THE UNIVERSITY of EDINBURGH

Edinburgh Research Explorer

Reflections on monadic lenses

Reflections on monadic lenses

Faris Abou-Saleh[1], James Cheney[2], Jeremy Gibbons[1], James McKinna[2], and
Perdita Stevens[2]

[1] University of Oxford, `firstname.lastname@cs.ox.ac.uk`
[2] University of Edinburgh, `firstname.lastname@ed.ac.uk`

Abstract. Bidirectional transformations (bx) have primarily been mod-
eled as pure functions, and do not account for the possibility of the
side-effects that are available in most programming languages. Recently
several formulations of bx that use monads to account for effects have
been proposed, both among practitioners and in academic research. The
combination of bx with effects turns out to be surprisingly subtle, leading
to problems with some of these proposals and increasing the complexity
of others. This paper reviews the proposals for monadic lenses to date,
and offers some improved definitions, paying particular attention to the
obstacles to naively adding monadic effects to existing definitions of pure
bx such as lenses and symmetric lenses, and the subtleties of equivalence
of symmetric bidirectional transformations in the presence of effects.

1 Introduction

Programming with multiple concrete representations of the same conceptual
information is a commonplace, and challenging, problem. It is commonplace
because data is everywhere, and not all of it is relevant or appropriate for every
task: for example, one may want to work with only a subset of one's full email
account on a mobile phone or other low-bandwidth device. It is challenging
because the most direct approach to mapping data across sources A and B is to
write separate functions, one mapping to B and one to A, following some (not
always explicit) specification of what it means for an A value and a B value to
be *consistent*. Keeping these transformations coherent with each other, and with
the specification, is a considerable maintenance burden, yet it remains the main
approach found in practice.

Over the past decade, a number of promising proposals to ease programming
such *bidirectional transformations* have emerged, including *lenses* (Foster et al.,
2007), bx based on consistency relations (Stevens, 2010), *symmetric lenses* (Hof-
mann et al., 2011), and a number of variants and extensions (e.g. (Pacheco
et al., 2014; Johnson and Rosebrugh, 2014)). Most of these proposals consist
of an interface with pure functions and some equational laws that characterise
good behaviour; the interaction of bidirectionality with other effects has received
comparatively little attention.

Some programmers and researchers have already proposed ways to combine
lenses and monadic effects (Diviánszky, 2013; Pacheco et al., 2014). Recently,

we have proposed symmetric notions of bidirectional computation based on *entangled state monads* (Cheney et al., 2014; Abou-Saleh et al., 2015a) and *coalgebras* (Abou-Saleh et al., 2015b). As a result, there are now several alternative proposals for bidirectional transformations with effects. While this diversity is natural and healthy, reflecting an active research area, the different proposals tend to employ somewhat different terminology, and the relationships among them are not well understood. Small differences in definitions can have disproportionate impact.

In this paper we summarize and compare the existing proposals, offer some new alternatives, and attempt to provide general and useful definitions of "monadic lenses" and "symmetric monadic lenses". Perhaps surprisingly, it appears challenging even to define the composition of lenses in the presence of effects, especially in the symmetric case. We first review the definition of pure asymmetric lenses and two prior proposals for extending them with monadic effects. These definitions have some limitations, and we propose a new definition of monadic lens that overcomes them.

Next we consider the symmetric case. The effectful bx and coalgebraic bx in our previous work are symmetric, but their definitions rely on relatively heavyweight machinery (monad transformers and morphisms, coalgebra). It seems natural to ask whether just adding monadic effects to symmetric lenses in the style of Hofmann et al. (Hofmann et al., 2011) would also work. We show that, as for asymmetric lenses, adding monadic effects to symmetric lenses is challenging, and give examples illustrating the problems with the most obvious generalisation. We then briefly discuss our recent work on symmetric forms of bx with monadic effects (Cheney et al., 2014; Abou-Saleh et al., 2015a,b). Defining composition for these approaches also turns out to be tricky, and our definition of monadic lenses arose out of exploring this space. The essence of composition of symmetric monadic bx, we now believe, can be presented most easily in terms of monadic lenses, by considering *spans*, an approach also advocated (in the pure case) by Johnson and Rosebrugh (2014).

Symmetric pure bx need to be equipped with a notion of equivalence, to abstract away inessential differences of representation of their "state" or "complement" spaces. As noted by Hofmann et al. (2011) and Johnson and Rosebrugh (2014), isomorphism of state spaces is unsatisfactory, and there are competing proposals for equivalence of symmetric lenses and spans. In the case of spans of monadic lenses, the right notion of equivalence seem even less obvious. We compare three, increasingly coarse, equivalences of spans based on isomorphism (following Abou-Saleh et al. (2015a)), span equivalence (following Johnson and Rosebrugh (2014)), and bisimulation (following Hofmann et al. (2011) and Abou-Saleh et al. (2015b)). In addition, we show a (we think surprising) result: in the pure case, span equivalence and bisimulation equivalence coincide.

In this paper we employ Haskell-like notation to describe and compare formalisms, with a few conventions: we write function composition $f \cdot g$ with a centred dot, and use a lowered dot for field lookup $x.f$, in contrast to Haskell's notation $f\ x$. Throughout the paper, we introduce a number of different repre-

sentations of lenses, and rather than pedantically disambiguating them all, we freely redefine identifiers as we go. We assume familiarity with common uses of monads in Haskell to encapsulate effects (following Wadler (1995)), and with the **do**-notation (following Wadler's monad comprehensions Wadler (1992)). Although some of these ideas are present or implicit in recent papers (Hofmann et al., 2011; Johnson and Rosebrugh, 2014; Cheney et al., 2014; Abou-Saleh et al., 2015a,b), this paper reflects our desire to clarify these ideas and expose them in their clearest form — a desire that is strongly influenced by Wadler's work on a wide variety of related topics (Wadler, 1992; King and Wadler, 1992; Wadler, 1995), and by our interactions with him as a colleague.

2   Asymmetric monadic lenses

Recall that a *lens* (Foster et al., 2007, 2012) is a pair of functions, usually called *get* and *put*:

$$\textbf{data } \alpha \rightsquigarrow \beta = Lens\ \{\,get :: \alpha \to \beta, put :: \alpha \to \beta \to \alpha\,\}$$

satisfying (at least) the following *well-behavedness* laws:

(GetPut) $put\ a\ (get\ a) = a$
(PutGet) $get\ (put\ a\ b) = b$

The idea is that a lens of type $A \rightsquigarrow B$ maintains a source of type A, providing a view of type B onto it; the well-behavedness laws capture the intuition that the view faithfully reflects the source: if we "get" a b from a source a and then "put" the same b value back into a, this leaves a unchanged; and if we "put" a b into a source a and then "get" from the result, we get b itself. Lenses are often equipped with a *create* function

$$\textbf{data } \alpha \rightsquigarrow \beta = Lens\ \{\,get :: \alpha \to \beta, put :: \alpha \to \beta \to \alpha, create :: \beta \to \alpha\,\}$$

satisfying an additional law:

(CreateGet) $get\ (create\ b) = b$

When the distinction is important, we use the term *full* for well-behaved lenses equipped with a *create* operation. It is easy to show that the source and view types of a full lens must either both be empty or both non-empty, and that the *get* operation of a full lens is surjective.

