

On Bunched Typing

Peter O’Hearn

Department of Computer Science
Queen Mary, University of London

Journal of Functional Programming, 2002


Abstract

We study a typing scheme derived from a semantic situation where a single category possesses several closed structures, corresponding to different varieties of function type. In this scheme typing contexts are trees built from two (or more) binary combining operations, or in short, *bunches*. Bunched typing and its logical counterpart, bunched implications, have arisen in joint work of the author and David Pym. The present paper gives a basic account of the type system, and then focusses on concrete models that illustrate how it may be understood in terms of resource access and sharing.

The most basic system has two context-combining operations, and the structural rules of Weakening and Contraction are allowed for one but not the other. This system includes a multiplicative, or substructural, function type \multimap alongside the usual (additive) function type \rightarrow ; it is dubbed the $\alpha\lambda$ -calculus after its binders, α for the additive binder and λ for the multiplicative, or linear, binder. We show that the features of this system are, in a sense, complementary to calculi based on linear logic; it is incompatible with an interpretation where a multiplicative function uses its argument once, but perfectly compatible with a reading based on sharing of resources. This *sharing interpretation* is derived from syntactic control of interference, a type-theoretic method of controlling sharing of storage, and we show how bunch-based management of Contraction can be used to provide a more flexible type system for interference control.

1 Introduction

In most type systems the context Γ in a typing judgement $\Gamma \vdash M : B$ is represented as a function from variables to types, or as a set or sequence of associations $x : A$ pairing identifiers with types. In this paper we study *bunched typing*, where the contexts are trees built from two or more combining operations. So, for example, we will have combining operations “ \cdot ” and “ $\cdot;$ ” which allow us to form contexts Γ, Δ and $\Gamma; \Delta$, and we will be able to nest “ \cdot ” and “ $\cdot;$ ” as in



The most promising possibility offered by bunched typing is that it gives us a flexible way to mix together calculi that treat the structural rules of Contraction, Weakening and Exchange in different ways; it is entirely possible to arrange matters so that one of the combining forms admits a structural rule when the other does not. Interest in such substructural type systems has arisen mainly as a result of work on linear logic, which has provided a novel way of understanding the structural rules in terms of *duplication and consumption* of data. We will show that bunched typing offers a different perspective, with (perhaps) surprising consequences. We show that the language we develop is even *incompatible* with a number-of-uses reading (which is characteristic of linear typing). We will argue, instead, that it should be understood in terms of *sharing*, rather than duplication: figuratively speaking, the bunch-based approach to structural rules is about who has access to what, rather than the number of times a piece of data is used.

Many variants on bunched typing are possible: the general case is to have a number of combining operations, each of which admits some combination of the structural rules, and perhaps with some interaction between the different forms of combination. The more pressing question, however, is why one might consider bunched typing at all, rather than what the general situation is. So, we will concentrate for the most part on a basic variant, which has two forms of combination, where “;” admits all of the structural rules and “,” admits Exchange only. There is no interaction between the two forms of combination. We obtain a calculus which combines simply-typed λ -calculus and a basic (multiplicative) linear λ -calculus: it is dubbed the $\alpha\lambda$ -calculus after its binders, α for the additive binder and λ for the multiplicative, or linear, binder.

Bunched typing may be understood from several theoretical viewpoints – proof theoretic, category theoretic, and semantic – and also from a specific application; sharing of storage in imperative programs. These sources serve to reinforce one another, and in next section we give an informal and leisurely survey of the perspectives offered by them. Readers who prefer a less leisurely approach can skip forward directly to the synopsis in Section 2.5.

Bunched typing and its Curry-Howard cousin, the logic BI of bunched implications, have been developed as part of joint work with David Pym; BI was introduced in a short paper by the two of us in 1999 [31]. The present paper gives a more detailed account of the type system, explaining how it arises, and various of its properties, from a particular point of view based on a “sharing interpretation” of connectives. This interpretation is suggested by Reynolds’s Syntactic Control of Interference [43], one of the main precursors of this work. Pym separately gives a more foundational treatment of both the type system and the logic [38].

Some of the material in this paper was presented, in preliminary form, in the the 1999 TLCA conference [26].

2 Routes to Bunched Typing

2.1 Sharing and Contraction

The work reported in this paper arose originally from a failed attempt to reconcile two substructural type systems, systems where the structural rule of Contraction is restricted:

$$\frac{\Gamma, x : A, y : A \vdash M : C}{\Gamma, z : A \vdash M[z/x, z/y] : C} \text{ Contraction.}$$

The background is that in 1978 Reynolds proposed syntactic control of interference (or, SCI), a type theoretic method of controlling aliasing and other shared variable interference in imperative programs [43]. In contemporary terminology, what

Reynolds used was an affine λ -calculus, where Contraction is absent and where the typing rule for function application requires that a procedure and its argument have disjoint free identifiers:

$$\frac{\Gamma \vdash M : A \multimap B \quad \Delta \vdash N : A}{\Gamma, \Delta \vdash MN : B}.$$

To understand how SCI works, it is crucial to draw a distinction between the notion of a variable, or identifier, and that of a storage cell or location that it might denote. The central statement of imperative programming is the assignment $x := e$, which overwrites the contents of a cell denoted by x . For example, a sequence of assignment statements $x := 1; y := 2$ sets (the cell denoted by) x to 1 and y to 2 if they denote different cells.

$$\begin{array}{cc} x & y \\ \boxed{1} & \boxed{2} \end{array}$$

But if x and y are aliased, which is to say denote the same cell, then the assignment to y in $x := 1; y := 2$ destroys the value placed in the cell previously by assignment to x .

$$\begin{array}{cc} x & y \\ \boxed{2} & \end{array}$$

To connect this discussion back to Contraction and function types, note that in

$$((\lambda x \lambda y . \dots x := 1; y := 2 \dots)z)z$$

if z denotes a cell, then that same cell will be passed to both x and y , resulting in aliasing. To enable this passing of the same cell to both y and z we have to have Contraction, either explicitly or as an admissible rule, in order to get two occurrences of z in an application $(Mz)z$. From this we can see that banishing Contraction abolishes aliasing, a basic example of sharing, at least in this example.

Meanwhile, in 1987 Girard introduced linear logic (or, LL), a logic that controls Contraction [16]. When its logical rules are used to type λ -terms, linear logic is evidently related to syntactic control. The typing rule for applying a linear function is

$$\frac{\Gamma \vdash M : A \multimap B \quad \Delta \vdash N : A}{\Gamma, \Delta \vdash MN : B}$$

and there is no Contraction (or Weakening). Furthermore, in SCI the rule of Contraction is not abolished altogether, but is allowed under restricted circumstances, for certain types that are labelled passive, and in linear typing Contraction is reintroduced under the control of a modality, “!”.

Thus, there is a tantalizing *formal* similarity between SCI and linear logic. But there is also a crucial *conceptual* difference: the reading of types, and resulting rationale for controlling contraction, is different in each case.

In SCI, the operative concept is *sharing*. This can be seen most clearly in the reading of the function type.

$A \multimap B$: functions that don’t share storage with their arguments.

Under this reading Contraction is understood as allowing shared access to storage. So limiting Contraction gives control over sharing.

In contrast, in LL the operative concepts are the *number of uses* of a datum, and *consumption*.

$A \multimap B$: functions that use their arguments exactly once.

The intuitive connection with consumption is that if a function uses its argument once, then it may never do so again; so one thinks of a linear function as consuming its argument. Under this reading Contraction is understood as duplicating a piece of data, rather than sharing. So limiting Contraction gives control over the number of times a piece of data can be used.

This conceptual discrepancy between LL and SCI was clear to the author in 1990 [27, 28, 29], but there was a central question left unresolved then: is the difference merely one of having two semantic interpretations of the same system (say, LL), or is a separate formal structure appropriate to each? Stated another way, does the distinction between copying and sharing have type theoretic significance, or is this just a case of having two models, where the same type theory is appropriate for both?

A hint of a possible way forward was contained in a curious property of models that had been found for SCI in the early 1990s [27, 29, 32], stated as follows in [30].

“The semantic model presented here possesses two kinds of exponential, one for the monoidal closed structure, and another, adjoint to \times for cartesian closed structure. This raises the question of whether interference control and uncontrolled Algol can coexist harmoniously in one system . . . An interesting point to note is that here the two kinds of closed structure coexist in the *same* category, so there is no need to pass to a separate category, such as a Kleisli category, to interpret the intuitionistic (i.e., Algol’s) function type.”

Given the natural structure that exists in the models, we are lead to ask: what is a typed λ -calculus corresponding to a category that admits two closed structures? This question leads us to bunched typing.

2.2 From Doubly-Closed Categories to Bunches

To see how bunches arise categorically, consider that an introduction rule for a function type typically corresponds to an adjunction. That is, a typing rule

$$\frac{\Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x . M : A \Rightarrow B}$$

corresponds to an isomorphism of maps of the corresponding shape in a closed category

$$\frac{\Gamma \otimes A \longrightarrow B}{\Gamma \longrightarrow (A \Rightarrow B)} .$$

Now, suppose that we have a *doubly closed category*, i.e., a single category equipped with two monoidal closed structures instead of only one:

$$\frac{\Gamma \wedge A \longrightarrow B}{\Gamma \longrightarrow (A \rightarrow B)} \quad \frac{\Gamma * A \longrightarrow B}{\Gamma \longrightarrow (A \multimap B)} .$$

To match this situation, we extend the syntax of typing contexts with an additional combining operation, semi-colon, which allows us to formulate introduction rules corresponding to the two adjunctions:

$$\frac{\Gamma; x : A \vdash M : B}{\Gamma \vdash \alpha x . M : A \rightarrow B} \quad \frac{\Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x . M : A \multimap B} .$$

This leads directly to the use of tree-like structures, or bunches, for typing contexts.

We will consider several variants. In each, “;” will admits structural rules of Weakening, Contraction, and Exchange and so \rightarrow will behave like the function type

in simply-typed λ -calculus or the implication in intuitionistic logic. The variants will arise by disallowing some of the structural rules for “;”. In the basic case we consider “;” will have Exchange, but not Weakening or Contraction, and this corresponds to the situation where $*$ is the tensor product of a symmetric monoidal structure.

Even at this preliminary stage, the categorical perspective allows us to crisply state the formal difference between bunch-based control over structurals and that obtained from linear logic, or linear typing. In models of linear logic two closed categories are involved, where one is often presented as a Kleisli category [48, 5, 4, 2]. For instance, in the original coherence space model there are indeed two function types, but \multimap is closed structure in the category of linear maps, while the additive \multimap , which can be represented as $!A \multimap B$, is closed for the category of stable maps. This does not provide a model of bunched typing, because in a doubly closed category we ask that the two closed structures reside in one and the same category.

Although theoretically clear, this categorical derivation of bunched structure is purely formal and does not, by itself, tell us much about the character of the resulting calculus; the view presented by categorical models is very abstract. More concretely, we have function types $A * B$ and $A \multimap B$, and we should ask: for what kinds of functions?

2.3 The Sharing Interpretation

The key to understanding the $\alpha\lambda$ -calculus is what we call the *sharing interpretation*. The background idea is of functional programming data such as functions, pairs, etc, but with an additional, intensional, notion of resources that computational entities are allowed to access. By *resource* we mean physical resources in a computer system, such as files, storage, or external devices. The reading of function types is as follows.

$A * B$: functions that have access to disjoint resources from their arguments.

$A \multimap B$: functions that have access to the same resources as their arguments.

Of course, the reading for $*$ is just the one mentioned above for SCI. The crucial point is that this can happily coexist with a direct reading of the additive function type.

Now, the bare statement of the interpretation is so direct that, at first glance, it may seem as if it must amount to the same thing as resource interpretations for other systems that control the structural rules. For, if we think of a context, roughly, as corresponding to a collection of resources, then the use of separate contexts in the elimination rule for the multiplicative function type $*$ directly expresses the disjointness mentioned in the informal interpretation, and the use of a common context in a rule for the additive corresponds to the sameness.

$$\frac{\Gamma \vdash M : A * B \quad \Delta \vdash N : A}{\Gamma, \Delta \vdash MN : B} \quad \frac{\Gamma \vdash M : A \multimap B \quad \Gamma \vdash N : A}{\Gamma \vdash M @ N : B}$$

However, there is an important point to notice: the reading places no constraint on how many times a $*$ -typed function uses its argument, it just cannot use arguments that share access to common resources. In fact, in Section 3.2 we will see a derivation of a multiplicative function that is compatible with the sharing reading, but that uses its argument twice.

Variations on the sharing interpretation are possible. For example, in a non-commutative situation there are two multiplicative function types, and the interpretation works by introducing dependency between procedure and argument.

$A \bullet B$: functions that may depend on resources of their arguments (but not vice versa).

$A \multimap B$: functions where the argument may depend on resources accessed by the function (but not vice versa).

This notion of dependency is intended to be spatial in nature. For example, “dependency” could mean that a pointer can go from the area of the store referenced by one datum to an area of store referenced by another.

