us, ds) = all\ ok \cdot tail \cdot scanl\ (\oplus)\ (i, us, ds) \ , \\ \text{where } (i, us, ds) \oplus x = (i + 1, (i + x : us), (i - x : ds)) \ .$$

One might wonder whether the “state” can be implemented using an actual state monad. Indeed, the following is among the theorems we have proved:

Theorem 2. *If state and non-determinism commute, we have that for all xs , st , (\oplus) , and ok ,*

$$filt\ (all\ ok \cdot tail \cdot scanl\ (\oplus)\ st)\ xs = \\ protect\ (put\ st \gg foldr\ (\odot)\ (return\ [])\ xs) \ , \\ \text{where } x \odot m = get \gg \lambda st \rightarrow guard\ (ok\ (st \oplus x)) \gg \\ put\ (st \oplus x) \gg ((x:) \$ m) \ .$$

The function $protect\ m = get \gg \lambda ini \rightarrow m \gg \lambda x \rightarrow put\ ini \gg return\ x$ saves the initial state and restores it after the computation. As for (\odot) , it assumes that the “state” passed around by *scanl* is stored in a monadic state, checks whether the new state $st \oplus x$ satisfies *ok*, and updates the state with the new value.

For Theorem 2 to hold, however, we need state and non-determinism to commute. It is so in the local state semantics, which can be proved using the non-determinism laws, (12), and (13).

Now that the safety check can be performed in a *foldr*, recalling that *perm* is an unfold, it is natural to try to fuse them into one. Indeed, it can be proved that, with (\oplus) , *ok*, and (\odot) defined above, we have $perm\ xs \gg foldr\ (\odot)\ (return\ []) = qBody\ xs$, where

$$qBody :: MStateNondet\ (Int, [Int], [Int])\ m \Rightarrow [Int] \rightarrow m\ [Int] \\ qBody\ [] = return\ [] \\ qBody\ xs = select\ xs \gg \lambda (x, ys) \rightarrow \\ get \gg \lambda st \rightarrow guard\ (ok\ (st \oplus x)) \gg \\ put\ (st \oplus x) \gg ((x:) \$ qBody\ ys) \ .$$

The proof also heavily relies on the commutativity between non-determinism and state.

To wrap up, having fused *perm* and safety checking into one phase, we may compute *queens* by:

$$queens :: MStateNondet\ (Int, [Int], [Int])\ m \Rightarrow Int \rightarrow m\ [Int] \\ queens\ n = protect\ (put\ (0, [], []) \gg qBody\ [0..n-1]) \ .$$

This is a backtracking algorithm that attempts to place queens column-by-column, proceeds to the next column if *ok* holds, and backtracks otherwise. The derivation from the specification to this program relies on a number of properties that hold in the local state semantics.

3.2 Transforming between Local State and Global State

For a monad with both non-determinism and state, the local state laws imply that each non-deterministic branch has its own state. This is not costly for states consisting of linked data structures, for example the state $(\text{Int}, [\text{Int}], [\text{Int}])$ in the n -queens problem. In some applications, however, the state might be represented by data structures, e.g. arrays, that are costly to duplicate: When each new state is only slightly different from the previous (say, the array is updated in one place each time), we have a wasteful duplication of information. Although this is not expected to be an issue for realistic sizes of the n -queens problem due to the relatively small state, one can imagine that for some problems where the state is very large, this can be a problem.

Global state semantics, on the other hand, has a more “low-level” feel to it. Because a single state is threaded through the entire computation without making any implicit copies, it is easier to reason about resource usage in this setting. So we might write our programs directly in the global state style. However, if we do this to a program that would be more naturally expressed in the local state style (such as our n -queens example), this will come at a great loss of clarity. Furthermore, as we shall see, although it is easier to reason about resource usage of programs in the global state setting, it is significantly more difficult to reason about their semantics. We could also write our program first in a local state style and then translate it to global state style. Doing this manually is a tedious and error-prone process that leaves us with code that is hard to maintain. A more attractive proposition is to design a systematic program transformation that takes a program written for local state semantics as input, and outputs a program that, when interpreted under global state semantics, behaves exactly the same as the original program interpreted under local state semantics.

In the remainder of the paper we define this program transformation and prove it correct. We believe that, in particular, the proof *technique* is of interest.

4 Non-Determinism with Global State

So far, we have evaded giving a precise axiomatic characterisation of global state semantics: although in Section 2 we provided an example implementation that matches our intuition of global state semantics, we haven’t provided a clear formulation of that intuition. We begin this section by finally stating the “global state law”, which characterises exactly the property that sets apart non-backtrackable state from backtrackable state.

In the rest of the section, we appeal to intuition and see what happens when we work with a global state monad: what pitfalls we may encounter, and what programming pattern we may use, to motivate the more formal treatment in Section 5.

4.1 The Global State Law

We have already discussed general laws for nondeterministic monads (laws (4) through (7)), as well as laws which govern the interaction between state and

nondeterminism in a local state setting (laws (13) and (12)). For global state semantics, an alternative law is required to govern the interactions between non-determinism and state. We call this the *global state law*.

$$\mathbf{put-or} : (put\ s \gg m) \parallel n = put\ s \gg (m \parallel n) \quad , \quad (15)$$

This law allows the lifting of a *put* operation from the left branch of a nondeterministic choice, an operation which does not preserve meaning under local state semantics. Suppose for example that $m = \emptyset$, then by (12) and (5), the left-hand side of the equation is equal to n , whereas by (5), the right-hand side of the equation is equal to $put\ s \gg n$.

By itself, this law leaves us free to choose from a large space of implementations with different properties. For example, in any given implementation, the programs $return\ x \parallel return\ y$ and $return\ y \parallel return\ x$ may be considered semantically identical, or they may be considered semantically distinct. The same goes for the programs $return\ x \parallel return\ x$ and $return\ x$, or the programs $(put\ s \gg return\ x) \parallel m$ and $(put\ s \gg return\ x) \parallel (put\ s \gg m)$. Additional axioms will be introduced as needed to cover these properties in Section 5.2.