Lenses have been investigated extensively; see for example Foster et al. (2012) for a recent tutorial overview. For the purposes of this paper, we just recall the definition of *composition* of lenses:

$$(;) :: (\alpha \rightsquigarrow \beta) \to (\beta \rightsquigarrow \gamma) \to (\alpha \rightsquigarrow \gamma)$$
$$l_1\ ;\ l_2 = Lens\ (l_2.get \cdot l_1.get)$$
$$\qquad\qquad (\lambda a\ c \to l_1.put\ a\ (l_2.put\ (l_1.get\ a)\ c))$$
$$\qquad\qquad (l_1.create \cdot l_1.create)$$

which preserves well-behavedness.

2.1   A naive approach

As a first attempt, consider simply adding a monadic effect μ to the result types of both *get* and *put*.

$$\textbf{data } [\alpha \leadsto_0 \beta]_\mu = MLens_0 \left\{ mget :: \alpha \to \mu\, \beta,\, mput :: \alpha \to \beta \to \mu\, \alpha \right\}$$

Such an approach has been considered and discussed in some recent Haskell libraries and online discussions (Diviánszky, 2013). A natural question arises immediately: what laws should a lens $l :: [A \leadsto_0 B]_M$ satisfy? The following generalisations of the laws appear natural:

$(\mathsf{MGetPut_0}) \quad \textbf{do } \{ b \leftarrow mget\ a;\ mput\ a\ b \} \quad = return\ a$
$(\mathsf{MPutGet_0}) \quad \textbf{do } \{ a' \leftarrow mput\ a\ b;\ mget\ a' \} = \textbf{do } \{ a' \leftarrow mput\ a\ b;\ return\ b \}$

that is, if we "get" b from a and then "put" the same b value back into a, this has the same effect as just returning a (and doing nothing else), and if we "put" a value b and then "get" the result, this has the same effect as just returning b after doing the "put". The obvious generalisation of composition from the pure case for these operations is:

$(;) :: [\alpha \leadsto_0 \beta]_\mu \to [\beta \leadsto_0 \gamma]_\mu \to [\alpha \leadsto_0 \gamma]_\mu$
$l_1\ ;\ l_2 = MLens_0\ (\lambda a \to \textbf{do } \{ b \leftarrow l_1.mget\ a;\ l_2.mget\ b \})$
$\qquad\qquad\qquad\quad (\lambda a\ c \to \textbf{do } \{ b \leftarrow l_1.mget\ a;\ b' \leftarrow l_2.mput\ b\ c;\ l_1.mput\ a\ b' \})$

This proposal has at least two apparent problems. First, the $(\mathsf{MGetPut_0})$ law appears to sharply constrain *mget*: indeed, if *mget a* has an irreversible side-effect then $(\mathsf{MGetPut_0})$ cannot hold. This suggests that *mget* must either be pure, or have side-effects that are reversible by *mput*, ruling out behaviours such as performing I/O during *mget*. Second, it appears difficult to compose these structures in a way that preserves the laws, unless we again make fairly draconian assumptions about μ. In order to show $(\mathsf{MGetPut_0})$ for the composition $l_1\ ;\ l_2$, it seems necessary to be able to commute $l_2.mget$ with $l_1.mget$ and we also need to know that doing $l_1.mget$ twice is the same as doing it just once. Likewise, to show $(\mathsf{MPutGet_0})$ we need to commute $l_2.mget$ with $l_1.mput$.

2.2   Monadic put-lenses

Pacheco et al. (2014) proposed a variant of lenses called *monadic putback-oriented lenses*. For the purposes of this paper, the putback-orientation of their approach is irrelevant: we focus on their use of monads, and we provide a slightly simplified version of their definition:

$$\textbf{data } [\alpha \leadsto_1 \beta]_\mu = MLens_1 \left\{ mget :: \alpha \to \beta,\, mput :: \alpha \to \beta \to \mu\, \alpha \right\}$$

The main difference from their version is that we remove the *Maybe* type constructors from the return type of *mget* and the first argument of *mput*. Pacheco

et al. state laws for these monadic lenses. First, they assume that the monad μ has a *monad membership* operation

$$(\in) :: \alpha \to \mu\, \alpha \to Bool$$

satisfying the following two laws:

$$\begin{array}{ll} (\in\text{-ID}) & x \in return\ x \iff True \\ (\in\text{-}\ggg) & y \in (m \ggg f) \iff \exists x\ .\ x \in m \wedge y \in (f\ x) \end{array}$$

Then the laws for $MLens_1$ (adapted from Pacheco et al. (2014, Prop. 3, p49)) are as follows:

$$\begin{array}{ll} (\mathsf{MGetPut_1}) & v = mget\ s \implies mput\ s\ v = return\ s \\ (\mathsf{MPutGet_1}) & s' \in mput\ s\ v' \implies v' = mget\ s' \end{array}$$

In the first law we correct an apparent typo in the original paper, as well as removing the *Just* constructors from both laws. By making *mget* pure, this definition avoids the immediate problems with composition discussed above, and Pacheco et al. outline a proof that their laws are preserved by composition. However, it is not obvious how to generalise their approach beyond monads that admit a sensible $\in$ operation.

Many interesting monads do have a sensible $\in$ operation (e.g. *Maybe*, $[\,]$). Pacheco et al. suggest that $\in$ can be defined for any monad as $x \in m \equiv (\exists h : h\, m = x)$, where h is what they call a "(polymorphic) algebra for the monad at hand, essentially, a of type $m\ a \to a$ for any type a." However, this definition doesn't appear satisfactory for monads such as IO, for which there is no such (pure) function: the $(\in\text{-ID})$ law can never hold in this case. It is not clear that we can define a useful $\in$ operation directly for IO either: given that $m :: IO\ a$ could ultimately return any a-value, it seems safe, if perhaps overly conservative, to define $x \in m = True$ for any x and m. This satisfies the $\in$ laws, at least, if we make a simplifying assumption that all types are inhabited, and indeed, it seems to be the only thing we could write in Haskell that would satisfy the laws, since we have no way of looking inside the monadic computation $m :: IO\ a$ to find out what its eventual return value is. But then the precondition of the $(\mathsf{MPutGet_1})$ law is always true, which forces the view space to be trivial. These complications suggest, at least, that it would be advantageous to find a definition of monadic lenses that makes sense, and is preserved under composition, for any monad.

2.3 Monadic lenses

We propose the following definition of monadic lenses for any monad M:

Definition 2.1 (monadic lens). A *monadic lens* from source type A to view type B in which the put operation may have effects from monad M (or "M-lens from A to B"), is represented by the type $[A \rightsquigarrow B]_M$, where

$$\mathbf{data}\ [\alpha \rightsquigarrow \beta]_\mu = MLens\ \{\, mget :: \alpha \to \beta,\, mput :: \alpha \to \beta \to \mu\, \alpha \,\}$$

(dropping the μ from the return type of *mget*, compared to the definition in Section 2.1). We say that M-lens l is *well-behaved* if it satisfies

> (MGetPut)   **do** $\{\, l.mput\ a\ (l.mget\ a)\,\} = return\ a$
> (MPutGet)   **do** $\{\, a' \leftarrow l.mput\ a\ b; k\ a'\ (l.mget\ a')\,\}$
> $= \mathbf{do}\ \{\, a' \leftarrow l.mput\ a\ b; k\ a'\ b\,\}$ $\diamond$

Note that in (MPutGet), we use a continuation $k :: \alpha \rightarrow \beta \rightarrow \mu\ \gamma$ to quantify over all possible subsequent computations in which a' and $l.mget\ a'$ might appear. In fact, using the laws of monads and simply-typed lambda calculus we can prove this law from just the special case $k = \lambda a\ b \rightarrow return\ (a, b)$, so in the sequel when we prove (MPutGet) we may just prove this case while using the strong form freely in the proof.