2.4 Bunches in Proof Theory

The discussion so far has charted a latter-day route to bunches, emphasizing an interplay between categorical properties and resource interpretations. Historically, bunches first arose in the 1970s for completely different reasons, as a result of a problem in the proof theory of relevant logics [13]. Since then, bunches have been a standard device used by relevantists’; e.g., [3, 39, 47]. (I am grateful to David Pym for making me aware of the relevant work.) Other uses of bunched contexts include the mixed linear logic of Ruet and Fages [46], and the dependent linear type theory of Ishtiaq and Pym [21].

The most famous property of relevant logics is their denial of Weakening

$$\frac{\Gamma \vdash C}{\Gamma, A \vdash C} \text{ Weakening.}$$

This denial is done in a bid to ensure that the premisses used in a proof are actually relevant to the conclusion. The problem is that, if one simply removes Weakening from standard sequent calculi, say for intuitionistic or classical logic, then some other, intuitively reasonable, laws are blocked as well. Principal among these is the law of distribution

$$A \wedge (B \vee C) \vdash (A \wedge B) \vee (A \wedge C).$$

A standard sequent calculus proof uses Weakening in the top steps:

$$\frac{\frac{\frac{A \vdash A}{A, B \vdash A} \quad \frac{B \vdash B}{A, B \vdash B}}{A, B \vdash A \wedge B} \quad \frac{\frac{A \vdash A}{A, C \vdash A} \quad \frac{C \vdash C}{A, C \vdash C}}{A, C \vdash (A \wedge C)}}{A, B \vdash (A \wedge B) \vee (A \wedge C)} \quad \frac{\frac{A, C \vdash (A \wedge B) \vee (A \wedge C)}{A, (B \vee C) \vdash (A \wedge B) \vee (A \wedge C)}}{A \wedge (B \vee C) \vdash (A \wedge B) \vee (A \wedge C)}$$

And there is no other proof of distribution, if Weakening is simply omitted from sequent calculi.

The most important relevant logics accept distribution for semantic reasons: if one reads \wedge as “and” and \vee as “or”, then distribution *must* follow. To address the proof-theoretic problem, of how to get distribution while restricting Weakening, novel sequent calculi were formulated by Dunn and Minc (see [13]). In the notation of the present paper, the “;” form of combination admits

$$\frac{\Delta(\Gamma) \vdash C}{\Delta(\Gamma; \Gamma') \vdash C} \text{ Weakening}$$

where Weakening for “;” can occur anywhere in a bunch. Then, the rules for \vee and \wedge mention “;” but not “,”, and the proof just given for distribution goes through simply by replacing “,” with “;”.

The flexibility of the bunch-based approach to the structural rules comes about from using one form of combination to describe rules for one collection of connectives, and the other combination for different connectives. Thus, proof theoretically the relevantists’ were able to cater for the different requirements of *extensional*, or

additive, connectives such as \wedge and \vee , and *intensional*, or *multiplicative*, connectives such as fusion and relevant implication. For example, the right rules for the two conjunctions are

$$\frac{\Gamma \vdash A \quad \Delta \vdash B}{\Gamma; \Delta \vdash A \wedge B} \quad \frac{\Gamma \vdash A \quad \Delta \vdash B}{\Gamma, \Delta \vdash A * B}$$

Although these rules are identical in form, the prohibition of Weakening for “,” means that they have significantly different ramifications.

The categorical derivation of bunched structure we described earlier complements the prior discovery of the relevantists’, but also provides a more theoretically cohesive rationale for it. Whereas bunches can be used to deal with technical problems (such as distribution), they are in a sense semantically inevitable from the point of view of doubly closed categories.

2.5 Synopsis

In the next section we present a pared down form of bunched typing, where the only type constructors are for function types. This function-only fragment is simple, but also displays the most important and unusual consequences of the approach. We make a comparison with linear typing, using a number of examples.

In Section 4 we do some basic work, verifying preservation of typing under substitution and reduction. We also spell out the interpretation of the system in its categorical models.

We have discussed the sharing interpretation above, and we will use it to provide intuitive justification for some of the examples treated in Section 3.2. But the interpretation is stated informally, and it is important to know that it is consistent with bunched structure and doubly closed categories. We tackle this issue in Section 5 by presenting several models, whose description reflects the informal interpretation closely, while at the same time exhibiting doubly closed structure.

The central technique for linking the formal properties of bunched typing to sharing is the *spatial approach to possible world semantics* [44, 29, 30, 33]. In this approach, a world is viewed as corresponding to an area of memory (or, more generally, to resource), and the semantics of types and terms is parameterized by worlds. A semantic expression describing the meaning of any given term will have several occurrences of possible world parameters within it. The spatial intuition captured by this form of semantics is that when two subexpressions of a semantic expression mention different worlds, the subexpressions refer to separate areas of storage, and consequently don’t interfere. The models in Section 5 are stripped-down versions of spatial possible world models.

After describing further properties of the categorical models in Section 6, we move on in Sections 7-10 to SCI. We show how the affine variant of the calculus (which admits Weakening but not Contraction for “,”) can be used to resolve problems with jumps and recursion in the original SCI. We use bunches to provide a flexible form of interference control, where sharing constraints can be switched on and off as one moves from more local to more global contexts.

To proceed with a minimum of distraction, in presenting this work we will avoid detailed questions about the relation between syntax and semantics, questions concerning coherence, completeness, and the like. We want to use the semantics mainly as a means to expose surprising or interesting properties of the system. On one level this seems fair enough, as the calculus is very close to the models, being derived from them. But a thorough analysis is crucial; this is provided in Pym’s monograph [38]. In any case, the presentation here will be self contained, and is largely complementary to that given in [38], to which we refer for material on completeness and coherence properties of the semantics.

3 Bunched Typing

3.1 The Basic System

The basic system is motivated by models as follows.

Definition 1 A cartesian doubly closed category, or *cartesian dcc in short*, is a category equipped with two symmetric monoidal closed structures $(I, *, -*)$ and $(1, \wedge, \rightarrow)$, where $1, \wedge$ is cartesian.

We assume an unspecified collection of primitive types.

TYPES

$A ::=$	ρ	primitive types
	$A -* A$	multiplicative function type
	$A \rightarrow A$	additive function type

BUNCHES

$\Gamma ::=$	$x : A$	identifier assumption
	I	multiplicative unit
	Γ, Γ	multiplicative combination
	1	additive unit
	$\Gamma; \Gamma$	additive combination

Bunches are subject to the restriction that no identifier may occur twice in the tree. This restriction determines implicit side conditions on some of the rules below. We write $\Gamma(\Delta)$ to indicate a bunch in which Δ appears as a subtree, and $\Gamma(\Delta')$ for the similar tree where Δ' replaces Δ . To describe Contraction below, we use $i(\Gamma)$ to denote the list of identifiers encountered one after the other in an inorder traversal of the tree Γ . $\Gamma \cong \Delta$ indicates that Γ and Δ are isomorphic as trees; i.e., one can be obtained from the other by a suitable renaming of identifiers.

We won't try to come up with a more compact representation of bunches using, say, sets or sequences instead of binary operators: The real point of bunches is to let us get the α - and λ -abstractions right. We use an equivalence on trees instead of worrying about representation.

COHERENT EQUIVALENCE: $\Gamma \equiv \Gamma'$.

\equiv is the smallest equivalence relation on bunches satisfying

- 1 Commutative monoid equations for 1 and $;$
- 2 Commutative monoid equations for I and $*$,
- 3 Congruence: if $\Delta \equiv \Delta'$ then $\Gamma(\Delta) \equiv \Gamma(\Delta')$

Note that “ $;$ ” and “ $*$ ” do not distribute over one another.

TYPING JUDGEMENTS

These are of the form

$$\Gamma \vdash M : A$$

where the terms M are defined in the following rules.

IDENTITY AND STRUCTURE

$$\frac{}{x : A \vdash x : A} Id \qquad \frac{\Gamma \vdash M : A}{\Delta \vdash M : A} \equiv \text{ (where } \Delta \equiv \Gamma \text{)}$$

$$\frac{\Gamma(\Delta) \vdash M : A}{\Gamma(\Delta; \Delta') \vdash M : A} W \qquad \frac{\Gamma(\Delta; \Delta') \vdash M : A}{\Gamma(\Delta) \vdash M[i(\Delta)/i(\Delta')] : A} C \text{ (where } \Delta \cong \Delta' \text{)}$$

FUNCTIONS

$$\frac{\Gamma; x : A \vdash M : B}{\Gamma \vdash \alpha x . M : A \rightarrow B} \rightarrow I \qquad \frac{\Gamma \vdash M : A \rightarrow B \quad \Delta \vdash N : A}{\Gamma; \Delta \vdash M @ N : B} \rightarrow E$$

$$\frac{\Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x . M : A \multimap B} \multimap I \qquad \frac{\Gamma \vdash M : A \multimap B \quad \Delta \vdash N : A}{\Gamma, \Delta \vdash MN : B} \multimap E$$

In the C rule we are using a multi-nary form of substitution, where $M[i(\Delta)/i(\Delta')]$ is M with each identifier in the list $i(\Delta')$ replaced by the identifier with the same list index in $i(\Delta)$.

The rules for \multimap and \rightarrow are identical in form, but the connectives behave differently because of the structural properties of “,” and “;”. For example, a rule for additive function application that shares contexts

$$\frac{\Gamma \vdash M : A \rightarrow B \quad \Gamma \vdash N : A}{\Gamma \vdash M @ N : B}$$

is derivable using Contraction (see Section 3.4 for further discussion of this). The corresponding rule for \multimap is not derivable.

So see how these rules are working, as a warming up example notice that, given a judgement $x : A \vdash x : A$, we cannot immediately use an introduction rule to type an identity function of type $A \multimap A$ or $A \rightarrow A$. To apply an introduction rule for a function type we must have a context of the form $\Gamma, x : A$ or $\Gamma; x : A$. So we need to use coherent equivalence first.

$$\frac{\overline{x : A \vdash x : A}}{1; x : A \vdash x : A} \qquad \frac{\overline{x : A \vdash x : A}}{I, x : A \vdash x : A}$$

$$\frac{}{1 \vdash \alpha x . x : A \rightarrow A} \qquad \frac{}{I \vdash \lambda x . x : A \multimap A}$$

As a second example, using coherent equivalence we can also mimic the isomorphisms

$$[1, A \rightarrow B] \cong [A, B] \cong [I, A \multimap B]$$

of hom sets in a dcc.

$$\frac{\overline{x : A \vdash M : B}}{1; x : A \vdash M : B} \qquad \frac{\overline{x : A \vdash M : B}}{I, x : A \vdash M : B}$$

$$\frac{}{1 \vdash \alpha x . M : A \rightarrow B} \qquad \frac{}{I \vdash \lambda x . M : A \multimap B}$$

$$\frac{1 \vdash M : A \rightarrow B \quad \overline{x : A \vdash x : A}}{1; x : A \vdash M @ x : B} \qquad \frac{I \vdash M : A \multimap B \quad \overline{x : A \vdash x : A}}{I, x : A \vdash Mx : B}$$

$$\frac{}{x : A \vdash M @ x : B} \qquad \frac{}{x : A \vdash Mx : B}$$

These derivations may make the difference between \multimap and \rightarrow appear rather thin, but they only show that *closed* terms in different contexts of one function type are convertible to the other. Furthermore, in $\alpha\lambda$ the putative judgements $A \multimap B \vdash ? : A \rightarrow B$ and $A \rightarrow B \vdash ? : A \multimap B$ are not inhabited by any terms. We confirm this in Section 6.2, where we give a model (Example 13) in which there are no maps between $A \multimap B$ and $A \rightarrow B$.

To see that Weakening for “;” is not admissible in the calculus, simply note that $x : A \vdash x : A$ is derivable but $x : A, y : B \vdash x : A$ is not. To see that Contraction for “,” is not admissible, note that $(f : A \multimap B; x : A), (f' : A \multimap B; x' : A) \vdash fx' : B$ is derivable, while $f : A \multimap B; x : A \vdash fx : B$ is not. We will confirm the non-admissibility of these rules by appealing to a semantic model in Remark 9 in Section 5.1.

3.2 Unusual Examples, and Comparison to Linear Typing

The α - λ -calculus allows for multiplicative functions that use their arguments many times. For example, in the following, a variable abstracted using λ , the multiplicative abstraction, appears multiple times in the body of the term.

$$\frac{\frac{\frac{\vdots}{x; f \vdash f@x : A \rightarrow B} \quad \frac{\vdots}{x; f \vdash x : A}}{x : A; f : A \rightarrow A \rightarrow B \vdash (f@x)@x : B} C, \rightarrow E}{x : A \vdash \alpha f . (f@x)@x : ((A \rightarrow A \rightarrow B) \rightarrow B)} \rightarrow I}{I, x : A \vdash \alpha f . (f@x)@x : ((A \rightarrow A \rightarrow B) \rightarrow B)} \equiv}{I \vdash \lambda x . \alpha f . (f@x)@x : A * ((A \rightarrow A \rightarrow B) \rightarrow B)} * I$$

Here, in the key, top-pictured, step we use the admissible rule for \rightarrow elimination (or equivalently we use $\rightarrow E$ followed by Contraction, with suitable renaming of premises).