4.2 Chaining Using Non-deterministic Choice

In backtracking algorithms that keep a global state, it is a common pattern to 1. update the current state to its next step, 2. recursively search for solutions, and 3. roll back the state to the previous step (regardless of whether a solution is found). To implement such pattern as a monadic program, one might come up with something like the code below:

$$modify\ next \gg search \gg= modReturn\ prev \quad ,$$

where *next* advances the state, *prev* undoes the modification of *next* ($prev \cdot next = id$), and *modify* and *modReturn* are defined by:

$$\begin{aligned} modify\ f &= get \gg= (put \cdot f) \quad , \\ modReturn\ f\ v &= modify\ f \gg return\ v \quad . \end{aligned}$$

Let the initial state be *st* and assume that *search* found three choices $m_1 \parallel m_2 \parallel m_3$. We wish that m_1 , m_2 , and m_3 all start running with state *next st*, and the state is restored to *prev (next st) = st* afterwards. Due to (6), however, it expands to

$$\begin{aligned} modify\ next \gg (m_1 \parallel m_2 \parallel m_3) \gg= modReturn\ prev = \\ modify\ next \gg ((m_1 \gg= modReturn\ prev) \parallel \\ (m_2 \gg= modReturn\ prev) \parallel \\ (m_3 \gg= modReturn\ prev)) \quad . \end{aligned}$$

With a global state, it means that m_2 starts with state *st*, after which the state is rolled back further to *prev st*. The computation m_3 starts with *prev st*, after which the state is rolled too far to *prev (prev st)*. In fact, one cannot guarantee

that *modReturn prev* is always executed — if *search* fails and reduces to \emptyset , *modReturn prev* is not run at all, due to (7).

We need a way to say that “*modify next* and *modReturn prev* are run exactly once, respectively before and after all non-deterministic branches in *solve*.” Fortunately, we have discovered a curious technique. Define

$$\begin{aligned} \textit{side} &:: \text{MNondet } m \Rightarrow m \ a \rightarrow m \ b \\ \textit{side } m &= m \gg \emptyset \ . \end{aligned}$$

Since non-deterministic branches are executed sequentially, the program

$$\textit{side } (\textit{modify next}) \parallel m_1 \parallel m_2 \parallel m_3 \parallel \textit{side } (\textit{modify prev})$$

executes *modify next* and *modify prev* once, respectively before and after all the non-deterministic branches, even if they fail. Note that *side m* does not generate a result. Its presence is merely for the side-effect of *m*, hence the name.

The reader might wonder: now that we are using (\parallel) as a sequencing operator, does it simply coincide with (\gg)? Recall that we still have left-distributivity (6) and, therefore, $(m_1 \parallel m_2) \gg n$ equals $(m_1 \gg n) \parallel (m_2 \gg n)$. That is, (\parallel) acts as “insertion points”, where future code followed by (\gg) can be inserted into! This is certainly a dangerous feature, whose undisciplined use can lead to chaos. However, we will exploit this feature and develop a safer programming pattern in the next section.

4.3 State-Restoring Operations

The discussion above suggests that one can implement backtracking, in a global-state setting, by using (\parallel) and *side* appropriately. We can even go a bit further by defining the following variation of *put*, which restores the original state when it is backtracked over:


$$\begin{aligned} \textit{put}_R &:: \text{MStateNondet } s \ m \Rightarrow s \rightarrow m \ () \\ \textit{put}_R \ s &= \textit{get} \gg \lambda s_0 \rightarrow \textit{put } s \parallel \textit{side } (\textit{put } s_0) \ . \end{aligned}$$


Fig. 2: Illustration of state-restoring put

To help build understanding for put_r , Figure 2 shows the flow of execution for the expression $(put_r t \gg ret x) \parallel ret y$. Initially, the state is s ; it gets modified to t at the $put t$ node after which the value x is output with the working state t . Then, because we found a result, we backtrack (since we're using global-state semantics, the state modification caused by $put t$ is not reversed), arriving in the *side* operation branch. The $put s$ operation is executed, which resets the state to s , and then the branch immediately fails, so we backtrack to the right branch of the topmost (\parallel) . There the value y is output with working state s .

For some further intuition about put_r , consider $put_r s \gg comp$ where $comp$ is some arbitrary computation:

$$\begin{aligned}
& put_r s \gg comp \\
= & (get \gg \lambda s_0 \rightarrow put s \parallel side (put s_0)) \gg comp \\
= & \{ \text{monad law, left-distributivity (6)} \} \\
& get \gg \lambda s_0 \rightarrow (put s \gg comp) \parallel (side (put s_0) \gg comp) \\
= & \{ \text{by (7) } \emptyset \gg comp = \emptyset, \text{ monad laws} \} \\
& get \gg \lambda s_0 \rightarrow (put s \gg comp) \parallel side (put s_0) .
\end{aligned}$$

Thanks to left-distributivity (6), $(\gg comp)$ is promoted into (\parallel) . Furthermore, the $(\gg comp)$ after $side (put s_0)$ is discarded by (7). In words, $put_r s \gg comp$ saves the current state, computes $comp$ using state s , and restores the saved state! The subscript R stands for “restore.” Note also that $(put_r s \gg m_1) \gg m_2 = put_r s \gg (m_1 \gg m_2)$ — the state restoration happens in the end.

The behaviour of put_r is rather tricky. It is instructive comparing

- (a) *return x*,
- (b) $put s \gg return x$,
- (c) $put_r s \gg return x$.

When run in initial state s_0 , they all yield x as the result. The final states after running (a), (b) and (c) are s_0 , s and s_0 , respectively. However, (c) does *not* behave identically to (a) in all contexts! For example, in the context $(\gg get)$, we can tell them apart: $return x \gg get$ returns s_0 , while $put_r s \gg return x \gg get$ returns s , even though the program yields final state s_0 .

We wish that put_r , when run with a global state, satisfies laws (8) through (13) — the state laws and the *local* state laws. If so, one could take a program written for a local state monad, replace all occurrences of put by put_r , and run the program with a global state. Unfortunately this is not the case: put_r does satisfy **put-put** (8) and **put-get** (10), but **get-put** (9) fails — $get \gg put_r$ and $return ()$ can be differentiated by some contexts, for example $(\gg put t)$. To see that, we calculate:

$$\begin{aligned}
& (get \gg put_r) \gg put t \\
= & (get \gg \lambda s \rightarrow get \gg \lambda s_0 \rightarrow put s \parallel side (put s_0)) \gg put t \\
= & \{ \text{get-get} \} \\
& (get \gg \lambda s \rightarrow put s \parallel side (put s)) \gg put t \\
= & \{ \text{monad laws, left-distributivity} \}
\end{aligned}$$

$$\begin{aligned}
& get \gg \lambda s \rightarrow (put\ s \gg put\ t) \parallel side\ (put\ s) \\
= & \{ \mathbf{put-put} \} \\
& get \gg \lambda s \rightarrow put\ t \parallel side\ (put\ s) .
\end{aligned}$$

Meanwhile, $return\ () \gg put\ t = put\ t$, which does not behave in the same way as $get \gg \lambda s \rightarrow put\ t \parallel side\ (put\ s)$ when $s \neq t$.