The ordinary asymmetric lenses are exactly the monadic lenses over $\mu = Id$; the laws then specialise to the standard equational laws. Monadic lenses where $\mu = Id$ are called *pure*, and we may refer to ordinary lenses as pure lenses also.

Definition 2.2. We can also define an operation that lifts a pure lens to a monadic lens:

> $lens2mlens :: Monad\ \mu \Rightarrow \alpha \rightsquigarrow \beta \rightarrow [\alpha \rightsquigarrow \beta]_\mu$
> $lens2mlens\ l = MLens\ (l.get)\ (\lambda a\ b \rightarrow return\ (l.put\ a\ b))$ $\diamond$

Lemma 2.3. If $l :: Lens\ \alpha\ \beta$ is well-behaved, then so is *lens2mlens l*. $\diamond$

Example 2.4. To illustrate, some simple pure lenses include:

> $id_l :: \alpha \rightsquigarrow \alpha$
> $id_l = Lens\ (\lambda a \rightarrow a)\ (\lambda_-\ a \rightarrow a)$
> $fst_l :: (\alpha, \beta) \rightsquigarrow \alpha$
> $fst_l = MLens\ fst\ (\lambda(s_1, s_2)\ s_1' \rightarrow (s_1', s_2))$

Many more examples of pure lenses are to be found in the literature (Foster et al., 2007, 2012), all of which lift to well-behaved monadic lenses. $\diamond$

As more interesting examples, we present asymmetric versions of the partial and logging lenses presented by Abou-Saleh et al. (2015a). Pure lenses are usually defined using total functions, which means that *get* must be surjective whenever A is nonempty, and *put* must be defined for all source and view pairs. One way to accommodate partiality is to adjust the return type of *get* to *Maybe b* or give *put* the return type *Maybe a* to allow for failure if we attempt to put a b-value that is not in the range of *get*. In either case, the laws need to be adjusted somehow. Monadic lenses allow for partiality without requiring such an ad hoc change. A trivial example is

> $constMLens :: \beta \rightarrow [\alpha \rightsquigarrow \beta]_{Maybe}$
> $constMLens\ b = MLens\ (const\ b)$
> $(\lambda a\ b' \rightarrow \mathbf{if}\ b \mathrel{==} b'\ \mathbf{then}\ Just\ a\ \mathbf{else}\ Nothing)$

which is well-behaved because both sides of (MPutGet) fail if the view is changed to a value different from b. Of course, this example also illustrates that the *mget* function of a monadic lens need not be surjective.

As a more interesting example, consider:

$$absLens :: [Int \leadsto Int]_{Maybe}$$
$$absLens = MLens\ abs$$
$$(\lambda a\ b \to \textbf{if}\ b < 0$$
$$\textbf{then}\ Nothing$$
$$\textbf{else}\ Just\ (\textbf{if}\ a < 0\ \textbf{then} - b\ \textbf{else}\ b))$$

In the *mget* direction, this lens maps a source number to its absolute value; in the reverse direction, it fails if the view b is negative, and otherwise uses the sign of the previous source a to determine the sign of the updated source.

The following *logging lens* takes a pure lens l and, whenever the source value a changes, records the previous a value.

$$logLens :: Eq\ \alpha \Rightarrow \alpha \leadsto \beta \to [\alpha \leadsto \beta]_{Writer\ \alpha}$$
$$logLens\ l = MLens\ (l.get)\ (\lambda a\ b \to$$
$$\textbf{let}\ a' = l.put\ a\ b\ \textbf{in do}\ \{\textbf{if}\ a \neq a'\ \textbf{then}\ tell\ a\ \textbf{else}\ return\ ();\ return\ a'\})$$

We presented a number of more involved examples of effectful symmetric bx in (Abou-Saleh et al., 2015a). They show how monadic lenses can employ user interaction, state, or nondeterminism to restore consistency. Most of these examples are equivalently definable as *spans* of monadic lenses, which we will discuss in the next section.

In practical use, it is usually also necessary to equip lenses with an *initialisation* mechanism. Indeed, as already mentioned, Pacheco et al.'s monadic put-lenses make the α argument optional (using *Maybe*), to allow for initialisation when only a β is available; we chose to exclude this from our version of monadic lenses above.

We propose the following alternative:

$$\textbf{data}\ [\alpha \leadsto \beta]_\mu = MLens\ \{\,mget :: \alpha \to \beta,$$
$$mput :: \alpha \to \beta \to \mu\ \alpha,$$
$$mcreate :: \beta \to \mu\ \alpha\,\}$$

and we consider such initialisable monadic lenses to be well-behaved when they satisfy the following additional law:

(MCreateGet) $\textbf{do}\ \{\,a \leftarrow mcreate\ b;\ k\ a\ (mget\ a)\,\} = \textbf{do}\ \{\,a \leftarrow mcreate\ b;\ k\ a\ b\,\}$

As with (MPutGet), this property follows from the special case $k = \lambda x\ y \to return\ (x, y)$, and we will use this fact freely.

This approach, in our view, helps keep the (GetPut) and (PutGet) laws simple and clear, and avoids the need to wrap *mput*'s first argument in *Just* whenever it is called.

Next, we consider composition of monadic lenses.

$$(;) :: Monad\ \mu \Rightarrow [\alpha \rightsquigarrow \beta]_\mu \rightarrow [\beta \rightsquigarrow \gamma]_\mu \rightarrow [\alpha \rightsquigarrow \gamma]_\mu$$
$$l_1\ ;\ l_2 = MLens\ (l_2.mget \cdot l_1.mget)\ mput\ mcreate\ \textbf{where}$$
$$mput\ a\ c = \textbf{do}\ \{\ b \leftarrow l_2.mput\ (l_1.mget\ a)\ c;\ l_1.mput\ a\ b\ \}$$
$$mcreate\ c = \textbf{do}\ \{\ b \leftarrow l_2.mcreate;\ l_1.mcreate\ \}$$

Note that we consider only the simple case in which the lenses share a common monad μ. Composing lenses with effects in different monads would require determining how to compose the monads themselves, which is nontrivial (King and Wadler, 1992; Jones and Duponcheel, 1993).

Theorem 2.5. If $l_1 :: [A \rightsquigarrow B]_M, l_2 :: [B \rightsquigarrow C]_M$ are well-behaved, then so is $l_1\ ;\ l_2$. $\diamondsuit$

3   Symmetric monadic lenses and spans

Hofmann et al. (2011) proposed *symmetric lenses* that use a *complement* to store (at least) the information that is not present in both views.

$$\textbf{data}\ \alpha \xleftrightarrow{\gamma} \beta = SLens\ \{\ put_\mathrm{R}\quad :: (\alpha, \gamma) \rightarrow (\beta, \gamma),$$
$$put_\mathrm{L} :: (\beta, \gamma) \rightarrow (\alpha, \gamma),$$
$$missing :: \gamma\ \}$$

Informally, put_R turns an α into a β, modifying a complement γ as it goes, and symmetrically for put_L; and *missing* is an initial complement, to get the ball rolling. Well-behavedness for symmetric lenses amounts to the following equational laws:

(PutRL) $\textbf{let}\ (b, c') = sl.put_\mathrm{R}\ (a, c)\ \textbf{in}\ sl.put_\mathrm{L}\ (b, c')$
$\qquad = \textbf{let}\ (b, c') = sl.put_\mathrm{R}\ (a, c)\ \textbf{in}\ (a, c')$
(PutLR) $\textbf{let}\ (a, c') = sl.put_\mathrm{L}\ (b, c)\ \textbf{in}\ sl.put_\mathrm{R}\ (a, c')$
$\qquad = \textbf{let}\ (a, c') = sl.put_\mathrm{L}\ (b, c)\ \textbf{in}\ (b, c')$

Furthermore, the composition of two symmetric lenses preserves well-behavedness, and can be defined as follows:

$$(;) :: (\alpha \xleftrightarrow{\sigma_1} \beta) \rightarrow (\beta \xleftrightarrow{\sigma_2} \gamma) \rightarrow (\alpha \xleftrightarrow{(\sigma_1, \sigma_2)} \gamma)$$
$$l_1\ ;\ l_2 = SLens\ put_\mathrm{R}\ put_\mathrm{L}\ (l_1.missing, l_2.missing)\ \textbf{where}$$
$$put_\mathrm{R}\ (a, (s_1, s_2)) = \textbf{let}\ (b, s_1') = put_\mathrm{R}\ (a, s_1)$$
$$(c, s_2') = put_\mathrm{R}\ (b, s_2)$$
$$\textbf{in}\ (c, (s_1', s_2'))$$
$$put_\mathrm{L}\ (c, (s_1, s_2)) = \textbf{let}\ (b, s_2') = put_\mathrm{L}\ (c, s_2)$$
$$(a, s_1') = put_\mathrm{L}\ (b, s_1)$$
$$\textbf{in}\ (a, (s_1', s_2'))$$

We can define an *identity* symmetric lens as follows:

$$id_{\mathrm{sl}} :: \alpha \xleftrightarrow{()} \alpha$$
$$id_{\mathrm{sl}} = SLens \; id \; id \; ()$$

It is natural to wonder whether symmetric lens composition satisfies identity and associativity laws making symmetric lenses into a category. This is complicated by the fact that the complement types of the composition $id_{\mathrm{sl}}; sl$ and of sl differ, so it is not even type-correct to ask whether $id_{\mathrm{sl}}; sl$ and sl are equal. To make it possible to relate the behaviour of symmetric lenses with different complement types, Hofmann *et al.* defined equivalence of symmetric lenses as follows:

Definition 3.1. Suppose $R \subseteq C_1 \times C_2$. Then $f \sim_R g$ means that for all c_1, c_2, x, if $(c_1, c_2) \in R$ and $(y, c_1') = f \; (x, c_1)$ and $(y', c_2') = g \; (y, c_2)$, then $y = y'$ and $(c_1', c_2') \in R$. $\diamondsuit$

Definition 3.2 (Symmetric lens equivalence). Two symmetric lenses $sl_1 :: X \xleftrightarrow{C_1} Y$ and $sl_2 :: X \xleftrightarrow{C_2} Y$ are considered *equivalent* ($sl_1 \equiv_{\mathrm{sl}} sl_2$) if there is a relation $R \subseteq C_1 \times C_2$ such that

1. $(sl_1.missing, sl_2.missing) \in R$,
2. $sl_1.put_{\mathrm{R}} \sim_R sl_2.put_{\mathrm{R}}$, and
3. $sl_1.put_{\mathrm{L}} \sim_R sl_2.put_{\mathrm{L}}$. $\diamondsuit$

Hofmann *et al.*show that $\equiv_{\mathrm{sl}}$ is an equivalence relation; moreover it is sufficiently strong to validate identity, associativity and congruence laws:

Theorem 3.3 ((Hofmann et al., 2011)). If $sl_1 :: X \xleftrightarrow{C_1} Y$ and $sl_2 :: Y \xleftrightarrow{C_2} Z$ are well-behaved, then so is $sl_1 \; ; \; sl_2$. In addition, composition satisfies the laws:

(Identity) $sl \; ; \; id_{\mathrm{sl}} \equiv_{\mathrm{sl}} sl \equiv_{\mathrm{sl}} id_{\mathrm{sl}} \; ; \; sl$
(Assoc) $sl_1 \; ; \; (sl_2 \; ; \; sl_3) \equiv_{\mathrm{sl}} (sl_1 \; ; \; sl_2) \; ; \; sl_3$
(Cong) $sl_1 \equiv_{\mathrm{sl}} sl_1' \wedge sl_2 \equiv_{\mathrm{sl}} sl_2' \Longrightarrow sl_1 \; ; \; sl_2 \equiv_{\mathrm{sl}} sl_1' \; ; \; sl_2'$ $\diamondsuit$

3.1   Naive monadic symmetric lenses

We now consider an obvious monadic generalisation of symmetric lenses, in which the put_{L} and put_{R} functions are allowed to have effects in some monad M:

Definition 3.4. A *monadic symmetric lens* from A to B with complement type C and effects M consists of two functions converting A to B and vice versa, each also operating on C and possibly having effects in M, and a complement value *missing* used for initialisation:

$$\mathbf{data} \; [\alpha \xleftrightarrow{\gamma} \beta]_\mu = SMLens \; \{ \, mput_{\mathrm{R}} \; :: (\alpha, \gamma) \to \mu \; (\beta, \gamma),$$
$$mput_{\mathrm{L}} \; :: (\beta, \gamma) \to \mu \; (\alpha, \gamma),$$
$$missing :: \gamma \, \}$$

Such a lens sl is called *well-behaved* if:

(PutRLM) $\mathbf{do} \; \{ (b, c') \leftarrow sl.mput_{\mathrm{R}} \; (a, c); sl.mput_{\mathrm{L}} \; (b, c') \}$
        $= \mathbf{do} \; \{ (b, c') \leftarrow sl.mput_{\mathrm{R}} \; (a, c); return \; (a, c') \}$
(PutLRM) $\mathbf{do} \; \{ (a, c') \leftarrow sl.mput_{\mathrm{L}} \; (b, c); sl.mput_{\mathrm{R}} \; (a, c') \}$
        $= \mathbf{do} \; \{ (a, c') \leftarrow sl.mput_{\mathrm{L}} \; (b, c); return \; (b, c') \}$ $\diamondsuit$

The above monadic generalisation of symmetric lenses appears natural, but it turns out to have some idiosyncrasies, similar to those of the naive version of monadic lenses we considered in Section 2.1.

Composition and well-behavedness Consider the following candidate definition of composition for monadic symmetric lenses:

$$(;) :: Monad\ \mu \Rightarrow [\alpha \xleftrightarrow{\sigma_1} \beta]_\mu \to [\beta \xleftrightarrow{\sigma_2} \gamma]_\mu \to [\alpha \xleftrightarrow{(\sigma_1,\sigma_2)} \gamma]_\mu$$

$$sl_1\ ;\ sl_2 = SMLens\ put_{\mathrm{R}}\ put_{\mathrm{L}}\ missing\ \textbf{where}$$
$$put_{\mathrm{R}}\ (a,(s_1,s_2)) = \textbf{do}\ \{(b,s_1') \leftarrow sl_1.mput_{\mathrm{R}}\ (a,s_1);$$
$$(c,s_2') \leftarrow sl_2.mput_{\mathrm{R}}\ (b,s_2);$$
$$return\ (c,(s_1',s_2'))\}$$
$$put_{\mathrm{L}}\ (c,(s_1,s_2)) = \textbf{do}\ \{(b,s_2') \leftarrow sl_2.mput_{\mathrm{L}}\ (c,s_2);$$
$$(a,s_1') \leftarrow sl_1.mput_{\mathrm{L}}\ (b,s_1);$$
$$return\ (a,(s_1',s_2'))\}$$
$$missing \qquad\qquad = (sl_1.missing, sl_2.missing)$$

which seems to be the obvious generalisation of pure symmetric lens composition to the monadic case. However, it does not always preserve well-behavedness.