This term seems wrong if one thinks of a number-of-uses reading. But it is justified by the sharing interpretation. To see why, consider that the subterm $f@x$ is of type $A \rightarrow B$. According to the sharing interpretation, it is allowed to share with its argument, in this case x , which is why $(f@x)@x$ is reasonable. The sharing interpretation would not support an application $(fx)x$ where f had type $A * A * B$. (It is important to realize that this term really is “using” x twice. Suppose that A is in fact a function type: if f is a function that accepts two functions, and applies them to different arguments, then $(f@x)@x$ would use x in two different ways.)

Similarly, we can have a multiplicative function that doesn’t use its argument at all.

$$\frac{\frac{\frac{y : B \vdash y : B}{x : A; y : B \vdash y : B} W}{x : A \vdash (\alpha y . y) : B \rightarrow B} \rightarrow I}{I, x : A \vdash (\alpha y . y) : B \rightarrow B} \equiv}{I \vdash \lambda x . (\alpha y . y) : A * (B \rightarrow B)} * I$$

It is instructive to compare with the corresponding types in linear type theory. For the first example, the type would be $A \multimap !(A \multimap !A \multimap B) \multimap B$. In trying to derive a term we could λ -abstract on $x : A$ and function parameter f . But then, to apply (the dereliction of) f to x , we would need to convert x to something of type $!A$, and we cannot do a conversion from A to $!A$ in general. Similarly, for the type $A \multimap !B \multimap B$ we can abstract on $x : A$ and $y : !B$, but we cannot throw x away.

What is happening here can perhaps be seen more clearly by reference to Barber and Plotkin’s DILL system [2], which is a particular formulation of linear typing, that admits a direct description of \rightarrow . In Barber and Plotkin’s setup, “;” is used as a marker, between intuitionistic and linear zones, and judgements are of the form $\Gamma; \Delta \vdash M : A$. Here, intuitionistic zone Γ and linear zone Δ are simply lists or sets, and the operative rules are

DILL TYPING

$$\begin{array}{c}
\frac{}{\Gamma, x : A; _ \vdash x : A} \text{Id} - \text{Int} \\
\frac{\Gamma, x : A; \Delta \vdash M : B}{\Gamma; \Delta \vdash \alpha x . M : A \rightarrow B} \rightarrow I \\
\frac{\Gamma; \Delta, x : A \vdash M : B}{\Gamma; \Delta \vdash \lambda x . M : A \multimap B} \multimap I \\
\frac{}{\Gamma; x : A \vdash x : A} \text{Id} - \text{Lin} \\
\frac{\Gamma; \Delta \vdash M : A \rightarrow B \quad \Gamma; _ \vdash N : A}{\Gamma; \Delta \vdash M @ N : B} \rightarrow E \\
\frac{\Gamma; \Delta \vdash M : A \multimap B \quad \Gamma; \Delta' \vdash N : A}{\Gamma; \Delta, \Delta' \vdash MN : B} \multimap E
\end{array}$$

For the first of our unusual examples, the main point in DILL is that the linear zone has to be empty in the argument of an additive application $M@N$. The reason, then, that $\lambda x . \alpha f . (f @ x) @ x$ is not typable in DILL is that x would have to be in the linear zone, as it is abstracted using λ , so the elimination rule for \rightarrow could not be used. For the second example, in DILL the linear zone must be empty when we introduce an identifier from the intuitionistic zone. As a result, $\lambda x . (\alpha y . y)$ is not typable, because x would have to be in the linear zone when y is introduced in the body.

We have given terms for certain judgements in $\alpha\lambda$ that are not inhabited in linear λ -calculi. Next, we give a judgement type that is inhabited in linear type systems, and which is not in $\alpha\lambda$.

In linear logic \multimap is convertible to \rightarrow : we always have $A \multimap B \vdash !A \multimap B$, using dereliction. In DILL we can represent this with the judgement

$$\vdash; f : A \multimap B \vdash \alpha x . fx : A \rightarrow B$$

We use an empty intuitionistic context for comparison here, as to use the intuitionistic zone would be tantamount to inserting an additional “!”. We remarked above that $*$ is not convertible to \rightarrow under bunched typing, but it is useful to see why this is so. If we try to derive the corresponding judgement in $\alpha\lambda$ we get stuck:

$$\frac{\frac{???}{f : A * B; x : A \vdash fx : A \rightarrow B}}{f : A * B \vdash \alpha x . fx : A \rightarrow B} \rightarrow I$$

The ??? here cannot be filled in, because under bunched typing f and x would have to be separated using “,” in order to use the $* E$ rule to apply the multiplicative function f . That there is no term inhabiting a judgement of this form will be verified in Section 6.2, again by exhibiting a model where there are no maps of the appropriate shape.

The discussion in this section shows that, in terms of inhabitation, $\alpha\lambda$ and linear type systems are *incomparable* extensions of multiplicative and simply-typed λ -calculi. Stated logically, intuitionistic linear logic (with “!”) and BI are incomparable extensions of multiplicative intuitionistic linear logic and intuitionistic logic. More importantly, each system has a conceptual justification for the point of view it takes on these judgements where they differ. For linear typing, this is provided by the number-of-uses reading, and also by the linear type structure of domain theory. For bunched typing, this rationale is provided by the sharing interpretation.

3.3 Variants and Extensions

Many variants of the basic system are possible; the general case is to have a number of closed structures on a given category. There is no theoretical reason to stop at two, and neither is there a technical reason why one of these structures should be cartesian. But the main reason why bunches are interesting is that they give a

particularly simple way to combine a substructural system, which on its own would be rather inexpressive, with a system of full strength additive connectives. By “full strength” we mean a function type or implication that is adjoint to a cartesian product or conjunction.

We briefly mention two of the variants, one where the substructural fragment is affine, and another that includes a non-commutative fragment. We also describe rules for products.

Adding Products. We have considered the function-only fragment. The rules for products, which internalize “,” and “;”, are as follows.

$$\frac{\Gamma \vdash M : A \quad \Delta \vdash N : B}{\Gamma; \Delta \vdash \langle M, N \rangle : A \wedge B} \wedge I \quad \frac{\Gamma \vdash M : A_1 \wedge A_2}{\Gamma \vdash \pi_i M : A_i} \wedge E \text{ (where } i \text{ is 1 or 2)}$$

$$\frac{\Gamma \vdash M : A \quad \Delta \vdash N : B}{\Gamma, \Delta \vdash M * N : A * B} *I \quad \frac{\Gamma(x : A, y : B) \vdash N : C \quad \Delta \vdash M : A * B}{\Gamma(\Delta) \vdash \mathbf{let} (x, y) = M \mathbf{in} N : C} *E$$

The full system of BI also contains units for these products, and coproducts [31]. See [38] for the corresponding term calculus rules.

The Affine Variant. The affine variant arises semantically by demanding that the units I and 1 be isomorphic. The models are *affine* dcc’s, which are cartesian dcc’s where I is terminal.

The affine variant extends the basic calculus as follows.

AFFINE COHERENT EQUIVALENCE adds

$$4 \quad I \equiv 1$$

to Coherent Equivalence.

CONVERTIBILITY OF “,” TO “;”

$$\frac{\Gamma(\Delta; \Delta') \vdash M : A}{\Gamma(\Delta, \Delta') \vdash M : A} \textit{Conv}$$

Weakening for “,” is derivable in the affine variant.

$$\frac{\Gamma(\Delta) \vdash M : A}{\Gamma(\Delta, \Delta') \vdash M : A} W,$$

$A \rightarrow B$ and $A * B$ are not convertible to one another in the basic $\alpha\lambda$, but in the affine variant we can go from the former to the latter.

$$\frac{\frac{\frac{\overline{f : A \rightarrow B \vdash f : A \rightarrow B}}{f : A \rightarrow B, x : A \vdash f : A \rightarrow B} W, \quad \frac{\overline{x' : A \vdash x' : A}}{f' : A \rightarrow B, x' : A \vdash x' : A} W,}{(f : A \rightarrow B, x : A); (f' : A \rightarrow B, x' : A) \vdash f @ x' : B} \rightarrow E}{\frac{f : A \rightarrow B, x : A \vdash f @ x : B}{f : A \rightarrow B \vdash \lambda x. f @ x : A * B} C} * I$$

An intuitive explanation of this conversion can be given in terms of syntactic control of interference (see Section 8.2). If f is a function that can be applied to any argument, then we can also use it in a context where it is only applied to arguments with which it doesn’t interfere.

A Non-commutative Variant. The non-commutative variant we consider combines non-commutative, commutative, and intuitionistic fragments. A model is a single category with: a monoidal biclosed structure; a symmetric monoidal closed structure; a cartesian closed structure. (We decline to formulate an acronym.) The biclosed part means that we have two function types $\bullet-$ and $- \bullet$ satisfying the isomorphisms

$$[B, A \bullet- C] \cong [A \bullet B, C] \cong [A, B - \bullet C].$$

where \bullet is the product of a (not necessarily symmetric) monoidal structure.

Syntactically, the bunches from the basic system are augmented by adding a new unit and combination:

$$\Gamma ::= \begin{array}{l} \vdots \\ | J \\ | \Gamma \bullet \Gamma \end{array} \quad \begin{array}{l} \text{previous clauses} \\ \text{non-commutative unit} \\ \text{non-commutative combination} \end{array}$$

where for coherent equivalence we require

4 Monoid equations for J and \bullet .

Notice that there is no commutativity. We can then add rules for the left-leaning and right-leaning function types.

$$\frac{x : A \bullet \Gamma \vdash M : B}{\Gamma \vdash \lambda_{\bullet} x. M : A \bullet- B} \bullet- I \quad \frac{\Gamma \vdash M : A \bullet- B \quad \Delta \vdash N : A}{\Delta \bullet \Gamma \vdash MN : B} \bullet- E$$

$$\frac{\Gamma \bullet x : A \vdash M : B}{\Gamma \vdash \lambda_{- \bullet} x. M : A - \bullet B} - \bullet I \quad \frac{\Gamma \vdash M : A - \bullet B \quad \Delta \vdash N : A}{\Gamma \bullet \Delta \vdash MN : B} - \bullet E$$

(We will not attempt to syntactically disambiguate the various forms of application.)

The way that this system mixes its three fragments is different from Polakow and Pfenning's [37] three-zone, non-commutative variant of DILL, similarly to how DILL and $\alpha\lambda$ are different, as discussed above. It appears that models of their system can be given using three categories, with appropriate mappings between them, just as DILL arises from models based on a pair of categories.

The relationship to the system of Ruet and Fages [46] (also, [41, 40]) is less obvious. Their system uses bunches to combine two multiplicative fragments: the non-commutative and the commutative. However, it does not use bunches to treat the additives, instead relying on a modality as in linear logic. Also, their multiplicatives are classical, in that there is a (dualizing) multiplicative negation. It appears that their system can be modelled using two categories, one a ccc and the other a category possessing simultaneously a monoidal biclosed structure and a (separate) symmetric monoidal closed structure, and a dualizing object (with additional properties).

3.4 Explicit versus Implicit Structural Rules

One might have expected a different formulation of $\alpha\lambda$, where the rules of Weakening and Contraction are removed, and the rules Id and $\rightarrow E$ are replaced with

$$\frac{}{\Gamma; x : A \vdash x : A} Id, \text{ revised} \quad \frac{\Gamma \vdash M : A \rightarrow B \quad \Gamma \vdash N : A}{\Gamma \vdash M @ N : B} \rightarrow E, \text{ revised}$$

These rules are derivable in the basic system. The revised rule for identifiers follows from Id , using Weakening for “;”. And the revised rule for additive function application follows from $\rightarrow E$ and Contraction.

Let us call this the *implicit system*. The simply-typed λ -calculus is usually presented in this implicit fashion, where Weakening is implicit in the rule for identifiers, and Contraction is implicit in the elimination rule for functions. This implicit approach does not work correctly in $\alpha\lambda$, owing to interactions between multiplicatives and additives.

To see the problem, consider the judgement

$$(f : A \multimap B, x : A); z : C \vdash fx : B$$

This is derivable in the basic system: first we derive $f : A \multimap B, x : A \vdash fx : B$, and then we use Weakening. But in the implicit system it cannot be derived, because when we apply the elimination rule for \multimap we get a context of the form Γ, Δ . And, $(f : A \multimap B, x : A); z : C$ is not equivalent to any bunch of the form Γ, Δ , where f appears in Γ and x in Δ . Since we can easily derive $f : A \multimap B, x : A \vdash fx : B$ in the implicit system, this shows that the implicit system does not admit Weakening for “;”.

There is a similar problem with $*$. For, if we add the rules for $*$ to the implicit system, then Contraction is not derivable. We may readily derive

$$(f : Z \multimap A * B, z : Z); (f' : Z \multimap A * B, z' : Z) \vdash \mathbf{let}(a, b) = fz \mathbf{in} f'z'$$

but not

$$f : Z \multimap A * B, z : Z \vdash \mathbf{let}(a, b) = fz \mathbf{in} fz$$

There is the option of building Contraction into the elim rule for $*$, and Weakening into the elim rule for \multimap , but this would be treating a symptom rather than a cause. In the absence of a less unsightly solution, the formulation with explicit structural rules is to be preferred.