In a global-state setting, the left-distributivity law (6) makes it tricky to reason about combinations of (\parallel) and (\gg) operators. Suppose we have a program $(m \parallel n)$, and we construct an extended program by binding a continuation f to it such that we get $(m \parallel n) \gg f$ (where f might modify the state). Under global-state semantics, the evaluation of the right branch is influenced by the state modifications performed by evaluating the left branch. So by (6), this means that when we get to evaluating the n subprogram in the extended program, it will do so with a different initial state (the one obtained after running $m \gg f$) compared to the initial state in the original program (the one obtained by running m). In other words, placing our program in a different context changed the meaning of one of its subprograms. So it is difficult to reason about programs compositionally in this setting—some properties hold only when we take the entire program into consideration.

It turns out that all properties we need do hold, provided that *all* occurrences of put are replaced by put_r —problematic contexts such as $put\ t$ above are thus ruled out. However, that “all put are replaced by put_r ” is a global property, and to properly talk about it we have to formally define contexts, which is what we will do in Section 5. Notice though, that simulation of local state semantics by judicious use of put_r does not avoid the unnecessary copying mentioned in Section 3.2, it merely makes it explicit in the program. We will address this shortcoming in Section 5.6.

5 Laws and Translation for Global State Monad

In this section we give a more formal treatment of the non-deterministic global state monad. Not every implementation of the global state law allows us to accurately simulate local state semantics though, so we propose additional laws that the implementation must respect. These laws turn out to be rather intricate. To make sure that there exists a model, an implementation is proposed, and it is verified in Coq that the laws and some additional theorems are satisfied.

The ultimate goal, however, is to show the following property: given a program written for a local-state monad, if we replace all occurrences of put by put_r , the resulting program yields the same result when run with a global-state monad. This allows us to painlessly port our previous algorithm to work with a global state. To show this we first introduce a syntax for nondeterministic and stateful monadic programs and contexts. Then we imbue these programs with global-state semantics. Finally we define the function that performs the translation just described, and prove that this translation is correct.

5.1 Programs and Contexts

<pre> data Prog a where Return :: a → Prog a ∅ :: Prog a (∥) :: Prog a → Prog a → Prog a Get :: (S → Prog a) → Prog a Put :: S → Prog a → Prog a </pre> <p style="text-align: center;">(a)</p>	<pre> run :: Prog a → Dom a ret̄ :: a → Dom a ∅̄ :: Dom a (∥̄) :: Dom a → Dom a → Dom a get̄ :: (S → Dom a) → Dom a put̄ :: S → Dom a → Dom a </pre> <p style="text-align: center;">(b)</p>
---	---

Fig. 3: (a) Syntax for programs. (b) Semantic domain.

In the previous sections we have been mixing syntax and semantics, which we avoid in this section by defining the program syntax as a free monad. This way we avoid the need for a type-level distinction between programs with local-state semantics and programs with global-state semantics. Figure 3(a) defines a syntax for nondeterministic, stateful, closed programs **Prog**, where the **Get** and **Put** constructors take continuations as arguments, and the ($\gg\equiv$) operator is defined as follows:

$$\begin{aligned}
(\gg\equiv) &:: \text{Prog } a \rightarrow (a \rightarrow \text{Prog } b) \rightarrow \text{Prog } b \\
\text{Return } x \gg\equiv f &= f \ x \\
\emptyset \gg\equiv f &= \emptyset \\
(m \parallel n) \gg\equiv f &= (m \gg\equiv f) \parallel (n \gg\equiv f) \\
\text{Get } k \gg\equiv f &= \text{Get } (\lambda s \rightarrow k \ s \gg\equiv f) \\
\text{Put } s \ m \gg\equiv f &= \text{Put } s \ (m \gg\equiv k) \ .
\end{aligned}$$

One can see that ($\gg\equiv$) is defined as a purely syntactical manipulation, and its definition has laws (6) and (7) built-in.

The meaning of such a monadic program is determined by a semantic domain of our choosing, which we denote with **Dom**, and its corresponding domain operators \overline{ret} , $\overline{\emptyset}$, \overline{get} , \overline{put} and $\overline{\parallel}$ (see figure 3(b)). The $run :: \text{Prog } a \rightarrow \text{Dom } a$ function “runs” a program **Prog** a into a value in the semantic domain **Dom** a :

$$\begin{aligned}
run \ (\text{Return } x) &= \overline{ret} \ x \\
run \ \emptyset &= \overline{\emptyset} \\
run \ (m_1 \parallel m_2) &= run \ m_1 \ \overline{\parallel} \ run \ m_2 \\
run \ (\text{Get } k) &= \overline{get} \ (\lambda s \rightarrow run \ (k \ s)) \\
run \ (\text{Put } s \ m) &= \overline{put} \ s \ (run \ m) \ .
\end{aligned}$$

Note that no $\gg\equiv$ operator is required to define run ; in other words, **Dom** need not be a monad. In fact, as we will see later, we will choose our implementation in such a way that there does not exist a bind operator for run .

5.2 Laws for Global State Semantics

We impose the laws upon Dom and the domain operators to ensure the semantics of a non-backtracking (global-state), nondeterministic, stateful computation for our programs. Naturally, we need laws analogous to the state laws and nondeterminism laws to hold for our semantic domain. As it is not required that a bind operator (\gg) $:: \text{Dom } a \rightarrow (a \rightarrow \text{Dom } b) \rightarrow \text{Dom } b$ be defined for the semantic domain (and we will later argue that it *cannot* be defined for the domain, given the laws we impose on it), the state laws ((8) through (11)) must be reformulated to fit the continuation-passing style of the semantic domain operators.

$$\overline{\text{put}} s (\overline{\text{put}} t p) = \overline{\text{put}} t p \quad , \quad (16)$$

$$\overline{\text{put}} s (\overline{\text{get}} k) = \overline{\text{put}} s (k s) \quad , \quad (17)$$

$$\overline{\text{get}} (\lambda s \rightarrow \overline{\text{put}} s m) = m \quad , \quad (18)$$

$$\overline{\text{get}} (\lambda s \rightarrow \overline{\text{get}} (\lambda t \rightarrow k s t)) = \overline{\text{get}} (\lambda s \rightarrow k s s) \quad . \quad (19)$$

Two of the nondeterminism laws—(6) and (7)—also mention the bind operator. As we have seen earlier, they are trivially implied by the definition of (\gg) for Prog . Therefore, we need not impose equivalent laws for the semantic domain (and in fact, we cannot formulate them given the representation we have chosen). Only the two remaining nondeterminism laws—(4) and (5)—need to be stated:

$$(m \bar{\parallel} n) \bar{\parallel} p = m \bar{\parallel} (n \bar{\parallel} p) \quad , \quad (20)$$

$$\bar{\emptyset} \bar{\parallel} m = m \bar{\parallel} \bar{\emptyset} = m \quad . \quad (21)$$