Example 3.5. Consider the following construction:

$$setBool \quad :: Bool \to [() \xleftrightarrow{()} ()]_{State\ Bool}$$
$$setBool\ b = SMLens\ m\ m\ ()\ \textbf{where}\ m\ \_ = \textbf{do}\ \{set\ b; return\ ((),())\}$$

The lens *setBool True* has no effect on the complement or values, but sets the state to *True*. Both *setBool True* and *setBool False* are well-behaved, but their composition (in either direction) is not: (PutRLM) fails for *setBool True*; *setBool False* because *setBool True* and *setBool False* share a single *Bool* state value.   ◇

Proposition 3.6. *setBool b* is well-behaved for $b \in \{\ True, False\ \}$, but *setBool True*; *setBool False* is not well-behaved.   ◇

Composition does preserve well-behavedness for commutative monads, i.e. those for which

$$\textbf{do}\ \{\ a \leftarrow x; b \leftarrow y; return\ (a,b)\} = \textbf{do}\ \{\ b \leftarrow y; a \leftarrow x; return\ (a,b)\}$$

but this rules out many interesting monads, such as *State* and *IO*.

3.2   Entangled state monads

The types of the $mput_{\mathrm{R}}$ and $mput_{\mathrm{L}}$ operations of symmetric lenses can be seen (modulo mild reordering) as stateful operations in the *state monad State* $\gamma\ \alpha = \gamma \to (\alpha, \gamma)$, where the state $\gamma = C$. This observation was also anticipated by Hofmann et al. In a sequence of papers, we considered generalising these operations and their laws to an arbitrary monad (Cheney et al., 2014; Abou-Saleh et al., 2015a,b). In our initial workshop paper, we proposed the following definition:

$$\textbf{data } [\alpha \mathbin{\approx} \beta]_\mu = SetBX \{ get_L :: \mu\, \alpha, set_L :: \alpha \rightarrow \mu\, (),$$
$$get_R :: \mu\, \beta, set_R :: \beta \rightarrow \mu\, () \}$$

subject to a subset of the *State* monad laws (Plotkin and Power, 2002), such as:

$(\mathsf{Get_L\,Set_L})$ $\textbf{do } \{ a \leftarrow get_L; set_L\, a \} = return\, ()$
$(\mathsf{Set_L\,Get_L})$ $\textbf{do } \{ set_L\, a; get_L \}$ $= \textbf{do } \{ set_L\, a; return\, a \}$

This presentation makes clear that bidirectionality can be viewed as a state effect in which two "views" of some common state are *entangled*. That is, rather than storing a pair of views, each independently variable, they are entangled, in the sense that a change to either may also change the other. Accordingly, the entangled state monad operations do *not* satisfy all of the usual laws of state: for example, the set_L and set_R operations do not commute.

    However, one difficulty with the entangled state monad formalism is that, as discussed in Section 2.1, effectful *mget* operations cause problems for composition. It turned out to be nontrivial to define a satisfactory notion of composition, even for the well-behaved special case where $\mu = StateT\ \sigma\ \nu$ for some ν, where $StateT$ is the *state monad transformer* (Liang et al., 1995), i.e. $StateT\ \sigma\ \nu\ \alpha = \sigma \rightarrow \nu\ (\alpha, \sigma)$. We formulated the definition of *monadic lenses* given earlier in this paper in the process of exploring this design space.

3.3   Spans of monadic lenses

Hofmann et al. (2011) showed that a symmetric lens is equivalent to a *span* of two ordinary lenses, and later work by Johnson and Rosebrugh (2014) investigated such *spans of lenses* in greater depth. Accordingly, we propose the following definition:

Definition 3.7 (Monadic lens spans). A span of monadic lenses ("M-lens span") is a pair of M-lenses having the same source:

$$\textbf{type } [\alpha \mathbin{\leftarrowtail} \sigma \mathbin{\rightsquigarrow} \beta]_\mu = Span\, \{ left :: [\sigma \rightsquigarrow \alpha]_\mu, right :: [\sigma \rightsquigarrow \beta]_\mu \}$$

We say that an M-lens span is well-behaved if both of its components are. $\Diamond$

We first note that we can extend either leg of a span with a monadic lens (preserving well-behavedness if the arguments are well-behaved):

$(\triangleleft)$ $:: Monad\ \mu \Rightarrow [\alpha_1 \rightsquigarrow \alpha_2]_\mu \rightarrow [\alpha_1 \leftarrowtail \sigma \rightsquigarrow \beta]_\mu \rightarrow [\alpha_2 \leftarrowtail \sigma \rightsquigarrow \beta]_\mu$
$ml \triangleleft sp = Span\, (sp.left\, ;\, ml)\, (sp.right)$
$(\triangleright)$ $:: Monad\ \mu \Rightarrow [\alpha \leftarrowtail \sigma \rightsquigarrow \beta_1]_\mu \rightarrow [\beta_1 \rightsquigarrow \beta_2]_\mu \rightarrow [\alpha \leftarrowtail \sigma \rightsquigarrow \beta_2]_\mu$
$sp \triangleright ml = Span\, sp.left\, (sp.right\, ;\, ml)$

    To define composition, the basic idea is as follows. Given two spans $[A \leftarrowtail S_1 \rightsquigarrow B]_M$ and $[B \leftarrowtail S_2 \rightsquigarrow C]_M$ with a common type B "in the middle", we want to form a single span from A to C. The obvious thing to try is to form a pullback of the two monadic lenses from S_1 and S_2 to the common type B, obtaining a span

Fig. 1. Composing spans of lenses

from some common state type S to the state types S_1 and S_2, and composing with the outer legs. (See Figure 1.) However, the category of monadic lenses doesn't have pullbacks (as Johnson and Rosebrugh note, this is already the case for ordinary lenses). Instead, we construct the appropriate span as follows.

$$(\bowtie) :: Monad\ \mu \Rightarrow [\sigma_1 \rightsquigarrow \beta]_\mu \to [\sigma_2 \rightsquigarrow \beta]_\mu \to [\sigma_1 \leftsquigarrow (\sigma_1 \bowtie \sigma_2) \rightsquigarrow \sigma_2]_\mu$$
$$l_1 \bowtie l_2 = Span\ (MLens\ fst\ put_{\mathrm{L}}\ create_{\mathrm{L}})\ (MLens\ snd\ put_{\mathrm{R}}\ create_{\mathrm{R}})\ \textbf{where}$$
$$put_{\mathrm{L}}\ (\_, s_2)\ s_1' = \textbf{do}\ \{ s_2' \leftarrow l_2.mput\ s_2\ (l_1.mget\ s_1');\ return\ (s_1', s_2') \}$$
$$create_{\mathrm{L}}\ s_1 \quad = \textbf{do}\ \{ s_2' \leftarrow l_2.mcreate\ (l_1.mget\ s_1);\ return\ (s_1, s_2') \}$$
$$put_{\mathrm{R}}\ (s_1, \_)\ s_2' = \textbf{do}\ \{ s_1' \leftarrow l_1.mput\ s_1\ (l_2.mget\ s_2');\ return\ (s_1', s_2') \}$$
$$create_{\mathrm{R}}\ s_1 \quad = \textbf{do}\ \{ s_1' \leftarrow l_1.mcreate\ (l_2.mget\ s_2);\ return\ (s_1', s_2) \}$$

where we write $S_1 \bowtie S_2$ for the type of *consistent* state pairs $\{(s_1, s_2) \in S_1 \times S_2 \mid l_1.mget\ (s_1) = l_2.mget\ (s_2)\}$. In the absence of dependent types, we represent this type as (S_1, S_2) in Haskell, and we need to check that the *mput* and *mcreate* operations respect the consistency invariant.