4 Basic Properties

Although the general idea of the $\alpha\lambda$ -calculus follows at once from doubly closed categorical structure, the detailed formulation does not. We even saw at the end of the previous section that it is entirely possible to formulate plausible-looking rules that are not quite right. In this section we examine some basic properties of $\alpha\lambda$, concentrating on on the basic system from Section 3.1. (We will not be ambitious here; this material is of the initial sanity check variety, and one could go much further.)

We first validate properties relating typing to substitution and reduction. Typing and reduction are areas where substructural type systems, which are surprisingly delicate, have encountered problems in the past, so it is appropriate that they be explored early. We then spell out how the calculus can be interpreted in any cartesian dcc.

4.1 Substitution and Reduction

Before tackling reduction, we need that each of the introduction rules for function types is reversible. This is a property we expect given the isomorphisms

$$[A * B, C] \cong [A, B \multimap C] \quad [A \wedge B, C] \cong [A, B \multimap C]$$

in a dcc.

Lemma 2 (Reversibility) *The inverses of $\multimap I$ and $\multimap E$ are admissible rules:*

$$\frac{\Gamma \vdash \alpha x. M : A \multimap B}{\Gamma; x : A \vdash M : B} \quad \frac{\Gamma \vdash \lambda x. M : A \multimap B}{\Gamma, x : A \vdash M : B}$$

Proof Sketch: In each case a derivation of the premiss must end with an application of the corresponding intro rule, followed by a sequence of applications of the structural rules (W, C, \equiv). The proof goes by induction on the length of this last part of the derivation; all cases are straightforward. ■

Recall here that “admissibility” means that if you can infer the judgement above the line, then you can also infer the judgement below. This does not, however, mean that there is a generic derivation from one to the other, and admissibility is a property that is not preserved under extensions to a language.

The substitution lemma is formulated for identifiers appearing arbitrarily deeply in a bunch.

Lemma 3 (Substitution Lemma) *The following is an admissible rule.*

$$\frac{\Gamma(x : A) \vdash M : B \quad \Delta \vdash N : A}{\Gamma(\Delta) \vdash M[N/x] : B}$$

Proof Sketch: As usual, a multi-nary version is proven in order to get a strong enough induction hypothesis:

$$\frac{\Gamma(x_1 : A_1 \mid \cdots \mid x_m : A_m) \vdash M : B \quad \Delta_1 \vdash N_1 : A_1, \dots, \Delta_m \vdash N_m : A_m}{\Gamma(\Delta_1 \mid \cdots \mid \Delta_m) \vdash M[N_1/x_1, \dots, N_m/x_m] : B}$$

where $\Gamma(\Delta_1 \mid \cdots \mid \Delta_m)$ indicates a bunch with multiple distinct sub-bunches.

The proof goes by induction on the derivation of $\Gamma(x_1 : A_1 \mid \cdots \mid x_m : A_m) \vdash M : B$, where the multi-nary aspect is used to deal with the case of Contraction. ■

The two kinds of function in $\alpha\lambda$ come associated with the usual reductions.

β -reductions	η -reductions
$(\alpha x. M)@N \triangleright M[N/x]$	$(\alpha x. M @ x) \triangleright M \ (x \notin \text{free}(M))$
$(\lambda x. M)N \triangleright M[N/x]$	$(\lambda x. Mx) \triangleright M \ (x \notin \text{free}(M))$

Proposition 4 (Subject Reduction) *If $\Gamma \vdash M : A$ and $M \triangleright N$ then $\Gamma \vdash N : A$.*

Proof: To prove the β case for λ , note that a derivation of $\Gamma \vdash (\lambda x. M)N : A$ must end in a use of $\rightarrow E$, followed by a number of applications of structural rules. The proof goes by induction on the length of this sequence.

In the base case we have

$$\frac{\Gamma \vdash (\lambda x. M) : A \rightarrow B \quad \Delta \vdash N : A}{\Gamma, \Delta \vdash (\lambda x. M)N : B}$$

as the last rule in the derivation. By the Reversibility Lemma we have $\Gamma, x : A \vdash M : B$ and we can then use the instance

$$\frac{\Gamma, x : A \vdash M : B \quad \Delta \vdash N : A}{\Gamma, \Delta \vdash M[N/x] : B}$$

of the substitution lemma.

All other cases (C, \equiv, W) are straightforward, with C using a typical commutativity property of substitution.

The β law for \rightarrow is similar, except that to handle the base case

$$\frac{\Gamma \vdash (\alpha x. M) : A \rightarrow B \quad \Delta \vdash N : A}{\Gamma; \Delta \vdash (\alpha x. M)@N : B}$$

we use the instance

$$\frac{\Gamma; x : A \vdash M : B \quad \Delta \vdash N : A}{\Gamma; \Delta \vdash M[N/x] : B}$$

of the Substitution Lemma. Notice how the single, nested, formulation of the lemma covers both “;” and “,” cases.

The proofs for η laws are straightforward, also relying on Reversibility. ■

Finally, we remark that the strong normalization theorem for $\alpha\lambda$ does not require any work on our part, because it follows at once from the corresponding result for simply-typed λ -calculus.

Proposition 5 (Strong Normalization) *There are no infinite reduction sequences, starting from any typable term.*

Proof Sketch: We can define a mapping of $\alpha\lambda$ into the simply-typed λ -calculus which sends both \rightarrow and \ast to the function type \rightarrow of λ -calculus. Any reduction in $\alpha\lambda$ then induces a reduction in λ -calculus, which is preserved by the mapping. There can therefore be no infinite reduction sequences in $\alpha\lambda$, or this would contradict the strong normalization theorem of typed λ -calculus. ■

4.2 Semantic Interpretation

Suppose we are given a cartesian dcc. An interpretation of $\alpha\lambda$ specifies an object $\llbracket \rho \rrbracket$ for each primitive type, which then extends to all types using the closed structures. A bunch is interpreted by mapping “;” to \ast , “,” to \wedge , and similarly for the units. Then, for any proof π of a judgement $\Gamma \vdash M : A$ we can define a map

$$\llbracket M \rrbracket_\pi : \llbracket \Gamma \rrbracket \longrightarrow \llbracket A \rrbracket$$

by induction on π .

To describe the interpretation, if $\Gamma(\cdot)$ is a bunch with a hole, then it determines a functor $\llbracket \Gamma \rrbracket(\cdot)$ on the dcc in question. When we write $\llbracket \Gamma \rrbracket(\llbracket \Delta \rrbracket)$ this indicates the action of the functor on objects, and the occurrences of $\llbracket \Gamma \rrbracket(\text{contr})$ and $\llbracket \Gamma \rrbracket(\pi_1)$ in the rules for Weakening and Contraction use the morphism part.

SEMANTIC INTERPRETATION

$$\begin{array}{c} \frac{}{\text{id} : \llbracket A \rrbracket \longrightarrow \llbracket A \rrbracket} \text{Id} \quad \frac{m : \llbracket \Delta \rrbracket \longrightarrow \llbracket A \rrbracket}{i; m : \llbracket \Delta \rrbracket \longrightarrow \llbracket A \rrbracket} \equiv (\text{canonical } i : \llbracket \Delta \rrbracket \longrightarrow \llbracket \Gamma \rrbracket) \\ \\ \frac{m : \llbracket \Gamma \rrbracket(\llbracket \Delta \rrbracket) \longrightarrow \llbracket A \rrbracket}{(\llbracket \Gamma \rrbracket(\pi_1)); m : \llbracket \Gamma \rrbracket(\llbracket \Delta \rrbracket \wedge \llbracket \Delta' \rrbracket) \longrightarrow \llbracket A \rrbracket} W \\ \\ \frac{(m : \llbracket \Gamma \rrbracket(\llbracket \Delta \rrbracket \wedge \llbracket \Delta' \rrbracket) \longrightarrow \llbracket A \rrbracket)}{(\llbracket \Gamma \rrbracket(\text{contr})); m : \llbracket \Gamma \rrbracket(\llbracket \Delta \rrbracket) \longrightarrow \llbracket A \rrbracket} C \text{ (where } \Delta \cong \Delta') \\ \\ \frac{m : \llbracket \Gamma \rrbracket \wedge \llbracket A \rrbracket \longrightarrow \llbracket B \rrbracket}{m^\ast : \llbracket \Gamma \rrbracket \longrightarrow \llbracket A \rrbracket \rightarrow \llbracket B \rrbracket} \rightarrow I \quad \frac{m : \llbracket \Gamma \rrbracket \longrightarrow \llbracket A \rrbracket \rightarrow \llbracket B \rrbracket \quad n : \llbracket \Gamma \rrbracket \longrightarrow \llbracket A \rrbracket}{m \wedge n; \text{app}_\rightarrow : \llbracket \Gamma \rrbracket \longrightarrow \llbracket B \rrbracket} \rightarrow E \\ \\ \frac{m : \llbracket \Gamma \rrbracket \ast \llbracket A \rrbracket \longrightarrow \llbracket B \rrbracket}{m^\circ : \llbracket \Gamma \rrbracket \longrightarrow \llbracket A \rrbracket \ast \llbracket B \rrbracket} \ast I \quad \frac{m : \llbracket \Gamma \rrbracket \longrightarrow \llbracket A \rrbracket \ast \llbracket B \rrbracket \quad n : \llbracket \Delta \rrbracket \longrightarrow \llbracket A \rrbracket}{(m \ast n); \text{app}_\ast : \llbracket \Gamma \rrbracket \ast \llbracket \Delta \rrbracket \longrightarrow \llbracket B \rrbracket} \ast E \end{array}$$

In the interpretations of Weakening and Contraction, π_1 is the second projection and contr is the doubling map $\langle \pi_1, \pi_2 \rangle : \llbracket \Delta \rrbracket \longrightarrow \llbracket \Delta \rrbracket \wedge \llbracket \Delta \rrbracket$ associated with cartesian structure. Note that $\llbracket \Delta \rrbracket$ and $\llbracket \Delta' \rrbracket$ are actually *equal* when Δ and Δ' are isomorphic; so this semantic clause is type correct.

For the other rules, i is a canonical isomorphism; $(\cdot)^*$ and $(\cdot)^\circ$ are the isomorphisms of hom sets obtained from the adjunctions for \rightarrow and \multimap ; app_{\rightarrow} and app_{\multimap} are the application maps obtained from the adjunctions; $\langle m, n \rangle$ is the pairing for the cartesian product; $m * n$ is the action of the $*$ functor on morphisms.

There are two coherence issues worth mentioning here. The first concerns the isomorphism $i : \llbracket \Delta \rrbracket \longrightarrow \llbracket \Gamma \rrbracket$. We intend that this is obtained from a proof that $\Delta \equiv \Gamma$ using commutative monoid laws: any such proof determines an isomorphism, using the coherent isomorphisms of symmetric monoidal categories [24]. We claim that applications of symmetry morphisms are explicitly disambiguated by the use of different identifiers for different types appearing in bunches, so that the morphism i is unique. (This appears to follow from the usual coherence results for symmetric monoidal categories, but we will not carry out a detailed proof.) In any case, this coherence issue can often be sidestepped in specific models by using an interpretation where equivalent bunches are semantically equal; an example will be given in Section 9.2.

The second coherence issue has to do with the order of application of the rules: there can be different proofs π and π' of the same judgement $\Gamma \vdash M : A$, for instance when a structural rule is applied before or after an elimination or intro rule. In this situation we would like to know in such a situation that the two proofs determine the same map $\llbracket M \rrbracket_\pi = \llbracket M \rrbracket_{\pi'}$.

A detailed study of the connection between syntax and semantics would involve a careful proof of coherence, together with soundness and completeness results connecting syntactic equality with equality in the models. We avoid this here, and instead refer the reader to Pym’s monograph for more information [38].

However, since we have established that any proof of a judgement $\Gamma \vdash M : A$ determines a map from $\llbracket \Gamma \rrbracket$ to $\llbracket A \rrbracket$ in a cartesian dcc, we can use the following fact with confidence (even prior to coherence or completeness issues).

Lemma 6 (Inhabitation Lemma) *If there exists a cartesian dcc and interpretation of primitive types in which the hom set $\llbracket \llbracket \Gamma \rrbracket, \llbracket A \rrbracket \rrbracket$ is empty, then there can be no M with $\Gamma \vdash M : A$ in $\alpha\lambda$.*

5 Models for the Sharing Interpretation

Thus far the sharing interpretation has been stated informally; in this section we give three models corresponding to it. In the models for the basic and non-commutative languages, we will give some somewhat lengthy examples, which show in some detail how the interpretations of types work. Our reason for doing this is that we are claiming that the sharing interpretation gives a consistent reading of bunched typing; to understand the sense in which it does, we need to do more than merely mention the models in passing. At the same time, though somewhat lengthy, these examples stop well short of being “applications”.

An affine model is described briefly. A variant of it is developed more fully in Sections 7-10, when we consider SCI.

The models in this section are instances of an abstract construction due to Day [12], which shows how to obtain a monoidal biclosed structure on the functor category $\mathbf{Set}^{\mathcal{C}^{op}}$, starting from a “promonoidal” structure on \mathcal{C} . Combined with the standard fact that $\mathbf{Set}^{\mathcal{C}^{op}}$ is cartesian closed, this construction gives us a host of models for bunched typing.