We also reformulate the global-state law (15):

$$\overline{\text{put}} s p \bar{\parallel} q = \overline{\text{put}} s (p \bar{\parallel} q) \quad . \quad (22)$$

It turns out that, apart from the **put-or** law, our proofs require certain additional properties regarding commutativity and distributivity which we introduce here:

$$\begin{aligned} \overline{\text{get}} (\lambda s \rightarrow \overline{\text{put}} (t s) p \bar{\parallel} \overline{\text{put}} (u s) q \bar{\parallel} \overline{\text{put}} s \bar{\emptyset}) = \\ \overline{\text{get}} (\lambda s \rightarrow \overline{\text{put}} (u s) q \bar{\parallel} \overline{\text{put}} (t s) p \bar{\parallel} \overline{\text{put}} s \bar{\emptyset}) \quad , \end{aligned} \quad (23)$$

$$\overline{\text{put}} s (\overline{\text{ret}} x \bar{\parallel} p) = \overline{\text{put}} s (\overline{\text{ret}} x) \bar{\parallel} \overline{\text{put}} s p \quad . \quad (24)$$

These laws are not considered general “global state” laws, because it is possible to define reasonable implementations of global state semantics that violate these laws, and because they are not exclusive to global state semantics.

The ($\bar{\parallel}$) operator is not, in general, commutative in a global state setting. However, we will require that the order in which results are computed does not matter. This might seem contradictory at first glance. To be more precise, we do not require the property $p \bar{\parallel} q = q \bar{\parallel} p$ because the subprograms p and q might perform non-commuting edits to the global state. But we do expect that programs without side-effects commute freely; for instance, we expect $\overline{\text{return}} x \bar{\parallel} \overline{\text{return}} y = \overline{\text{return}} y \bar{\parallel} \overline{\text{return}} x$. In other words, collecting all the

results of a nondeterministic computation is done with a set-based semantics in mind rather than a list-based semantics, but this does not imply that the order of state effects does not matter.

In fact, the notion of commutativity we wish to impose is still somewhat stronger than just the fact that results commute: we want the $(\overline{\parallel})$ operator to commute with respect to any pair of subprograms whose modifications of the state are ignored—that is, immediately overwritten—by the surrounding program. This property is expressed by law (23). An example of an implementation which does not respect this law is one that records the history of state changes.

In global-state semantics, \overline{put} operations cannot, in general, distribute over $(\overline{\parallel})$. However, an implementation may permit distributivity if certain conditions are met. Law (24) states that a \overline{put} operation distributes over a nondeterministic choice if the left branch of that choice simply returns a value. This law has a particularly striking implication: it disqualifies any implementation for which a bind operator $(\overline{\gg})::\text{Dom } a \rightarrow (a \rightarrow \text{Dom } b) \rightarrow \text{Dom } b$ can be defined! Consider for instance the following program:

$$\overline{put } x (\overline{ret } w \overline{\parallel} \overline{get } \overline{ret }) \overline{\gg} \lambda z \rightarrow \overline{put } y (\overline{ret } z) .$$

If (24) holds, this program should be equal to

$$(\overline{put } x (\overline{ret } w) \overline{\parallel} \overline{put } x (\overline{get } \overline{ret})) \overline{\gg} \lambda z \rightarrow \overline{put } y (\overline{ret } z) .$$

However, it is proved in Figure 4 that the first program can be reduced to $\overline{put } y (\overline{ret } w \overline{\parallel} \overline{ret } y)$, whereas the second program is equal to $\overline{put } y (\overline{ret } w \overline{\parallel} \overline{ret } x)$, which clearly does not always have the same result.

To gain some intuition about why law (24) prohibits a bind operator, consider that the presence or absence of a bind operator influences what equality of programs means. Our first intuition might be that we consider two programs equal if they produce the same outputs given the same inputs. But this is too narrow a view: for two programs to be considered equal, they must also behave the same under *composition*; that is, we must be able to replace one for the other within a larger whole, without changing the meaning of the whole. The bind operator allows us to compose programs *sequentially*, and therefore its existence implies that, for two programs to be considered equal, they must also behave identically under sequential composition. Under local-state semantics, this additional requirement coincides with other notions of equality: we can't come up with a pair of programs which both produce the same outputs given the same inputs, but behave differently under sequential composition. But under global-state semantics, we can come up with such counterexamples: consider the subprograms of our previous example $\overline{put } x (\overline{ret } w \overline{\parallel} \overline{get } \overline{ret})$ and $(\overline{put } x (\overline{ret } w) \overline{\parallel} \overline{put } x (\overline{get } \overline{ret}))$. Clearly we expect these two programs to produce the exact same results in isolation, yet when they are sequentially composed with the program $\lambda z \rightarrow \overline{put } y (\overline{ret } z)$, their different nature is revealed (by (6)).

It is worth remarking that introducing either one of these additional laws disqualify the example implementation given in Section 2.2 (even if it is adapted

$$\begin{array}{ll}
\overline{put} \ x \ (\overline{ret} \ w \ \overline{\parallel} \ \overline{get} \ \overline{ret}) \ \overline{\gg} \ \lambda z \rightarrow \overline{put} \ y \ (\overline{ret} \ z) & (\overline{put} \ x \ (\overline{ret} \ w) \ \overline{\parallel} \ \overline{put} \ x \ (\overline{get} \ \overline{ret})) \\
= \{ \text{definition of } (\overline{\gg}) \} & \overline{\gg} \ \lambda z \rightarrow \overline{put} \ y \ (\overline{ret} \ z) \\
\overline{put} \ x \ (\overline{put} \ y \ (\overline{ret} \ w) \ \overline{\parallel} \ \overline{get} \ (\lambda s \rightarrow \overline{put} \ y \ (\overline{ret} \ s))) & = \{ \text{definition of } (\overline{\gg}) \} \\
= \{ \text{by (22) and (16)} \} & \overline{put} \ x \ (\overline{put} \ y \ (\overline{ret} \ w)) \\
\overline{put} \ y \ (\overline{ret} \ w \ \overline{\parallel} \ \overline{get} \ (\lambda s \rightarrow \overline{put} \ y \ (\overline{ret} \ s))) & \overline{\parallel} \ \overline{put} \ x \ (\overline{get} \ (\lambda s \rightarrow \overline{put} \ y \ (\overline{ret} \ s))) \\
= \{ \text{by (24)} \} & = \{ \text{by (16) and (17)} \} \\
\overline{put} \ y \ (\overline{ret} \ w) \ \overline{\parallel} \ \overline{put} \ y \ (\overline{get} \ (\lambda s \rightarrow \overline{put} \ y \ (\overline{ret} \ s))) & \overline{put} \ y \ (\overline{ret} \ w) \ \overline{\parallel} \ \overline{put} \ x \ (\overline{put} \ y \ (\overline{ret} \ x)) \\
= \{ \text{by (17) and (16)} \} & = \{ \text{by (16)} \} \\
\overline{put} \ y \ (\overline{ret} \ w) \ \overline{\parallel} \ \overline{put} \ y \ (\overline{ret} \ y) & \overline{put} \ y \ (\overline{ret} \ w) \ \overline{\parallel} \ \overline{put} \ y \ (\overline{ret} \ x) \\
= \{ \text{by (24)} \} & = \{ \text{by (24)} \} \\
\overline{put} \ y \ (\overline{ret} \ w \ \overline{\parallel} \ \overline{ret} \ y) & \overline{put} \ y \ (\overline{ret} \ w \ \overline{\parallel} \ \overline{ret} \ x)
\end{array}$$