Lemma 3.8. If $ml_1 :: [S_1 \rightsquigarrow B]_M$ and $ml_2 :: [S_2 \rightsquigarrow B]_M$ are well-behaved then so is $ml_1 \bowtie ml_2 :: [S_1 \leftsquigarrow (S_1 \bowtie S_2) \rightsquigarrow S_2]_\mu$. $\diamondsuit$

Note that (MPutGet) and (MCreateGet) hold by construction and do not need the corresponding properties for l_1 and l_2, but these properties are needed to show that consistency is established by *mcreate* and preserved by *mput*.

We can now define composition as follows:

$$(;) :: Monad\ \mu \Rightarrow [\alpha \leftsquigarrow \sigma_1 \rightsquigarrow \beta]_\mu \to [\beta \leftsquigarrow \sigma_2 \rightsquigarrow \gamma]_\mu \to [\alpha \leftsquigarrow (\sigma_1 \bowtie \sigma_2) \rightsquigarrow \gamma]_\mu$$
$$sp_1\ ;\ sp_2 = sp_1.left \triangleleft (sp_1.right \bowtie sp_2.left) \triangleright sp_2.right$$

The well-behavedness of the composition of two well-behaved spans is immediate because $\triangleleft$ and $\triangleright$ preserve well-behavedness of their arguments:

Theorem 3.9. If $sp_1 :: [A \leftsquigarrow S_1 \rightsquigarrow B]_M$ and $sp_2 :: [B \leftsquigarrow S_2 \rightsquigarrow C]_M$ are well-behaved spans of monadic lenses, then their composition $sp_1 ; sp_2$ is well-behaved.
$\diamondsuit$

Given a span of monadic lenses $sp :: [A \leftsquigarrow S \rightsquigarrow B]_M$, we can construct a monadic symmetric lens $sl :: [A \overset{Maybe\ S}{\longleftrightarrow} B]_M$ as follows:

$$span2smlens\ (left, right) = SMLens\ mput_{\mathrm{R}}\ mput_{\mathrm{L}}\ Nothing\ \textbf{where}$$
$$mput_{\mathrm{R}}\ (a, Just\ s) \quad = \textbf{do}\ \{ s' \leftarrow left.mput\ s\ a;\ return\ (right.mget\ s', Just\ s') \}$$

$$mput_{\mathrm{R}}\ (a, Nothing) = \mathbf{do}\ \{\, s' \leftarrow left.mcreate\ a; return\ (right.mget\ s', Just\ s')\,\}$$
$$mput_{\mathrm{L}}\ (b, Just\ s)\ \ = \mathbf{do}\ \{\, s' \leftarrow right.mput\ s\ b; return\ (left.mget\ s', Just\ s')\,\}$$
$$mput_{\mathrm{L}}\ (b, Nothing) = \mathbf{do}\ \{\, s' \leftarrow right.mcreate\ b; return\ (left.mget\ s', Just\ s')\,\}$$

Essentially, these operations use the span's *mput* and *mget* operations to update one side and obtain the new view value for the other side, and use the *mcreate* operations to build the initial S state if the complement is *Nothing*.

Well-behavedness is preserved by the conversion from monadic lens spans to *SMLens*, for arbitrary monads M:

Theorem 3.10. If $sp :: [A \leftarrowtail S \rightsquigarrow B]_M$ is well-behaved, then *span2smlens sp* is also well-behaved. $\diamondsuit$

Given $sl :: [A \xleftrightarrow{C} B]_M$, let $S \subseteq A \times B \times C$ be the set of *consistent triples* (a, b, c), that is, those for which $sl.mput_{\mathrm{R}}\ (a, c) = return\ (b, c)$ and $sl.mput_{\mathrm{L}}\ (b, c) = return\ (a, c)$. We construct $sp :: [A \leftarrowtail S \rightsquigarrow B]_M$ by

$$smlens2span\ sl = Span\ (MLens\ get_L\ put_{\mathrm{L}}\ create_{\mathrm{L}})\ (MLens\ get_R\ put_{\mathrm{R}}\ create_{\mathrm{R}})$$
$$\mathbf{where}$$
$$get_L\ (a, b, c)\ \ \ \ = a$$
$$put_{\mathrm{L}}\ (a, b, c)\ a' = \mathbf{do}\ \{\, (b', c') \leftarrow sl.mput_{\mathrm{R}}\ (a', c); return\ (a', b', c')\,\}$$
$$create_{\mathrm{L}}\ a\ \ \ \ \ \ \ \ = \mathbf{do}\ \{\, (b, c) \leftarrow sl.mput_{\mathrm{R}}\ (a, sl.missing); return\ (a, b, c)\,\}$$
$$get_R\ (a, b, c)\ \ \ \ = b$$
$$put_{\mathrm{R}}\ (a, b, c)\ b' = \mathbf{do}\ \{\, (a', c') \leftarrow sl.mput_{\mathrm{L}}\ (b', c); return\ (a', b', c')\,\}$$
$$create_{\mathrm{R}}\ b\ \ \ \ \ \ \ \ = \mathbf{do}\ \{\, (a, c) \leftarrow sl.mput_{\mathrm{L}}\ (b, sl.missing); return\ (a, b, c)\,\}$$

However, *smlens2span* may not preserve well-behavedness even for simple monads such as *Maybe*, as the following counterexample illustrates.

Example 3.11. Consider the following monadic symmetric lens construction:

$$fail :: [() \xleftrightarrow{()} ()]_{Maybe}$$
$$fail = SMLens\ Nothing\ Nothing\ ()$$

This is well-behaved but *smlens2span fail* is not. In fact, the set of consistent states of *fail* is empty, and each leg of the induced span is of the following form:

$$failMLens :: MLens\ Maybe\ \emptyset\ ()$$
$$failMLens = MLens\ (\lambda_- \rightarrow ())\ (\lambda_-\ () \rightarrow Nothing)\ (\lambda_- \rightarrow Nothing)$$

which fails to satisfy (MGetPut). $\diamondsuit$

For pure symmetric lenses, *smlens2span* does preserve well-behavedness.

Theorem 3.12. If $sl :: SMLens\ Id\ C\ A\ B$ is well-behaved, then *smlens2span sl* is also well-behaved, with state space S consisting of the consistent triples of sl. $\diamondsuit$

To summarize: spans of monadic lenses are closed under composition, and correspond to well-behaved symmetric monadic lenses. However, there are well-behaved symmetric monadic lenses that do not map to well-behaved spans. It seems to be an interesting open problem to give a direct axiomatisation of the symmetric monadic lenses that are essentially spans of monadic lenses (and are therefore closed under composition).

4   Equivalence of spans

Hofmann et al. (2011) introduced a bisimulation-like notion of equivalence for pure symmetric lenses, in order to validate laws such as identity, associativity and congruence of composition. Johnson and Rosebrugh (2014) introduced a definition of equivalence of spans and compared it with symmetric lens equivalence. We have considered equivalences based on isomorphism (Abou-Saleh et al., 2015a) and bisimulation (Abou-Saleh et al., 2015b). In this section we consider and relate these approaches in the context of spans of M-lenses.