We present the models here in an elementary fashion, using direct descriptions of the dcc structure; it is not even necessary to know (or remember) the definitions of cartesian closed structure in functor categories. In two of the three cases the parametrizing category will just be a set, or a discrete category, and $\mathbf{Set}^{\mathcal{C}^{op}} = \mathbf{Set}^{\mathcal{C}}$

thus a product category. In the other case \mathcal{C} is a poset. Our intention in doing this is to use very concrete, and even simple-minded, models, in order to make the connection back to sharing in a straightforward way.

5.1 Resource Separation: The Basic Disjointness Model

Let \mathcal{F} denote the collection of finite subsets of a given infinite set Loc . We think of an element $X \in \mathcal{F}$ as a possible world, that determines a finite collection of resources or, more concretely, locations in computer memory. We will use the product category $\mathbf{Set}^{\mathcal{F}}$ as a model of $\alpha\lambda$. For an object A and element $a \in AX$, we regard a as a computational entity that has access to X . The model in this section is for the basic version of $\alpha\lambda$ described in Section 3.1; we refer to it as the “basic disjointness model”.

The crucial operation on worlds is disjoint combination. It is a partial operation, because we only combine those finite sets that are disjoint; intuitively, $X * Y$ indicates separation, where the component worlds X and Y determine distinct resources.

$$\begin{aligned} X * Y &= X \cup Y, \text{ when } X \cap Y = \{\}; \\ X * Y &= \textit{undefined}, \text{ when } X \cap Y \neq \{\}. \end{aligned}$$

This definition makes $\mathcal{F}, \{\}, *$ a partial commutative monoid. This means that the commutative monoid laws hold up to an equivalence $e \simeq e'$ on expressions built using $*$ which says that both sides are defined and equal or both undefined.

The cartesian closed structure in $\mathbf{Set}^{\mathcal{F}}$ is determined pointwise.

$$\begin{aligned} (A \rightarrow B)X &= AX \Rightarrow BX \\ (A \wedge B)X &= Ax \times BX \\ 1X &= \{a\} \end{aligned}$$

Here, \Rightarrow and \times are function space and product in \mathbf{Set} .

Notice how the pointwise definition of \rightarrow corresponds closely to the informal reading in the sharing interpretation, where an additive function and its argument have access to the same resources. The additive function type has a strongly local character, where an application of a function stays located at a given world.

To describe the multiplicative function type, we use a multiplicative form of indexed product. If $A(X, Y)$ is an expression (in the metalanguage) containing parameters X, Y for distinct worlds, then

$$\Pi_Y \# A(X, Y)$$

is the product, indexed over finite sets disjoint from X . To be precise, an element is a function that accepts a world Y disjoint from X and produces an element of $A(X, Y)$. We often refer to the Y parameter as being “fresh”, to briefly indicate its disjointness from X .

The multiplicative function type then quantifies over fresh worlds.

$$(A * B)X = \Pi_Y \# A(Y) \Rightarrow B(X * Y)$$

Because of the disjointness requirement on $\Pi_Y \#$, the $X * Y$ component in this expression is always defined.

In this definition the absence of X in $A(Y)$ mirrors the informal description of multiplicative functions as disjoint from their arguments. An element $p \in (A * B)X$ accepts fresh world Y and element $a \in AY$ as arguments, and produces $p[Y]a \in B(X * Y)$: The “resources” for p are X , while those for a are Y , and these are separate in the result type by virtue of their positions in the combined world $X * Y$.

This illustrates the spatial way of reading the semantic expressions referred to in the Introduction.

We give an example to further illustrate the sharing aspect. Consider the inclusion $L : \mathcal{F} \rightarrow \mathbf{Set}$. For each finite set X , we think of $LX = X$ as a set of names, or locations. Let N be the constant object, which is the natural numbers at every component, and define

$$S = L \rightarrow (1 \vee (N \wedge L))$$

where \vee is the pointwise-defined coproduct. Because of the pointwise definition of \rightarrow we have that $SX = X \Rightarrow \{a\} + (N \times X)$. We regard an element $s \in SX$ as a representation of a portion of a computer store, where each $x \in X$ is a pointer to a linked list (possibly with loops).

Now consider any function $f \in ((S \wedge L) \rightarrow ((S \wedge L) \multimap S))X$. f accepts $(s, x) \in SX \times X$ and $(s', y) \in SY \times Y$, for fresh finite set Y , as arguments, and produces a state in $S(X * Y)$ as a final result. From the point of view of $S(X * Y)$, there is no overlap between x and y , or between the other pointers in the list pointed to by x and those pointed to by y . Thus, we can view f as a procedure that accepts two linked lists as arguments, with the proviso that the two input lists are defined using disjoint collections of pointers. This kind of proviso is often required in the statement of correctness of an algorithm that, say, removes the elements of one list that appear in the other.

On the other hand, consider the type $L \rightarrow (L \multimap (S \rightarrow S))$. A function of this type would accept two pointers to linked lists as arguments, and the two pointer arguments would again have to be distinct, but now they could point to lists that overlap in the store.

No particular practical significance is claimed for this example; it is offered just as an illustration of how \multimap and \rightarrow can express sharing properties.

In this model the multiplicative unit is I where $I(\{\}) = \{*\}$, and $I(X) = \{\}$ for all other X .

Before defining $*$ it is useful to observe that a multi-map characterization of maps out of $A * B$ is forced by the definition of \multimap . That is, if we are to have the isomorphism $\mathbf{Set}^{\mathcal{F}}[A * B, C] \cong \mathbf{Set}^{\mathcal{F}}[A, B \multimap C]$, then we must obtain the following.

*Maps $p : A * B \rightarrow C$ out of a tensor are in bijection with families of functions*

$$\bar{p}[X][Y] : AX \times BY \rightarrow C(X * Y) ,$$

indexed over disjoint finite sets X and Y .

The idea in terms of sharing is that the components of $*$ are assigned different resources (this is in line with the form of semantics devised by Reynolds for syntactic control of interference [33]).

Given this characterization, we can give a simple description of the application map $app_{\multimap} : (A \multimap B) * A \rightarrow B$. It is nothing other than

$$app_{\multimap} [X][Y]\langle p, a \rangle = p[Y]a.$$

Also, the exponential transpose, which takes $m : A * B \rightarrow C$ to $m^\circ : A \rightarrow (B \multimap C)$ is just

$$m^\circ[X]a[Y]b = m[X][Y]\langle a, b \rangle.$$

The actual definition of $*$ is straightforward:

$$(A * B)X = \{Y, Z, a \in AY, b \in BZ \mid Y * Z = X\}$$

Here, the condition $Y * Z = X$ requires that $Y * Z$ be defined; so an element of $(A * B)X$ consists of a splitting of the current world, together with two entities of types A and B having access to the respective components of the splitting.

The application and exponentiation maps for \rightarrow are immediate, given the pointwise definition of \rightarrow . All told, we have all of the structure necessary to model $\alpha\lambda$.

Proposition 7 $\mathbf{Set}^{\mathcal{F}}$ is a cartesian dcc.

From the point of view of this model the judgement

$$I \vdash \lambda x . \alpha f . (f @ x) @ x : A * ((A \rightarrow A \rightarrow B) \rightarrow B)$$

from Section 3.2 is utterly unsurprising. It determines an element $p \in A * ((A \rightarrow A \rightarrow B) \rightarrow B)\{\}$ (where we indulge in a confusion between types and objects in $\mathbf{Set}^{\mathcal{F}}$). This function p accepts a fresh world X and $a \in AX$, and produces a function $p[X]a \in ((A \rightarrow A \rightarrow B) \rightarrow B)X$. By the pointwise definition of \rightarrow , this is a function of type $(AX \Rightarrow AX \Rightarrow BX) \Rightarrow BX$ in \mathbf{Set} , and it is the expected function that maps f to $(fa)a$.

Remark 8 A broadly similar development can also be carried out in $\mathbf{Set}^{\mathcal{B}}$, where \mathcal{B} is the category of finite sets and bijections. In this category, we model disjointness using the (total) operation $+$ on finite sets which takes disjoint union by tagging its components. For this to work, however, the use of non-identity bijections is crucial: it gives rise to associativity, unity and symmetry isomorphisms that make $(\{\}, +)$ a symmetric monoidal structure on \mathcal{B} . We could not use \mathcal{F} with $+$, because \mathcal{F} is a discrete category, and does not have the morphisms needed to make $(\{\}, +)$ monoidal on that category.

Remark 9 It is important to see that there is no hidden Weakening or Contraction for “,” lurking in the examples of terms that use their arguments two or zero times. In fact, we can see that these rules are absent $\mathbf{Set}^{\mathcal{F}}$ in a very strong sense; there are not even any candidate maps of the required types to model them, let alone maps with the proper properties.

To model Contraction we would need maps of shape $A \rightarrow A * A$. But there are no maps $L \rightarrow L * L$, where L is the inclusion from \mathcal{F} to \mathbf{Set} . To see why, given $a \in L\{a\}$ we would have to produce an element in $(L * L)\{a\}$, but this set is empty. The reason is that if $X * Y = \{a\}$ then either X or Y must be $\{\}$ and so, by the definition of $*$, a tuple in $(L * L)\{a\}$ would have to identify an element of $L\{\}$. But $L\{\}$ is empty so there can be no such tuple.

To model Weakening, we would need maps $A \rightarrow I$, for all A . But there are no maps $1 \rightarrow I$.

These remarks confirm the non-admissibility of Weakening and Contraction for “,” referred to in Section 3.1.

5.2 An Affine Model

Strictly speaking, the sharing interpretation is stated as for the basic version of $\alpha\lambda$. The reading for the *affine* variant, which admits Weakening for the multiplicative combination “,”, is obtained by changing the interpretation of the additive function type to say that functions *may* share resources with their arguments.

Let \mathcal{FS} (for \mathcal{F} -sub) denote the poset whose elements are the same as \mathcal{F} , and whose objects are ordered by subset inclusion. We will use the functor category $\mathbf{Set}^{\mathcal{FS}}$. The $*$ operation on worlds from the previous subsection satisfies the following monotonicity property: if $X * Y$ is defined, and $Z \sqsubseteq Y$, then $Z * Y$ is

defined and $X * Z \subseteq Y$. In this sense, $*$ gives \mathcal{FS} the structure of an *ordered* partial commutative monoid.

The first modification that needs to be made to the basic disjointness model is in the additive function type. There is a general formula for this type in a functor category, but it will be useful to adopt a special representation, which is tuned to the properties of \mathcal{FS} .

Specifically, if $X \subseteq Z$ then there is a unique Y such that $X * Y = Z$. This allows us to quantify only over sets disjoint from the current world when defining \rightarrow , instead of over all supersets of it.

$$(A \rightarrow B)(X) = \Pi_Y \# A(X * Y) \Rightarrow B(X * Y), \text{ natural in } Y.$$

Here, the notation $\Pi_Y \#$ is the multiplicative indexed product defined in the last subsection. Naturality in Y means that the equality $p[Z](A(X * Y \subseteq X * Z)a) = A(X * Y \subseteq X * Z)(p[Y]a)$ holds when $Z \supseteq Y$ is disjoint from X .

We think of the presence of X in the argument type $A(X * Y)$ as indicating the possibility of sharing between procedure and argument. Notice, however, that by use of subset inclusion an element $p \in (A \rightarrow B)X$ can actually be applied to an element that “comes from” a world where X is not present, such as $p[Y]a$ where $a = A(Y \subseteq X * Y)a'$ for some $a' \in AY$. So we regard the formula for the additive type as saying that the procedure *may* share with its argument.

The definition of the morphism part of $A \rightarrow B$ is essentially as in Section 9.1.

The multiplicative function type has the same definition as before, with naturality added:

$$(A * B)X = \Pi_Y \# A(Y) \Rightarrow B(X * Y), \text{ natural in } Y.$$

Recall the judgement $f : A \rightarrow B \vdash \lambda x. f @ x : A * B$, that converts an additive to a multiplicative function in the affine language. In this model the conversion takes a natural transformation $A(X + -) \rightarrow B(X + -)$ and composes on the left with the map $A \rightarrow A(X + -)$ that sends $a \in AY$ to $A(\text{inr})a \in A(X + Y)$, where inr is the right injection. Here, an additive function in world X is applied to an argument $a \in AY$ that doesn't happen to depend on X .

To complete the definition of the model we must define $A * B$ for functors A and B . First we set up a preorder.

- Elements: tuples $\langle Y, Z, a \in AY, b \in BZ \rangle$ where $Y * Z$ is defined and $Y * Z \subseteq X$.
- Order: $\langle Y, Z, a \in AY, b \in BZ \rangle \sqsubseteq \langle Y', Z', a' \in AY', b' \in BZ' \rangle$ if $Y \subseteq Y'$, $Z \subseteq Z'$, $a' = A(Y \subseteq Y')a$ and $b' = B(Z \subseteq Z')b$.