(a) (b)

Fig. 4: Proof that law (24) implies that a bind operator cannot exist for the semantic domain.

for the continuation-passing style of these laws). As the given implementation records the order in which results are yielded by the computation, law (23) cannot be satisfied. And the example implementation also forms a monad, which means it is incompatible with law (24).

Machine-Verified Proofs From this point forward, we provide proofs mechanized in Coq for many theorems. When we do, we mark the proven statement with a check mark (✓).

5.3 An Implementation of the Semantic Domain

We present an implementation of Dom that satisfies the laws of section 5.2, and we provide machine-verified proofs to that effect. In the following implementation, we let Dom be the union of $\text{M } s \ a$ for all a and for a given s .

The implementation is based on a multiset or **Bag** data structure. In the mechanization, we implement $\text{Bag } a$ as a function $a \rightarrow \text{Nat}$.

```

type Bag a
singleton :: a → Bag a
emptyBag :: Bag a
sum      :: Bag a → Bag a → Bag a

```

We model a stateful, nondeterministic computation with global state semantics as a function that maps an initial state onto a bag of results, and a final state. Each result is a pair of the value returned, as well as the state at that point in the computation. The use of an unordered data structure to return the results of the computation is needed to comply with law (23).

In Section 5.2 we mentioned that, as a consequence of law (24) we must design the implementation of our semantic domain in such a way that it is impossible to define a bind operator $\overline{\gg} :: \text{Dom } a \rightarrow (a \rightarrow \text{Dom } b) \rightarrow \text{Dom } b$ for it. This is the case for our implementation: we only retain the final result of the branch without any information on how to continue the branch, which makes it impossible to define the bind operator.

type $\text{M } s \ a = s \rightarrow (\text{Bag } (a, s), s)$

$\overline{\emptyset}$ does not modify the state and produces no results. $\overline{\text{ret}}$ does not modify the state and produces a single result.

$$\begin{aligned} \overline{\emptyset} &:: \text{M } s \ a \\ \overline{\emptyset} &= \lambda s \rightarrow (\text{emptyBag}, s) \\ \overline{\text{ret}} &:: a \rightarrow \text{M } s \ a \\ \overline{\text{ret}} \ x &= \lambda s \rightarrow (\text{singleton } (x, s), s) \end{aligned}$$

$\overline{\text{get}}$ simply passes along the initial state to its continuation. $\overline{\text{put}}$ ignores the initial state and calls its continuation with the given parameter instead.

$$\begin{aligned} \overline{\text{get}} &:: (s \rightarrow \text{M } s \ a) \rightarrow \text{M } s \ a \\ \overline{\text{get}} \ k &= \lambda s \rightarrow k \ s \ s \\ \overline{\text{put}} &:: s \rightarrow \text{M } s \ a \rightarrow \text{M } s \ a \\ \overline{\text{put}} \ s \ k &= \lambda _ \rightarrow k \ s \end{aligned}$$

The $\overline{\parallel}$ operator runs the left computation with the initial state, then runs the right computation with the final state of the left computation, and obtains the final result by merging the two bags of results.

$$\begin{aligned} \overline{\parallel} &:: \text{M } s \ a \rightarrow \text{M } s \ a \rightarrow \text{M } s \ a \\ (xs \ \overline{\parallel} \ ys) \ s &= \mathbf{let} \ (ansx, s') = xs \ s \\ &\quad (ansy, s'') = ys \ s' \\ &\mathbf{in} \ (\text{sum } ansx \ ansy, s'') \end{aligned}$$

Lemma 1. *This implementation conforms to every law introduced in Section 5.2. ✓*

5.4 Contextual Equivalence

With our semantic domain sufficiently specified, we can prove analogous properties for programs interpreted through this domain. We must take care in how we reformulate these properties however. It is certainly not sufficient to merely copy the laws as formulated for the semantic domain, substituting **Prog** data constructors for semantic domain operators as needed; we must keep in mind that a term in **Prog** a describes a syntactical structure without ascribing meaning to it. For example, one cannot simply assert that $\text{Put } x \ (\text{Put } y \ p)$ is *equal*

```

data Ctx e1 a e2 b where
  □ :: Ctx e a e a
  COr1 :: Ctx e1 a e2 b → OProg e2 b
        → Ctx e1 a e2 b
  COr2 :: OProg e2 b → Ctx e1 a e2 b
        → Ctx e1 a e2 b
  CPut :: (Env e2 → S) → Ctx e1 a e2 b
        → Ctx e1 a e2 b
  CGet :: (S → Bool) → Ctx e1 a (S : e2) b
        → (S → OProg e2 b) → Ctx e1 a e2 b
  CBind1 :: Ctx e1 a e2 b → (b → OProg e2 c)
          → Ctx e1 a e2 c
  CBind2 :: OProg e2 a → Ctx e1 b (a : e2) c
          → Ctx e1 b e2 c
data Env (l :: [*]) where
  Nil :: Env '[]
  Cons :: a → Env l → Env (a : l)
type OProg e a = Env e → Prog a

```

(a)

(b)

Fig. 5: (a) Environments and open programs. (b) Syntax for contexts.

to `Put y p`, because although these two programs have the same semantics, they are not structurally identical. It is clear that we must define a notion of “semantic equivalence” between programs. We can map the syntactical structures in `Prog a` onto the semantic domain `Dom a` using `run` to achieve that. Yet wrapping both sides of an equation in `run` applications is not enough as such statements only apply at the top-level of a program. For instance, while `run (Put x (Put y p)) = run (Put y p)` is a correct statement, we cannot prove `run (Return w [] Put x (Put y p)) = run (Return w [] Put y p)` from such a law.