Definition 4.1 (Isomorphism Equivalence). Two M-lens spans $sp_1 :: [A \leftharpoonup S_1 \rightsquigarrow B]_M$ and $sp_2 :: [A \leftharpoonup S_2 \rightsquigarrow B]_M$ are isomorphic ($sp \equiv_i sp'$) if there is an isomorphism $h :: S_1 \to S_2$ on their state spaces such that $h \, ; sp_2.left = sp_1.left$ and $h \, ; sp_2.right = sp_1.right$. $\diamondsuit$

Note that any isomorphism $h :: S_1 \to S_2$ can be made into a (monadic) lens; we omit the explicit conversion.

We consider a second definition of equivalence, inspired by Johnson and Rosebrugh (2014), which we call *span equivalence*:

Definition 4.2 (Span Equivalence). Two M-lens spans $sp_1 :: [A \leftharpoonup S_1 \rightsquigarrow B]_M$ and $sp_2 :: [A \leftharpoonup S_2 \rightsquigarrow B]_M$ are related by $\curvearrowright$ if there is a full lens $h :: S_1 \rightsquigarrow S_2$ such that $h \, ; sp_2.left = sp_1.left$ and $h \, ; sp_2.right = sp_1.right$. The equivalence relation $\equiv_s$ is the least equivalence relation containing $\curvearrowright$. $\diamondsuit$

One important consideration emphasized by Johnson and Rosebrugh is the need to avoid making all compatible spans equivalent to the "trivial" span $[A \leftharpoonup \emptyset \rightsquigarrow B]_M$. To avoid this problem, they imposed conditions on h: its *get* function must be surjective and *split*, meaning that there exists a function c such that $h.get \cdot c = id$. We chose instead to require h to be a full lens. This is actually slightly stronger than Johnson and Rosebrugh's definition, at least from a constructive perspective, because h is equipped with a specific choice of $c = create$ satisfying $h.get \cdot c = id$, that is, the (CreateGet) law.

We have defined span equivalence as the reflexive, symmetric, transitive closure of $\curvearrowright$. Interestingly, even though span equivalence allows for an arbitrary sequence of (pure) lenses between the respective state spaces, it suffices to consider only spans of lenses. To prove this, we first state a lemma about the ($\bowtie$) operation used in composition. Its proof is straightforward equational reasoning.

Lemma 4.3. Suppose $l_1 :: A \rightsquigarrow B$ and $l_2 :: C \rightsquigarrow B$ are pure lenses. Then $(l_1 \bowtie l_2).left \, ; l_1 = (l_1 \bowtie l_2).right \, ; l_2$. $\diamondsuit$

Theorem 4.4. Given $sp_1 :: [A \leftharpoonup S_1 \rightsquigarrow B]_M$ and $sp_2 :: [A \leftharpoonup S_2 \rightsquigarrow B]_M$, if $sp_1 \equiv_s sp_2$ then there exists $sp :: S_1 \leftharpoonup S \rightsquigarrow S_2$ such that $sp.left; sp_1.left = sp.right; sp_2.left$ and $sp.left \, ; sp_1.right = sp.right \, ; sp_2.right$. $\diamondsuit$

Proof. Let sp_1 and sp_2 be given such that $sp_1 \equiv_s sp_2$. The proof is by induction on the length of the sequence of $\curvearrowright$ or $\curvearrowleft$ steps linking sp_1 to sp_2.

If $sp_1 = sp_2$ then the result is immediate. If $sp_1 \curvearrowright sp_2$ then we can complete a span between S_1 and S_2 using the identity lens. For the inductive case, suppose that the result holds for sequences of up to $n \curvearrowright$ or $\curvearrowleft$ steps, and suppose $sp_1 \equiv_s sp_2$ holds in $n \curvearrowright$ or $\curvearrowleft$ steps. There are two cases, depending on the direction of the first step. If $sp_1 \curvearrowleft sp_3 \equiv_s sp_2$ then by induction we must have a pure span sp between S_3 and S_2 and $sp_1 \curvearrowleft sp_3$ holds by virtue of a lens $h :: S_3 \to S_1$, so we can simply compose h with $sp.left$ to obtain the required span between S_1 and S_2. Otherwise, if $sp_1 \curvearrowright sp_3 \equiv_s sp_2$ then by induction we must have a pure span sp between S_3 and S_2 and we must have a lens $h :: S_1 \to S_3$, so we use Lemma 4.3 to form a span $sp_0 :: S_1 \leftsquigarrow (S_1 \bowtie S_3) \rightsquigarrow S_3$ and extend $sp_0.right$ with $sp.right$ to form the required span between S_1 and S_3. □

Thus, span equivalence is a doubly appropriate name for $\equiv_s$: it is an equivalence of spans witnessed by a (pure) span.

Finally, we consider a third notion of equivalence, inspired by the natural bisimulation equivalence for coalgebraic bx (Abou-Saleh et al., 2015b):

Definition 4.5 (Base map). Given M-lenses $l_1 :: [S_1 \rightsquigarrow V]_M$ and $l_2 :: [S_2 \rightsquigarrow V]_M$, we say that $h : S_1 \to S_2$ is a *base map* from l_1 to l_2 if

$$
\begin{aligned}
l_1.mget \ s &= l_2.mget \ (h \ s) \\
\mathbf{do} \ \{s \leftarrow l_1.mput \ s \ v; return \ (h \ s)\} &= l_2.mput \ (h \ s) \ v \\
\mathbf{do} \ \{s \leftarrow l_1.mcreate \ v; return \ (h \ s)\} &= l_2.mcreate \ v
\end{aligned}
$$

Similarly, given two M-lens spans $sp_1 :: [A \leftsquigarrow S_1 \rightsquigarrow B]_M$ and $sp_2 :: [A \leftsquigarrow S_2 \rightsquigarrow B]_M$ we say that $h :: S_1 \to S_2$ is a base map from sp_1 to sp_2 if h is a base map from $sp_1.left$ to $sp_2.left$ and from $sp_1.right$ to $sp_2.right$. ◇

Definition 4.6 (Bisimulation equivalence). A *bisimulation* of M-lens spans $sp_1 :: [A \leftsquigarrow S_1 \rightsquigarrow B]_M$ and $sp_2 :: [A \leftsquigarrow S_2 \rightsquigarrow B]_M$ is a M-lens span $sp :: [A \leftsquigarrow R \rightsquigarrow B]_M$ where $R \subseteq S_1 \times S_2$ and fst is a base map from sp to sp_1 and snd is a base map from sp to sp_2. We write $sp_1 \equiv_b sp_2$ when there is a bisimulation of spans sp_1 and sp_2. ◇

Figure 2 illustrates the three equivalences diagrammatically.

Proposition 4.7. Each of the relations $\equiv_i$, $\equiv_s$ and $\equiv_b$ are equivalence relations on compatible spans of M-lenses and satisfy (Identity), (Assoc) and (Cong). ◇

Theorem 4.8. $sp_1 \equiv_i sp_2$ implies $sp_1 \equiv_s sp_2$, but not the converse. ◇

Proof. The forward direction is obvious; for the reverse direction, consider

$$
\begin{aligned}
&h :: Bool \rightsquigarrow () \\
&h = Lens \ (\lambda_- \to ()) \ (\lambda a \ () \to a) \ (\lambda() \to True) \\
&sp_1 :: [() \leftsquigarrow () \rightsquigarrow ()]_\mu \\
&sp_1 = Span \ idMLens \ idMLens \\
&sp_2 = (h \ ; sp_1.left, h \ ; sp_2.right)
\end{aligned}
$$

Clearly $sp_1 \equiv_s sp_2$ by definition and all three structures are well-behaved, but h is not an isomorphism: any $k :: () \rightsquigarrow Bool$ must satisfy $k.get \ () = True$ or $k.get \ () = False$, so $(h ; k).get = k.get \cdot h.get$ cannot be the identity function. □

Fig. 2. (a) Isomorphism equivalence ($\equiv_i$), (b) span equivalence ($\equiv_s$), and (c) bisimulation ($\equiv_b$) equivalence. In (c), the dotted arrows are base maps; all other arrows are (monadic) lenses.