Two tuples are then declared equivalent if they have a common parent under this order. Writing $[\cdot]$ for equivalence classes,

$$(A * B)X = \{[\langle Y, Z, a \in AY, b \in BZ \rangle] \mid Y * Z \text{ is defined and } Y * Z \subseteq X\}.$$

This definition is complex, but is just an instance of Day's tensor product, which can be described compactly using a coend formula:

$$(A * B)X = \int^{Y, Z} AY \times BZ \times C[X, Y * Z].$$

The unit of $*$ is the terminal object 1 , where $1X = \{a\}$ is constantly the one-point set, and so the model validates Weakening.

Fact 10 $\text{Set}^{\mathcal{FS}}$ is an affine dcc.

Comparing to the work in the previous section, we can try to use the inclusion functor $L : \mathcal{T} \rightarrow \mathbf{Set}$ as a variant on the functor used to illustrate the sharing interpretation. But L has something of a different character in the affine model. It would not be as reasonable to think of $s \in (L \rightarrow (1 \vee (N \wedge L)))X$ as a state, because s would have to accept other worlds Y , and potentially $y \in LY$, as arguments. So the development above, for the basic disjointness model, does not carry through well to the affine case. However, a thorough account of the sharing aspect of a variant of the affine model is given later, in the context of SCI.

5.3 A Non-commutative Model

In the non-commutative model the commutative multiplicatives $*$ and \multimap will continue to express absence of sharing. To this we add non-commutative operators \bullet , $\bullet\multimap$ and $\multimap\bullet$ which express a directional form of sharing.

Let \mathcal{W} denote the set of binary relations $X \subseteq Loc \times Loc$, for a fixed set Loc . A relation describes a constraint on the shape of the computer store, where $\ell X \ell'$ means that there can be a pointer from ℓ to ℓ' .

The non-commutative product of worlds, $X \bullet Y$, will describe a situation where there can be pointers from (the domain of) Y back into X , but not vice versa. To describe this we make two definitions:

- $domX = \{\ell \mid \exists \ell'. \ell X \ell'\}$ is the domain of relation X ;
- $X \gg Y$ holds just if $domX \cap domY = \{\}$ and $\ell X \ell' \Rightarrow \ell' \notin domY$.

When $X \gg Y$ holds, the union of relations X and Y allows ℓ and ℓ' to be related, where $\ell \in domY$ and $\ell' \in domX$, but not the converse. This leads us to

- $X \bullet Y = X \cup Y$, when $X \gg Y$;
- $X \bullet Y = \text{undefined}$, when $\neg(X \gg Y)$.

There are two natural choices for the commutative product $*$.

1. *Shallow Non-interference*: $X * Y = X \cup Y$, when $domX \cap domY = \{\}$.
2. *Deep Non-interference*: $X * Y = X \cup Y$, when $(domX \cup codX) \cap (domY \cup codY) = \{\}$.

In either case (with the operation undefined in other cases), we obtain that $(\mathcal{W}, I, *)$ is a partial commutative monoid. (We resist the temptation to formulate a language with two separate commutative monoidal fragments.)

The definitions of \multimap and $*$ in $\mathbf{Set}^{\mathcal{W}}$ are similar to the ones in the basic disjointness model, and omitted.

To describe the non-commutative function types, we first define a non-commutative, multiplicative indexed product in the metalanguage. If $A(X, Y)$ is an expression containing parameters X, Y for worlds, where $X \gg Y$, then

$$\Pi_Y \gg A(X, Y)$$

is the product, indexed over worlds \gg -related to Y . That is, an element is a function that accepts a world Y where $X \gg Y$ holds, and produces an element of $A(X, Y)$. Similarly, an element of

$$\Pi_X \ll A(X, Y)$$

is a function that accepts a world X where $X \gg Y$ and produces an element of $A(X, Y)$.

Then the two function types are:

$$\begin{aligned} (A \bullet - B)X &= \Pi_Y \gg A(Y) \Rightarrow B(Y \bullet X) \\ (A - \bullet B)X &= \Pi_Y \ll A(Y) \Rightarrow B(X \bullet Y). \end{aligned}$$

For \bullet ,

$$(A \bullet B)X = \{Y, Z, a \in AY, b \in BZ \mid Y \bullet Z = X\}.$$

As in the basic disjointness model there multi-map characterization of maps out of $A \bullet B$, except that the characterization for \bullet works with pairs of world subject to the constraint that $X \gg Y$; the application and transpose maps are also straightforward. So we state:

Proposition 11 $\mathbf{Set}^{\mathcal{W}}$ is a model of the non-commutative variant: It is

1. monoidal biclosed $(\bullet, \bullet -, \bullet \bullet)$,
2. symmetric monoidal closed $(*, -*)$, and
3. cartesian closed (\wedge, \rightarrow) .

(The cartesian closed structure is inherited pointwise from \mathbf{Set} .)

The connection between the sharing interpretation and the definitions of $-*$ and \rightarrow established in the discussion of the disjointness model go through just as well for the model of this section; so we concentrate on the directionality of sharing expressible using the non-commutative operators. (We stress that it is important that this reading does not invalidate that for the other connectives, especially the reading for the additive \rightarrow .)

First, we define states similarly as in the disjointness model, but with two differences.

$$SX = \{s \in \text{dom}X \rightarrow (\{a\} \vee (N \wedge \text{Loc})) \mid sl = \langle n, \ell' \rangle \text{ implies } \ell X \ell'\}.$$

The first difference is the use of the constraint $\ell X \ell'$ determined by the relation: the store must be compatible with the given store shape.

The second difference is that locations in the domain and range of a state s are treated differently, because the former must be in $\text{dom}X$ while the latter are taken from all locations. We think of $\text{dom}X$ as the collection of known or active locations at a given world. The use of Loc enables a situation where one location points to another, where that other's contents is unknown: we have dangling references. For example, in the world $\{\langle \ell, \ell' \rangle\}$ the store $[\ell \mapsto \langle 3, \ell' \rangle]$ is valid, where we do not know what ℓ' points to.

Dangling references play a crucial role in the treatment of \bullet . For example, in the composite world $Z = Y \bullet \{\langle \ell, \ell' \rangle\}$ we use the dangling reference ℓ' in $\{\langle \ell, \ell' \rangle\}$ to “reach back” into Y . In this composite world there cannot be any pairs of the form $\langle \ell'', \ell \rangle$, so in a state $s \in SZ$ there cannot be any pointers into ℓ . In particular, ℓ can be the head of a linked list, but not a non-head node.

Now we want to see how the non-commutative function types work in this model. Recall the sharing interpretation for $- \bullet$:

$A - \bullet B$: functions where the argument may depend on resources accessed by the function (but not vice versa).

To make this concrete we work with an object of cells, as well as with states. Since we regard the domain of a relation as the collection of known cells, we set $\mathbf{cell}X = \text{dom}X$ (and $\mathbf{cell}\perp = \{\}$).

We are going to describe a map

$$\text{stack-alloc} : (\mathbf{cell} - \bullet (S \rightarrow S)) \longrightarrow (S \rightarrow S)$$

where $\text{stack-alloc}(\lambda_{-\bullet} x.C)$

- allocates a new location, initializes it to a , and binds it to x ;
- executes C ;
- de-allocates the new location on block exit.

This may seem surprising, since in the presence of pointers stack allocation is not generally possible. For, if you allocate a pointer, and make some other pointer point to it, then the new pointer cannot be deallocated without creating a dangling reference. (And, dangling references are a rich source of program errors.)

Ultimately, this works because the right-leaning function type allows us to express a kind of dependency: x may point to other, older, pointers but not the reverse. Thus, deallocating x on block exit will not create any dangling pointers. And, if we know that C does not itself create a dangling pointer – say, if there are no facilities for freeing or disposing a pointer in the language – then this form of stack allocation of pointers will be completely safe.

To nail this down, first note that a function $f \in (\mathbf{cell} \multimap (S \rightarrow S))X$ will accept a world Y and cell ℓ and then give us back

$$f[Y]\ell : S(X \bullet Y) \longrightarrow S(X \bullet Y)$$

What we need to do is choose ℓ to be a new location: then the definition of \bullet will ensure that ℓ cannot be pointed to from X . We also need to choose the relational constraint for Y , and for this it makes sense to let ℓ point to anything in X .

To formalize this discussion we require:

- a location $\mathit{newloc}(X) \notin \mathit{dom}X$.

This location can be chosen using some enumeration of Loc . Next, we define

$$\mathit{newworld}(X) = \{(\mathit{newloc}(X), \ell) \mid \ell \in \mathit{dom}X \vee \ell = \mathit{newloc}(X)\}.$$

The composite world $X \bullet \mathit{newworld}(X)$ describes a situation where a new location can point to locations in X , but not conversely. Then

$$\mathit{stack-alloc}[X]fs = \mathit{chop}(f[\mathit{newworld}(X)]\mathit{newloc}(X) [s \mid \mathit{newloc}(X) \mapsto a])$$

where chop takes a state in $X \bullet \mathit{newworld}(X)$ and removes $\mathit{newloc}(X)$ from its domain, giving us a state in X . This chopping operation does not create any new dangling references; in particular, if there is no dangling in s' , then there will be none in $\mathit{chop}(s')$. The point here is that the final state must obey the constraint described by $X \bullet \mathit{newworld}(X)$. The directional sharing information expressed by \multimap and \bullet should be clear from this example.

On the other hand, suppose we were to try the same thing with \bullet –

$$(\mathbf{cell} \bullet (S \rightarrow S)) \longrightarrow (S \rightarrow S)$$

Then we could make the same definition as above, by redefining $\mathit{newworld}(X)$ it could not point into X) and using $\mathit{newworld}(X) \bullet X$. But then we would no longer be guaranteed of safety of deallocation, because chopping the new location could create a dangling reference.

As before, no particular practical significance is claimed for this example: To go further one would want to allow some form of heap allocation, or one might even regard the elements of worlds as *regions* rather than single locations.

6 More on Categorical Models

In this section we look more closely at some properties of the categorical models. The reader who is more interested in seeing $\alpha\lambda$ in action can safely skip forward to the next section.

6.1 An Obstruction

Although \multimap and \rightarrow are not in general convertible to one another, the dcc isomorphisms

$$[1, A \rightarrow B] \cong \mathcal{C}[A, B] \cong [I, A \multimap B]$$

do place a demand on non-degenerate models. We can make this precise by establishing an obstacle to the existence of non-degenerate dcc's, which rules out categories such as **Set** or the category of predomains. To state this, recall that a category with a terminal object is well pointed if for any parallel maps $f, g : X \rightarrow Y$ there is $e : 1 \rightarrow X$ such that $e; f \neq e; g$.

Proposition 12 *If a cartesian dcc is well pointed then it is degenerate in at least one of the following two senses:*

- (a) *it is a preorder (at most one map in any hom set), or*
- (b) *the units 1 and I of the monoidal structures coincide (up to isomorphism).*

Proof: Suppose \mathcal{C} well pointed and not a preorder. We show that that $I \cong 1$.

To see that I is weakly terminal, since \mathcal{C} is not a preorder there are two unequal maps $A \rightarrow B$ for some A and B . By adjointness we obtain two maps $I \rightarrow A \multimap B$ which, by well pointedness, can be distinguished by a map $1 \rightarrow I$. For any object D we can compose that map with $D \rightarrow 1$, thus showing that I is weakly terminal.

For uniqueness of the map $D \rightarrow I$ we make use of the following two facts.

- (i) [15] If \mathcal{C} is a well pointed category then there is only one natural endomorphism on the identity functor $id_{\mathcal{C}} : \mathcal{C} \rightarrow \mathcal{C}$.
- (ii) [14] If $(\mathcal{C}, *, I)$ is a monoidal category then there is an injective function from $\mathcal{C}[I, I]$ into $\mathcal{C}^{\mathcal{C}}[id_{\mathcal{C}}, id_{\mathcal{C}}]$.

It follows that the identity is the only endomorphism on I in \mathcal{C} . Now, suppose (toward contradiction) that there are two maps $D \rightarrow I$. Then there would be a map $1 \rightarrow D$ distinguishing them by well pointedness, and composing gives us two maps $f, g : 1 \rightarrow I$. By well pointedness we know that \mathcal{C} is equivalent to a (not necessarily full) subcategory of **Set**, the category of its points (whose objects are hom sets $\mathcal{C}[1, A]$), where 1 corresponds to a one-point set. From this it is evident that composing on the left with the unique map $h : I \rightarrow 1$ we obtain two different endomorphisms $(h; f) \neq (h; g) : I \rightarrow I$. But we saw above that there could only be one such endomorphism, so we obtain a contradiction. So there can be at most one map $D \rightarrow I$, and thus I is terminal and isomorphic to 1. ■

A preliminary version of this paper (from October, 1997) contained the erroneous claim that in a well-pointed dcc $*$ and \multimap would collapse to \wedge and \rightarrow as well. However, Martin Hofmann has constructed well-pointed affine dcc's in which the products and function types are indeed distinct [19].

Nonetheless, the proposition does establish an obstacle to the search for models by a number of standard techniques. For example, realizability models are often given using partial equivalence relations over a partial combinatory algebra; the maps are those functions on equivalence classes that can be tracked by an element in the algebra. But such categories are well pointed, so this construction cannot be immediately used to give non-degenerate realizability models of $\alpha\lambda$. Indeed, the problem of finding a convincing realizability interpretation of $\alpha\lambda$ remains open.