So the concept of semantic equivalence in itself is not sufficient; we require a notion of “contextual semantic equivalence” of programs which allows us to formulate properties about semantic equivalence which hold in any surrounding context. Figure 5(a) provides the definition for single-hole contexts `Ctx`. A context `C` of type `Ctx e1 a e2 b` can be interpreted as a function that, given a program that returns a value of type `a` under environment `e1` (in other words: the type and environment of the hole), produces a program that returns a value of type `b` under environment `e2` (the type and environment of the whole program). Filling the hole with `p` is denoted by `C[p]`. The type of environments, `Env` is defined using heterogeneous lists (Figure 5(b)). When we consider the notion of programs in contexts, we must take into account that these contexts may introduce variables which are referenced by the program. The `Prog` datatype however represents only closed programs. Figure 5(b) introduces the `OProg` type to represent “open” programs, and the `Env` type to represent environments. `OProg e a` is defined as the type of functions that construct a *closed* program of type `Prog a`, given an environment of type `Env e`. Environments, in turn, are defined as heterogeneous lists. We also define a function for mapping open programs onto the semantic domain.

```

orun :: OProg e a → Env e → Dom a
orun p env = run (p env) .

```

We can then assert that two programs are contextually equivalent if, for *any* context, running both programs wrapped in that context will yield the same result:

$$m_1 =_{\text{GS}} m_2 \triangleq \forall C. \text{orun} (C[m_1]) = \text{orun} (C[m_2]) \ .$$

We can then straightforwardly formulate variants of the state laws, the non-determinism laws and the *put-or* law for this global state monad as lemmas. For example, we reformulate law (16) as

$$\text{Put } s \ (\text{Put } t \ p) =_{\text{GS}} \text{Put } t \ p \ .$$

Proofs for the state laws, the nondeterminism laws and the *put-or* law then easily follow from the analogous semantic domain laws.

More care is required when we want to adapt law (24) into the **Prog** setting. At the end of Section 5.2, we saw that this law precludes the existence of a bind operator in the semantic domain. Since a bind operator for **Progs** exists, it might seem we're out of luck when we want to adapt law (24) to **Progs**. But because **Progs** are merely syntax, we have much more fine-grained control over what equality of programs means. When talking about the semantic domain, we only had one notion of equality: two programs are equal only when one can be substituted for the other in any context. So if, in that setting, we want to introduce a law that does not hold under a particular composition (in this case, sequential composition), the only way we can express that is to say that that composition is impossible in the domain. But the fact that **Prog** is defined purely syntactically opens up the possibility of defining multiple notions of equality which may exist at the same time. In fact, we have already introduced syntactic equality, semantic equality, and contextual equality. It is precisely this choice of granularity that allows us to introduce laws which only hold in programs of a certain form (non-contextual laws), while other laws are much more general (contextual laws). A direct adaptation of law (24) would look something like this:

$$\begin{aligned} \forall C. C \text{ is bind-free} &\Rightarrow \text{orun} (C[\text{Put } s \ (\text{Return } x \ \parallel \ p)]) \\ &= \text{orun} (C[\text{Put } s \ (\text{Return } x) \ \parallel \ \text{Put } s \ p]) \ . \end{aligned}$$

In other words, the two sides of the equation can be substituted for one another, but only in contexts which do not contain any binds. However, in our mechanization, we only prove a more restricted version where the context must be empty (in other words, what we called semantic equivalence), which turns out to be enough for our purposes.

$$\text{run} (\text{Put } s \ (\text{Return } x \ \parallel \ p)) = \text{run} (\text{Put } s \ (\text{Return } x) \ \parallel \ \text{Put } s \ p) \ .\checkmark \quad (25)$$

5.5 Simulating Local-State Semantics

We simulate local-state semantics by replacing each occurrence of **Put** by a variant that restores the state, as described in Section 4.3. This transformation is implemented by the function *trans* for closed programs, and *otrans* for open programs:

$$\begin{aligned}
trans &:: \text{Prog } a \rightarrow \text{Prog } a \\
trans (\text{Return } x) &= \text{Return } x \\
trans (p \parallel q) &= trans p \parallel trans q \\
trans \emptyset &= \emptyset \\
trans (\text{Get } p) &= \text{Get } (\lambda s \rightarrow trans (p s)) \\
trans (\text{Put } s p) &= \text{Get } (\lambda s' \rightarrow \text{Put } s (trans p) \parallel \text{Put } s' \emptyset) \text{ ,} \\
otrans &:: \text{OProg } e a \rightarrow \text{OProg } e a \\
otrans p &= \lambda env \rightarrow trans (p env) \text{ .}
\end{aligned}$$

We then define the function *eval*, which runs a transformed program (in other words, it runs a program with local-state semantics).

$$\begin{aligned}
eval &:: \text{Prog } a \rightarrow \text{Dom } a \\
eval &= run \cdot trans \text{ .}
\end{aligned}$$

We show that the transformation works by proving that our free monad equipped with *eval* is a correct implementation for a nondeterministic, stateful monad with local-state semantics. We introduce notation for “contextual equivalence under simulated backtracking semantics”:

$$m_1 =_{\text{LS}} m_2 \triangleq \forall C. eval (C[m_1]) = eval (C[m_2]) \text{ .}$$

For example, we formulate the statement that the *put-put* law (16) holds for our monad as interpreted by *eval* as

$$\text{Put } s (\text{Put } t p) =_{\text{LS}} \text{Put } t p \text{ .}\checkmark$$

Proofs for the nondeterminism laws follow trivially from the nondeterminism laws for global state. The state laws are proven by promoting *trans* inside, then applying global-state laws. For the proof of the *get-put* law, we require the property that in global-state semantics, *Put* distributes over (\parallel) if the left branch has been transformed (in which case the left branch leaves the state unmodified). This property only holds at the top-level.

$$run (\text{Put } x (trans m_1 \parallel m_2)) = run (\text{Put } x (trans m_1) \parallel \text{Put } x m_2) \text{ .}\checkmark \quad (26)$$

Proof of this lemma depends on law (24).

Finally, we arrive at the core of our proof: to show that the interaction of state and nondeterminism in this implementation produces backtracking semantics. To this end we prove laws analogous to the local state laws (13) and (12)

$$m \gg \emptyset =_{\text{LS}} \emptyset \text{ ,}\checkmark \quad (27)$$

$$m \gg (\lambda x \rightarrow f_1 x \parallel f_2 x) =_{\text{LS}} (m \gg f_1) \parallel (m \gg f_2) \text{ .}\checkmark \quad (28)$$

We provide machine-verified proofs for these theorems. The proof for (27) follows by straightforward induction. The inductive proof (with induction on *m*) of law (28) requires some additional lemmas.