Theorem 4.9. Given $sp_1 :: [A \leftsquigarrow S_1 \rightsquigarrow B]_M$, $sp_2 :: [A \leftsquigarrow S_2 \rightsquigarrow B]_M$, if $sp_1 \equiv_s sp_2$ then $sp_1 \equiv_b sp_2$. $\diamondsuit$

Proof. For the forward direction, it suffices to show that a single $sp_1 \curvearrowright sp_2$ step implies $sp_1 \equiv_b sp_2$, which is straightforward by taking R to be the set of pairs $\{(s_1, s_2) \mid l_1.get\ s_1 = s_2\}$, and constructing an appropriate span $sp : A \leftsquigarrow R \rightsquigarrow B$. Since bisimulation equivalence is transitive, it follows that $sp_1 \equiv_s sp_2$ implies $sp_1 \equiv_b sp_2$ as well. $\square$

In the pure case, we can also show a converse:

Theorem 4.10. Given $sp_1 :: A \leftsquigarrow S_1 \rightsquigarrow B$, $sp_2 :: A \leftsquigarrow S_2 \rightsquigarrow B$, if $sp_1 \equiv_b sp_2$ then $sp_1 \equiv_s sp_2$. $\diamondsuit$

Proof. Given R and a span $sp :: A \leftsquigarrow R \rightsquigarrow B$ constituting a bisimulation $sp_1 \equiv_b sp_2$, it suffices to construct a span $sp' = (l, r) :: S_1 \leftsquigarrow R \rightsquigarrow S_2$ satisfying $l\ ;\ sp_1.left = r\ ;\ sp_2.left$ and $l\ ;\ sp_1.right = r\ ;\ sp_2.right$. $\square$

This result is surprising because the two equivalences come from rather different perspectives. Johnson and Rosebrugh introduced a form of span equivalence, and showed that it implies bisimulation equivalence. They did not explicitly address the question of whether this implication is strict. However, there are some differences between their presentation and ours; the most important difference is the fact that we assume lenses to be equipped with a create function, while they consider lenses without create functions but sometimes consider spans of lenses to be "pointed", or equipped with designated intitial state values. Likewise, Abou-Saleh et al. (2015b) considered bisimulation equivalence for coalgebraic bx over pointed sets (i.e. sets equipped with designated initial values). It remains to be determined whether Theorem 4.10 transfers to these settings.

We leave it as an open question to determine whether $\equiv_b$ is equivalent to $\equiv_s$ for spans of monadic lenses (we conjecture that they are not), or whether an analogous result to Theorem 4.10 carries over to symmetric lenses (we conjecture that it does).

5   Conclusions

Lenses are a popular and powerful abstraction for bidirectional transformations. Although they are most often studied in their conventional, pure form, practical applications of lenses typically grapple with side-effects, including exceptions, state, and user interaction. Some recent proposals for extending lenses with monadic effects have been made; our proposal for (asymmetric) monadic lenses improves on them because M-lenses are closed under composition for any fixed monad M. Furthermore, we investigated the symmetric case, and showed that *spans* of monadic lenses are also closed under composition, while the obvious generalisation of pure symmetric lenses to incorporate monadic effects is not closed under composition. Finally, we presented three notions of equivalence for spans of monadic lenses, related them, and proved a new result: bisimulation and span equivalence coincide for pure spans of lenses. This last result is somewhat surprising, given that Johnson and Rosebrugh introduced (what we call) span equivalence to overcome perceived shortcomings in Hofmann et al.'s bisimulation-based notion of symmetric lens equivalence. Further investigation is necessary to determine whether this result generalizes.

These results illustrate the benefits of our formulation of monadic lenses and we hope they will inspire further research and appreciation of bidirectional programming with effects.

Acknowledgements

Bibliography

Abou-Saleh, F., Cheney, J., Gibbons, J., McKinna, J., Stevens, P.: Notions of bidirectional computation and entangled state monads. In: MPC. pp. 187–214. No. 9129 in LNCS, Springer-Verlag (2015a)

Abou-Saleh, F., McKinna, J., Gibbons, J.: Coalgebraic aspects of bidirectional computation. In: BX 2015. CEUR-WS, vol. 1396, pp. 15–30 (2015b)

Cheney, J., McKinna, J., Stevens, P., Gibbons, J., Abou-Saleh, F.: Entangled state monads. In: Terwilliger and Hidaka (2014)

Diviánszky, P.: LGtk API correction. `http://people.inf.elte.hu/divip/LGtk/CorrectedAPI.html` (April 2013)

Foster, J.N., Greenwald, M.B., Moore, J.T., Pierce, B.C., Schmitt, A.: Combinators for bidirectional tree transformations: A linguistic approach to the view-update problem. TOPLAS 29(3), 17 (2007)

Foster, N., Matsuda, K., Voigtländer, J.: Three complementary approaches to bidirectional programming. In: SSGIP'10. LNCS, vol. 7470, pp. 1–46. Springer (2012), `http://dx.doi.org/10.1007/978-3-642-32202-0_1`

Hofmann, M., Pierce, B.C., Wagner, D.: Symmetric lenses. In: POPL. pp. 371–384. ACM (2011)

Johnson, M., Rosebrugh, R.: Spans of lenses. In: Terwilliger and Hidaka (2014)

Jones, M.P., Duponcheel, L.: Composing monads. Tech. Rep. RR-1004, DCS, Yale (1993)

King, D.J., Wadler, P.: Combining monads. In: Proceedings of the 1992 Glasgow Workshop on Functional Programming. pp. 134–143 (1992)

Liang, S., Hudak, P., Jones, M.P.: Monad transformers and modular interpreters. In: POPL. pp. 333–343 (1995)

Pacheco, H., Hu, Z., Fischer, S.: Monadic combinators for "putback" style bidirectional programming. In: PEPM. pp. 39–50. ACM (2014), `http://doi.acm.org/10.1145/2543728.2543737`

Plotkin, G.D., Power, J.: Notions of computation determine monads. In: FOS-SACS. LNCS, vol. 2303, pp. 342–356. Springer (2002)

Stevens, P.: Bidirectional model transformations in QVT: Semantic issues and open questions. SoSyM 9(1), 7–20 (2010)

Terwilliger, J., Hidaka, S. (eds.): BX Workshop. `http://ceur-ws.org/Vol-1133/#bx` (2014)

TLCBX Project: A theory of least change for bidirectional transformations. `http://www.cs.ox.ac.uk/projects/tlcbx/`, `http://groups.inf.ed.ac.uk/bx/` (2013–2016)

Wadler, P.: Comprehending monads. Mathematical Structures in Computer Science 2(4), 461–493 (1992), `http://dx.doi.org/10.1017/S0960129500001560`

Wadler, P.: Monads for functional programming. In: Jeuring, J., Meijer, E. (eds.) Advanced Functional Programming. LNCS, vol. 925, pp. 24–52. Springer (1995)