6.2 Other Models

Our first example is degenerate, in that it is a poset (and even a boolean algebra).

Example 13 Consider the two-element boolean algebra $B = \{f, t\}$. It can be viewed as a degenerate (posetal) ccc, where $f \sqsubseteq t$, $1 = t$ and \wedge and \rightarrow are given by the truth tables for conjunction and implication. The product poset $B \times B$ inherits this ccc structure in a pointwise fashion, and it has symmetric monoidal closed structure given by

$$\begin{aligned} I &= (t, f) \\ (A_0, A_1) * (B_0, B_1) &= ((A_0 \wedge B_0) \vee (A_1 \wedge B_1), (A_0 \wedge B_1) \vee (A_1 \wedge B_0)) \\ (A_0, A_1) \multimap (B_0, B_1) &= ((A_0 \rightarrow B_0) \wedge (A_1 \rightarrow B_1), \\ &\quad (A_0 \rightarrow B_1) \wedge (A_1 \rightarrow B_0)) \end{aligned}$$

Coproduct structure is given by join in $B \times B$.

We can use this model to confirm the remark from Section 3.2 that \multimap and \rightarrow are not convertible to one another in the $\alpha\lambda$ -calculus. To see this, note that $((f, t) \rightarrow (t, f)) = (t, f)$ and $((f, t) \multimap (t, f)) = (f, t)$. This, combined with the fact that there are no maps between (f, t) and (t, f) in either direction, implies that there are no maps from $((f, t) \rightarrow (t, f))$ to $((f, t) \multimap (t, f))$ or back. \square

A number of other naturally occurring examples arise from Day’s construction, including higher-dimensional automata [17], complexity models [19], and logical interpretations which can be viewed as posetal dcc’s [11, 20].

A final example is given by the category **Cat** of small categories.

Example 14 **Cat** is cartesian closed, with product of categories and the one object category giving finite products and the functor category A^B giving the additive exponent. **Cat** also has another closed structure, where $A \multimap B$ is the category whose objects are functors and whose morphisms are “transformations,” i.e. families of maps but without naturality constraints. The symmetric monoidal structure is given by Gray’s tensor product [18] with the one object category as unit. So **Cat** is an affine, bicartesian dcc. These are the only symmetric monoidal closed structures on **Cat** [14]. \square

6.3 On Adding “!” to $\alpha\lambda$, or BI

In Section 3.2 we showed how $\alpha\lambda$ and linear λ -calculi mix additive and multiplicative function types in a fundamentally different way. In this section we would like to probe this issue further by asking: what happens if we add “!” to $\alpha\lambda$? To study this question we will not formulate explicit syntactic rules for “!”, but rather will work exclusively at the semantic level. Also, we will consider models that include coproduct types, so will work with bicartesian dcc’s (cartesian dcc’s that have coproducts).

We have two reasons for asking this question. First, it further illuminates the differences between the two systems. Second, it is a first step towards understanding whether it is possible to have a type system that combines the merits of linear and bunched typing.

We begin by noting a basic fact.

Proposition 15 *There is model of $\alpha\lambda$ for which the decomposition $!A \multimap B \cong A \rightarrow B$ is impossible. That is, there is no functor “!” which can decompose the additive function type into the multiplicative one, in that specific model.*

Proof: Consider the affine model from Section 5.2. We claim that there is no functor (or even function on objects) $! : \mathbf{Set}^{\mathcal{I}} \rightarrow \mathbf{Set}^{\mathcal{I}}$ admitting $!A \multimap B \cong A \rightarrow B$. To see why, consider the constant functor 2 which delivers the two element set $\{t, f\}$. Then

$$(A \multimap 2)X = \mathbf{Set}^{\mathcal{I}}[A, 2(X + -)] = \mathbf{Set}^{\mathcal{I}}[A, 2]$$

is independent of X , and so $A \multimap 2$ is a constant functor. On the other hand, $(A \rightarrow 2)X = \mathbf{Set}^{\mathcal{I}}[A(X + -), 2]$ depends on X , and is not necessarily (isomorphic to) a constant functor. For instance, if L is the inclusion functor from \mathcal{I} into \mathbf{Set} , then $(L \rightarrow 2)\{\}$ has two elements, corresponding to the two constant functions into $\{t, f\}$. But, $(L \rightarrow 2)\{a, b\}$ has elements that are not in the range of $(L \rightarrow 2)(f : \{\} \hookrightarrow \{a, b\})$. One such maps a to t and b to f (and all other inputs to, say, f). Therefore, no matter what “!” we try to pick, $!L \multimap 2$ will be a constant functor, while $L \rightarrow 2$ is not, so they cannot be isomorphic. ■

To the categorically-inclined reader this result will not be a surprise. But it does underline the fact that a dcc is not simply a model of linear logic in disguise. In fact, we have yet to find an interesting model of $\alpha\lambda$ that does admit such a decomposition. This is not an obstacle to the existence of *some* “!” satisfying the required properties for linear logic. It just shows that, in general, we cannot expect such a “!” to decompose the additive function type that exists in the $\alpha\lambda$ model.

This leaves open the possibility, then, of having a category that is simultaneously a model of $\alpha\lambda$ and a model of linear logic. Here, we will take “model of linear logic” to mean a monoidal closed category, with products and coproducts, and equipped with a “monoidal comonad”; these are the models of intuitionistic linear logic, as presented in [7]. So we explicitly define:

Definition 16 *A dcc with “!” is a cartesian dcc with coproducts (a “bicartesian dcc”), with a monoidal comonad structure (where the monoidal structure used is that for $(*, I)$).*

These are the minimum conditions we would expect from a model for a system combining $\alpha\lambda$ with “!”.

The affine model of Section 5.2 provides an example. There, for “!” we choose the functor where $!AX = A\{\}$. We omit the further data needed to describe a monoidal comonad, and simply state:

Proposition 17 $\mathbf{Set}^{\mathcal{F}S}$ *with the indicated structure is an $\alpha\lambda$ with “!” model.*

This gives us a model with an additive function type $A \rightarrow B$, the exponent in the functor category, together with an additional function type $!A \multimap B$ gotten by decomposition. To see how different the decomposed function type is, consider $p \in (!A \multimap B)X$. This gives, for any world Y , a function from $A\{\}$ to $B(X + Y)$ which is, by naturality, completely determined by a function from $A\{\}$ to BX . So, such a function effectively accepts only “resource unconscious” arguments. In fact, we can see that this comonad essentially arises from an adjunction between $\mathbf{Set}^{\mathcal{I}}$ and \mathbf{Set} . The left adjoint takes any functor A to $A\{\}$ and the right adjoint takes a set to the constant functor on it.

So, we see that it is possible to add “!” to $\alpha\lambda$, and that doing so does not completely collapse the system (at least, we have seen how to do so semantically). But if we add “!” it is natural to ask how the resulting system relates to linear logic. The system is clearly an extension of linear logic, but it is not a conservative extension of it because of the following.

Fact 18 *Any dcc with “!” satisfies distribution:*

$$A \wedge (B \vee C) \cong (A \vee B) \wedge (A \vee C).$$

However, distribution fails in models of linear logic (for instance, in the coherence space model).

This fact follows at once from the requirement that $A \wedge (-)$ be a left adjoint (because of \rightarrow), which hence must preserve coproducts.

This suggests that, while it is possible to add “!” to $\alpha\lambda$, the resulting system does not retain the merits of linear logic (though it might have other merits). The previous proposition rules out the most important models of linear type theory that have been given in the literature, including coherence space and the strict-function model from domain theory.

Ideally, we would like a way to combine linear and bunched typing in a way that simultaneously accounts for sharing as in bunched typing and consumption as in linear logic. (Examples of the consumptive aspect of linear typing include [22, 49, 33, 6].) Here we have discussed models that consist of a single category, possessing all of the properties required to model both linear logic or type theory and BI or $\alpha\lambda$, and concluded that the essence of linear logic is lost in such models. There is another way that one might try to combine linear and bunched typing, based on a pair of categories. In the pair-of-category models of linear logic one asks for a symmetric monoidal category and a separate cartesian (perhaps closed) category, with a monoidal adjunction between them [4, 2]. To obtain a combined linear/bunched type system one might start with a symmetric monoidal closed category and a separate cartesian dcc. By observing such a separation, it might be possible to develop a calculus that supports number-of-uses and sharing interpretations at the same time, where the multiplicatives in the smcc have to do with consumption and those in the dcc with sharing. The main problem, besides having convincing specific models, is to determine the right conditions on the means of passing from one category to the other, and the corresponding syntactic rules.

7 Interference Control

Now we switch gears and show a detailed use of $\alpha\lambda$. In this section we describe syntactic control of interference and Idealized Algol, two imperative languages defined by Reynolds in the late 70s and early 80s. The following section shows how SCI and IA can be combined into a single language, whose type system is based on the affine $\alpha\lambda$ -calculus. This combined language overcomes a problem with recursion in the original SCI [43, 45]. After that, we indicate how $\alpha\lambda$ can be used to treat jumps, another problem area in the original SCI.

Experience suggests that SCI can be difficult to understand if presented too quickly. Therefore, we will include a number of small examples, and some informal discussion, in this section. The main focus, again, is on the connection between structural rules and sharing.

7.1 Basic SCI

We work with a version of SCI whose types are as follows.

$$\begin{array}{ll} \rho ::= \mathbf{exp} \mid \mathbf{cell} \mid \mathbf{comm} & \text{primitive types} \\ \theta ::= \rho \mid \theta \wedge \theta' \mid \theta \multimap \theta' & \text{types} \end{array}$$

The primitive type **exp** is the type of natural number-valued expressions, **comm** is the type of commands, and **cell** is the type of storage cells, or locations.

AFFINE λ -CALCULUS

$$\begin{array}{c}
\frac{}{x : \theta \vdash x : \theta} \text{Id} \qquad \frac{\Gamma \vdash M : \theta}{\Delta \vdash M : \theta} \text{Ex (where } \Delta \text{ is a permutation of } \Gamma) \\
\\
\frac{\Gamma \vdash M : \theta'}{\Gamma, x : \theta \vdash M : \theta'} \text{W} \\
\\
\frac{\Gamma, x : \theta \vdash M : \theta'}{\Gamma \vdash \lambda x : \theta. M : \theta \multimap \theta'} \multimap I \qquad \frac{\Gamma \vdash M : \theta \multimap \theta' \quad \Delta \vdash N : \theta}{\Gamma, \Delta \vdash MN : \theta'} \multimap E \\
\\
\frac{\Gamma \vdash M : \theta \quad \Gamma \vdash N : \theta'}{\Gamma \vdash \langle M, N \rangle : \theta \wedge \theta'} \wedge I \qquad \frac{\Gamma \vdash M : \theta_1 \wedge \theta_2}{\Gamma \vdash \pi_i M : \theta_i} \wedge E \text{ (where } i \text{ is 1 or 2)}
\end{array}$$

A typing context Γ here is a list of assumptions $x : \theta$ pairing identifiers with types, with the proviso that no identifier appears twice.

The crucial rule is $\multimap E$, where the use of distinct contexts Γ and Δ prevents the procedure and argument from sharing identifiers (the proviso that no identifier appears twice in a context puts an implicit constraint on Γ, Δ). Because of this, Contraction is not admissible in this setup, though the rule of Weakening

$$\frac{\Gamma \vdash M : \theta'}{\Gamma, x : \theta \vdash M : \theta'} \text{Weakening}$$

is. In fact, an equivalent way to present the system is to include Weakening explicitly, along with a rule

$$\frac{}{x : \theta \vdash x : \theta} \text{Id}'$$

for identifiers that does not include the dummy assumption Γ .

SCI-SPECIFIC RULES

$$\begin{array}{c}
\frac{\Gamma \vdash M : \mathbf{comm} \quad \Delta \vdash N : \mathbf{comm}}{\Gamma, \Delta \vdash M \parallel N : \mathbf{comm}} \qquad \frac{\Gamma \vdash M : \mathbf{comm} \quad \Gamma \vdash N : \mathbf{comm}}{\Gamma \vdash M; N : \mathbf{comm}} \\
\\
\frac{}{\Gamma \vdash 17 : \mathbf{exp}} \qquad \frac{\Gamma \vdash N_1 : \mathbf{exp} \quad \Gamma \vdash N_i : \mathbf{comm}, i = 2, 3}{\Gamma \vdash \text{if } N_1 = 0 \text{ then } N_2 \text{ else } N_3 : \mathbf{comm}} \\
\\
\frac{x : \theta \vdash M : \theta}{\vdash \mathbf{rec } x. M : \theta} \qquad \frac{\Gamma, x : \mathbf{cell} \vdash M : \mathbf{comm}}{\Gamma \vdash \mathbf{new } x. M : \mathbf{comm}} \\
\\
\frac{\Gamma \vdash M : \mathbf{cell}}{\Gamma \vdash M : \mathbf{exp}} \qquad \frac{\Gamma \vdash M : \mathbf{cell} \quad \Gamma \vdash N : \mathbf{exp}}{\Gamma \vdash M := N : \mathbf{comm}}
\end{array}$$

We have included a rule for implicit dereferencing, which converts a term of type **cell** to one of type **exp**. Most of the other rules should be familiar; we mention only that **new** allocates a fresh cell (which is put on the runtime stack). We have not listed typical arithmetic operations.