For the case $m = m_1 \parallel m_2$, we require the property that, at the top-level of a global-state program, (\parallel) is commutative if both its operands are state-restoring. Formally:

$$run (trans\ p \parallel trans\ q) = run (trans\ q \parallel trans\ p) \ .\checkmark \quad (29)$$

The proof of this property motivated the introduction of law (23).

The proof for both the $m = \text{Get } k$ and $m = \text{Put } s\ m'$ cases requires that Get distributes over (\parallel) at the top-level of a global-state program if the left branch is state restoring.

$$run (\text{Get } (\lambda s \rightarrow trans\ (m_1\ s) \parallel (m_2\ s))) = \quad (30)$$

$$run (\text{Get } (\lambda s \rightarrow trans\ (m_1\ s)) \parallel \text{Get } m_2) \ .\checkmark \quad (31)$$

And finally, we require that the $trans$ function is, semantically speaking, idempotent, to prove the case $m = \text{Put } s\ m'$.

$$run (trans (trans\ p)) = run (trans\ p) \ .\checkmark \quad (32)$$

5.6 Backtracking with a Global State Monad

Although we can now interpret a local state program through translation to a global state program, we have not quite yet delivered on our promise to address the space usage issue of local state semantics. From the definition of put_r it is clear that we simply make the implicit copying of the local state semantics explicit in the global state semantics. As mentioned in Section 4.2, rather than using put , some algorithms typically use a pair of commands $modify\ next$ and $modify\ prev$, with $prev \cdot next = id$, to respectively update and roll back the state. This is especially true when the state is implemented using an array or other data structure that is usually not overwritten in its entirety. Following a style similar to put_r , this can be modelled by:

$$\begin{aligned} modify_r &:: \text{MStateNondet } s\ m \Rightarrow (s \rightarrow s) \rightarrow (s \rightarrow s) \rightarrow m\ () \\ modify_r\ next\ prev &= modify\ next \parallel side\ (modify\ prev) \ . \end{aligned}$$

Unlike put_r , $modify_r$ does not keep any copies of the old state alive, as it does not introduce a branching point where the right branch refers to a variable introduced outside the branching point. Is it safe to use an alternative translation, where the pattern $get \gg (\lambda s \rightarrow put\ (next\ s) \gg m)$ is not translated into $get \gg (\lambda s \rightarrow put_r\ (next\ s) \gg trans\ m)$, but rather into $modify_r\ next\ prev \gg trans\ m$? We explore this question by extending our Prog syntax with an additional Modify_r construct, thus obtaining a new Prog_M syntax:

data $\text{Prog}_M\ a$ **where**

...

$$\text{Modify}_r :: (S \rightarrow S) \rightarrow (S \rightarrow S) \rightarrow \text{Prog}_M\ a \rightarrow \text{Prog}_M\ a$$

We assume that $prev \cdot next = id$ for every $Modify_R \ next \ prev \ p$ in a $Prog_M \ a$ program.

We then define two translation functions from $Prog_M \ a$ to $Prog \ a$, which both replace Put s with put_R s along the way, like the regular $trans$ function. The first replaces each $Modify_R$ in the program by a direct analogue of the definition given above, while the second replaces it by $Get \ (\lambda s \rightarrow Put \ (next \ s) \ (trans_2 \ p))$:

$$\begin{aligned}
trans_1 &:: Prog_M \ a \rightarrow Prog \ a \\
&\dots \\
trans_1 \ (Modify_R \ next \ prev \ p) &= Get \ (\lambda s \rightarrow Put \ (next \ s) \ (trans_1 \ p)) \\
&\quad \parallel Get \ (\lambda t \rightarrow Put \ (prev \ t) \ \emptyset) \\
trans_2 &:: Prog_M \ a \rightarrow Prog \ a \\
&\dots \\
trans_2 \ (Modify_R \ next \ prev \ p) &= Get \ (\lambda s \rightarrow put_R \ (next \ s) \ (trans_2 \ p)) \\
&\quad \textbf{where} \ put_R \ s \ p = Get \ (\lambda t \rightarrow Put \ s \ p \parallel Put \ t \ \emptyset)
\end{aligned}$$

It is clear that $trans_2 \ p$ is the exact same program as $trans \ p'$, where p' is p but with each $Modify_R \ next \ prev \ p$ replaced by $Get \ (\lambda s \rightarrow Put \ (next \ s) \ p)$.

We then prove that these two transformations lead to semantically identical instances of $Prog \ a$.

Lemma 2. $run \ (trans_1 \ p) = run \ (trans_2 \ p)$. \checkmark

This means that, if we make some effort to rewrite parts of our program to use the $Modify_R$ construct rather than Put , we can use the more efficient translation scheme $trans_1$ to avoid introducing unnecessary copies.

n-Queens using a global state To wrap up, we revisit the n -queens puzzle. Recall that, for the puzzle, the operator that alters the state (to check whether a chess placement is safe) is defined by

$$(i, us, ds) \oplus x = (1 + i, (i + x) : us, (i - x) : ds) \ .$$

By defining $(i, us, ds) \ominus x = (i - 1, tail \ us, tail \ ds)$, we have $(\ominus x) \cdot (\oplus x) = id$. One may thus compute all solutions to the puzzle, in a scenario with a shared global state, by $run \ (queens_R \ n)$, where

$$\begin{aligned}
queens_R \ n &= put \ (0, [], []) \gg qBody \ [0..n-1] \ , \\
qBody \ [] &= return \ [] \\
qBody \ xs &= select \ xs \ \gg \ \lambda(x, ys) \rightarrow \\
&\quad (get \ \gg \ (guard \cdot ok \cdot (\oplus x))) \gg \\
&\quad modify_R \ (\oplus x) \ (\ominus x) \ \gg \ ((x:) \ \$ \ qBody \ ys) \ , \\
\textbf{where} \ (i, us, ds) \oplus x &= (1 + i, (i + x) : us, (i - x) : ds) \\
(i, us, ds) \ominus x &= (i - 1, tail \ us, \quad tail \ ds) \\
ok \ (_, u : us, d : ds) &= (u \notin us) \wedge (d \notin ds) \ .
\end{aligned}$$

6 Related Work

6.1 Prolog

Prolog is a prominent example of a system that exposes nondeterminism with local state to the user, but is itself implemented in terms of a single global state.

Warren Abstract Machine The folklore idea of undoing modifications upon backtracking is a key feature of many Prolog implementations, in particular those based on the Warren Abstract Machine (WAM) [1]. The WAM's global state is the program heap and Prolog programs modify this heap during unification only in a very specific manner: following the union-find algorithm, they overwrite cells that contain self-references with pointers to other cells. Undoing these modifications only requires knowledge of the modified cell's address, which can be written back in that cell during backtracking. The WAM has a special stack, called the trail stack, for storing these addresses, and the process of restoring those cells is called *untrailing*.

The 4-Port Box Model While trailing happens under the hood, there is a folklore Prolog programming pattern for observing and intervening at different points in the control flow of a procedure call, known as the *4-port box model*. In this model, upon the first entrance of a Prolog procedure it is *called*; it may yield a result and *exits*; when the subsequent procedure fails and backtracks, it is asked to *redo* its computation, possibly yielding the next result; finally it may *fail*. Given a Prolog procedure p implemented in Haskell, the following program prints debugging messages when each of the four ports are used:

$$\begin{aligned} & (\text{putStr "call" } \parallel \text{ side } (\text{putStr "fail"})) \gg \\ & p \gg \lambda x \rightarrow \\ & (\text{putStr "exit" } \parallel \text{ side } (\text{putStr "redo"})) \gg \text{return } x \end{aligned}$$

This technique was applied in the monadic setting by Hinze [6], and it has been our inspiration for expressing the state restoration with global state.

6.2 Reasoning About Side Effects

There are many works on reasoning and modelling side effects. Here we cover those that have most directly inspired this paper.

Axiomatic Reasoning Our work was directly inspired by Gibbons and Hinze's proposal to reason axiomatically about programs with side effects, and their axiomatic characterisation of local state in particular [4]. We have extended their work with an axiomatic characterisation of global state and on handling the former in terms of the latter. We also provide models that satisfy the axioms, whereas their paper mistakenly claims that one model satisfies the local state axioms and that another model is monadic.

Algebraic Effects Our formulation of implementing local state with global state is directly inspired by the effect handlers approach of Plotkin and Pretnar [12]. By making the free monad explicit our proofs benefit directly from the induction principle that Bauer and Pretnar establish for effect handler programs [2].

While Lawvere theories were originally Plotkin’s inspiration for studying algebraic effects, the effect handlers community has for a long time paid little attention to them. Yet, recently Lukšič and Pretnar [10] have investigated a framework for encoding axioms (or “effect theories”) in the type system: the type of an effectful function declares the operators used in the function, as well as the equalities that handlers for these operators should comply with. The type of a handler indicates which operators it handles and which equations it complies with. This type system would allow us to express at the type-level that our handler interprets local state in terms of global state.

7 Conclusions

Starting from Gibbons and Hinze’s observation [4] that the interaction between state and nondeterminism can be characterized axiomatically in multiple ways, we explored the differences between local state semantics (as characterised by Gibbons and Hinze) and global state semantics (for which we gave our own non-monadic characterisation).

In global state semantics, we find that we may use (\parallel) to simulate sequencing, and that the idea can be elegantly packaged into commands like put_r and $modify_r$. The interaction between global state and non-determinism turns out to be rather tricky. For a more rigorous treatment, we enforce a more precise separation between syntax and semantics and, as a side contribution of this paper, propose a *global state law*, plus some additional laws, which the semantics should satisfy. We verified (with the help of the Coq proof assistant) that there is an implementation satisfying these laws.

Using the n -queens puzzle as an example, we showed that one can end up in a situation where a problem is naturally expressed with local state semantics, but the greater degree of control over resources that global state semantics offers is desired. We then describe a technique to systematically transform a monadic program written against the local state laws into one that, when interpreted under global state laws, produces the same results as the original program. This transformation can be viewed as a handler (in the algebraic effects sense): it implements the interface of one effect in terms of the interface of another. We also verified the correctness of this transformation in Coq.

Acknowledgements We would like to thank Matija Pretnar, the members of IFIP WG 2.1, the participants of Shonan meeting 146 and the MPC reviewers for their insightful comments. We would also like to thank the Flemish Fund for Scientific Research (FWO) for their financial support.

References

1. Ait-Kaci, H.: Warren's Abstract Machine: A Tutorial Reconstruction (01 1991)
2. Bauer, A., Pretnar, M.: An effect system for algebraic effects and handlers. *Logical Methods in Computer Science* **10**(4) (2014), [https://doi.org/10.2168/LMCS-10\(4:9\)2014](https://doi.org/10.2168/LMCS-10(4:9)2014)
3. Gale, Y.: ListT done right alternative. https://wiki.haskell.org/ListT_done_right_alternative (2007)
4. Gibbons, J., Hinze, R.: Just do it: simple monadic equational reasoning. In: Danvy, O. (ed.) *International Conference on Functional Programming*. pp. 2–14. ACM Press (2011)
5. Gill, A., Kmett, E.: The monad transformer library. <https://hackage.haskell.org/package/mtl> (2014)
6. Hinze, R.: Prolog's control constructs in a functional setting - axioms and implementation. *Int. J. Found. Comput. Sci.* **12**(2), 125–170 (2001), <https://doi.org/10.1142/S0129054101000436>
7. Hutton, G., Fulger, D.: Reasoning about effects: seeing the wood through the trees. In: *Symposium on Trends in Functional Programming* (2007)
8. Kiselyov, O.: Laws of monadplus. <http://okmij.org/ftp/Computation/monads.html#monadplus> (2015)
9. Kiselyov, O., Ishii, H.: Freer monads, more extensible effects. In: Reppy, J.H. (ed.) *Symposium on Haskell*. pp. 94–105. ACM Press (2015)
10. Lukšič, v., Pretnar, M.: Local algebraic effect theories (2019), submitted
11. Moggi, E.: Computational lambda-calculus and monads. In: Parikh, R. (ed.) *Logic in Computer Science*. pp. 14–23. IEEE Computer Society Press (1989)
12. Plotkin, G., Pretnar, M.: Handlers of algebraic effects. In: Castagna, G. (ed.) *European Symposium on Programming*. pp. 80–94. No. 5502 in *Lecture Notes in Computer Science* (2009)
13. Volkov, N.: The list-t package. <http://hackage.haskell.org/package/list-t> (2014)
14. Wadler, P.: Monads for functional programming. In: Broy, M. (ed.) *Program Design Calculi: Marktoberdorf Summer School*. pp. 233–264. Springer-Verlag (1992)
15. Wu, N., Schrijvers, T., Hinze, R.: Effect handlers in scope. In: Voigtländer, J. (ed.) *Symposium on Haskell*. pp. 1–12. ACM Press (2012)