Now let us reconsider the example from the Introduction, which leads to aliasing:

$$((\lambda x \lambda y. \dots x := 1; y := 2 \dots)z)z.$$

This term does not typecheck in SCI because the function $((\lambda x \lambda y. \dots x := 1; y := 2 \dots)z)$ and argument z share the free identifier z : there is no way to apply the elimination rule for \rightarrow .

The parallel composition $M \parallel N$ is included alongside $M; N$ for contrast. If interference control is working properly then we would expect, because of the use of disjoint contexts, that the commands M and N refer to distinct areas of storage

in $M \parallel N$. As a result, its overall effect should be determinate, and it should be semantically equivalent to the sequential composition $M; N$ (when $M \parallel N$ type-checks). For example, $x := x + 1 \parallel y := 2$ is perfectly determinate, as long as x and y denote distinct cells. But $x := x + 1 \parallel x := 2$, which would be indeterminate, is ruled out by the typing rule for \parallel .

Conversely, if interference control is not working properly, then we would expect this to be seen in $M \parallel N$. For example, $x := 1 \parallel y := 2$ would be indeterminate if x and y were aliases.

The restricted rule for recursion, where x is the only contextual variable in the premiss, is what one expects for affine typing. If M had free identifiers other than x then a fixed-point unwinding $\mathbf{rec} x.M \triangleright M[\mathbf{rec} x.M/x]$ could violate affine typing. This can be seen also with a fixed-point combinator $Y(M)$ where $M : A \multimap A$: an unwinding to $M(Y(M))$ would violate the disjointness property of procedure calls, if M was not closed.

As we mentioned in the Introduction, the original SCI allowed a restricted form of Contraction for *passive* types, which are types of values that may read from, but not write to, the store. Passivity is discussed briefly in Section 8.3.

7.2 The Sharing Interpretation of SCI

We saw above how abolishing Contraction eliminates one instance of aliasing. More generally, the absence of aliasing is subsumed under the

DISJOINTNESS POLICY: distinct identifiers never interfere.

In the language here we take “interfere” to mean “refer to common storage.”

The disjointness policy impacts the meaning of function types, while the meaning of products remains more standard:

$A \multimap B$: functions that don’t interfere with their arguments;
 $A \wedge B$: pairs that may interfere with one another.

SCI did not, originally, have a multiplicative product. The reading for it would be

$A * B$: pairs whose components don’t interfere with one another.

But a form of this product is already present in the comma in typing contexts, in that in a judgement

$$x_1 : A_1, \dots, x_n : A_n \vdash M : B$$

the disjointness policy expresses the same non-interference property as for $*$.

It is important to realize how the sharing interpretation is an unusual reading of the affine λ -calculus. Often, the idea in the affine calculus is that a function uses its argument at most once, so that for instance in a function of type $A \wedge B \multimap C$ either the A or the B component may be used, but not both. But according to SCI’s reading, it is perfectly reasonable for a function p of such a type to use either or both components of a pair $\langle a, b \rangle$ supplied to it as an argument, and either of these elements could be used many times. The only constraint is that p doesn’t interfere with $\langle a, b \rangle$.

For example, in SCI we can write a function

$$(\lambda c : \mathbf{comm} \wedge \mathbf{comm} . \pi_1 c ; \pi_2 c ; \pi_1 c) : \mathbf{comm} \wedge \mathbf{comm} \multimap \mathbf{comm}$$

that uses the first component of a pair twice and the second component once.

The sharing reading also helps to understand the typing of **if**. In the number-of-uses reading, in **if** $N_1 = 0$ **then** N_2 **else** N_3 one would expect to use one context

for N_1 , and a separate context for N_2 and N_3 . But the conditional essentially corresponds to a constant of type $\mathbf{exp} \wedge \mathbf{comm} \wedge \mathbf{comm} \multimap \mathbf{comm}$ in SCI and there is no inconsistency if all the N_i 's share the same context. In imperative programming this sharing is often wanted, so that information can pass from the condition into the branches.

Now the affine λ -calculus certainly does not *force* the sharing reading. But it is consistent with it. The pure affine calculus is actually too small for this “many uses” aspect to be seen; the additional constants of SCI are where it comes out. The pure $\alpha\lambda$ -calculus, in contrast, already admits multiplicative functions that use their arguments many times, as we saw in Section 3.2.

7.3 IA

IA (for Idealized Algol) is similar to SCI, except that it uses the simply-typed λ -calculus in place of the affine λ -calculus. Formally, it is obtained from Basic SCI by removing the rule for \parallel and adding Contraction and a new rule for recursion

$$\frac{\Gamma, x : \theta, y : \theta \vdash M : \theta'}{\Gamma, z : \theta \vdash M[z/x, z/y] : \theta'} \text{ Contraction} \quad \frac{\Gamma, x : \theta \vdash M : \theta}{\Gamma \vdash \mathbf{rec } x. M : \theta}$$

Instead of adding Contraction, we could equivalently banish the disjointness requirement in the $\multimap E$ rule. IA violates the disjointness policy, as now a term $((\lambda x \lambda y. \dots x := 1; y := 2 \dots)z)z$ where distinct identifiers x and y interfere is typable.

For future reference (Proposition 19), in IA we also rename \multimap to \rightarrow , to emphasize that it uses simply-typed λ -calculus.

7.4 A Limitation

Many programs one would typically write (in a language, or language fragment, like IA without references or pointers) do in fact satisfy the disjointness policy of SCI. But a problem with recursion was raised by Reynolds [43]: If a recursive procedure contains a free identifier which uses storage in an active way (by changing it), then in the body of the procedure this free identifier and the procedure being defined will interfere (violating the disjointness policy). Technically, this problem is avoided in the affine type system in this section by restricting the rule for recursion, so that a recursive procedure cannot have any free identifiers.

An example of this limitation is the Towers of Hanoi program, where disks are moved between pegs.

```

procedure movemany(k, a, b, c : int)
  if k > 0 then
    movemany(k - 1, a, c, b);
    moveone(a, b);
    movemany(k - 1, c, b, a)

```

The procedure *moveone* can work by printing a message to the screen, or by recording a move in a global data structure.

Technically, since *moveone* is free in the body of the procedure, we cannot use the restricted rule for recursion to type it. Desugaring the recursion, we would have to type $\mathbf{rec } \mathbf{movemany}. \lambda k a b c. \mathit{body}$ in a context that contains *moveone*, where the recursion rule requires an empty context.

More conceptually, *moveone* and *movemany* interfere in the body of the procedure if *moveone* contains side effects. Other examples of this form may be found in objects where one of the methods is recursive.

It is possible to write a recursive version of *movemany* in SCI by passing *moveone* as a dummy argument, and instantiating a curried version of the procedure with the actual *movemany*. But this seems unnecessarily complex, as the given definition of *movemany* is simple and clear enough as is; as a result, it does seem to be desirable to be able to turn off interference control in local contexts, as long as we can turn it back on again in a broader context. The “problem” would be exacerbated when programming an object that uses several cells to maintain a local state.

8 An Enveloping Language

8.1 SCI+

Now we consider a language, SCI+, that has primitive operations similar to those in IA and SCI, but which uses the affine $\alpha\lambda$ -calculus as its type system.

The types are given by the following grammar.

$$\begin{array}{ll} \rho ::= \mathbf{exp} \mid \mathbf{cell} \mid \mathbf{comm} & \text{primitive types} \\ \theta ::= \rho \mid \theta \wedge \theta' \mid \theta \rightarrow \theta' \mid \theta \multimap \theta' & \text{types} \end{array}$$

The primitive types are as in SCI and Idealized Algol, and we include both of the function types of $\alpha\lambda$, with the rules from the affine variant as in Section 3.3. We also include the rules for cartesian products.

$$\frac{\Gamma \vdash M : A \quad \Gamma \vdash N : B}{\Gamma \vdash \langle M, N \rangle : A \wedge B} \wedge I \quad \frac{\Gamma \vdash M : A_1 \wedge A_2}{\Gamma \vdash \pi_i M : A_i} \wedge E \text{ (where } i \text{ is 1 or 2)}$$

SCI+-SPECIFIC TYPING RULES.

$$\begin{array}{c} \frac{\Gamma \vdash M : \mathbf{comm} \quad \Delta \vdash N : \mathbf{comm}}{\Gamma, \Delta \vdash M \parallel N : \mathbf{comm}} \\ \\ \frac{}{\Gamma \vdash 17 : \mathbf{exp}} \\ \\ \frac{\Gamma; x : \theta \vdash M : \theta}{\Gamma \vdash \mathbf{rec } x. M : \theta} \\ \\ \frac{\Gamma \vdash M : \mathbf{cell}}{\Gamma \vdash M : \mathbf{exp}} \end{array} \quad \begin{array}{c} \frac{\Gamma \vdash M : \mathbf{comm} \quad \Gamma \vdash N : \mathbf{comm}}{\Gamma \vdash M; N : \mathbf{comm}} \\ \\ \frac{\Gamma \vdash N_1 : \mathbf{exp} \quad \Gamma \vdash N_i : \mathbf{comm}, i = 2, 3}{\Gamma \vdash \mathbf{if } N_1 = 0 \mathbf{ then } N_2 \mathbf{ else } N_3 : \mathbf{comm}} \\ \\ \frac{\Gamma, x : \mathbf{cell} \vdash M : \mathbf{comm}}{\Gamma \vdash \mathbf{new } x. M : \mathbf{comm}} \\ \\ \frac{\Gamma \vdash M : \mathbf{cell} \quad \Gamma \vdash N : \mathbf{exp}}{\Gamma \vdash M := N : \mathbf{comm}} \end{array}$$

All of these rules except for **rec** are, textually, exactly the same as rules in IA and SCI. The difference is that now the comma has a different meaning than in IA, in that it refers to the multiplicative combination. If we read the IA “,” as “;” in SCI+, and the SCI “,” as “;”, then we have omitted the rule for recursion from SCI, and rule for **new** from IA. Let us see that the omitted rules are derivable.

The IA **new** would be

$$\frac{\Gamma; x : \mathbf{cell} \vdash M : \mathbf{comm}}{\Gamma \vdash \mathbf{new } x. M : \mathbf{comm}}.$$

We can derive this at once using the SCI+ rule for **new** and the inference

$$\frac{\Gamma; x : \mathbf{cell} \vdash M : \rho}{\Gamma, x : \mathbf{cell} \vdash M : \mathbf{comm}}$$

which is an instance of the *Conv* rule of affine $\alpha\lambda$.

The recursion rule we have given is the one appropriate to IA. It has the SCI rule as a special case, using $\Gamma = 1$ and a coherent equivalence. The use of “;” instead of “,” in this rule is crucial.

Given these remarks it is not difficult to show the following.

Proposition 19 1. *SCI+ has IA as a sublanguage. That is, if*

$$x_1 : A_1, \dots, x_n : A_n \vdash M : B$$

in IA then

$$x_1 : A_1^*; \dots; x_n : A_n^* \vdash M^* : B^*$$

in SCI+, where $(\cdot)^$ maps λ to α , MN to $M^\circ @ N^\circ$, and everything else (inductively) to itself.*

2. *SCI+ has SCI as a sublanguage. That is, if*

$$x_1 : A_1, \dots, x_n : A_n \vdash M : B$$

is derivable in SCI then it is also derivable in SCI+.

8.2 The Sharing Interpretation

There is thus a syntactic sense in which SCI+ is an enveloping language, but this in itself is unremarkable. It is still conceivable that the larger language has features that are inconsistent with the essence of IA or SCI, destroying some crucial aspect of one of smaller languages.

The sharing interpretation of $\alpha\lambda$ describes the sense in which the larger language preserves the essence of SCI; the readings of \multimap and \wedge are exactly as in SCI. The reading of \rightarrow is one that is appropriate to IA. To sum up:

$A \wedge B$	pairs that may access a common portion of the store
$A \multimap B$	procedures that don't share store with arguments
$A \rightarrow B$	procedures that may share store with arguments

The resulting sense in which the $\alpha\lambda$ -calculus allows detection of interference is that whenever we see a sequence $\alpha x \lambda y$ or $\lambda x \lambda y$ we know that x and y don't interfere. So, non-interference can be inferred (in a fail-safe manner) from a simple inspection of a context. The one difference is that in Basic SCI this determination is context free. It is context sensitive in SCI+ because when we see $\alpha x \alpha y$ or $\lambda x \alpha y$ we don't know if x and y interfere or not.

Let us revisit the Towers of Hanoi in light of this interpretation. The *movemany* procedure can now be typed without difficulty, because we have IA as a sublanguage. It is instructive, however, to look at the way the typing works, as it illustrates the way we can move between SCI-style and IA-style typing in SCI+.

Using the rule for recursion we can type (with a little syntactic sugar)

```

moveone : exp → exp → comm
⊢ rec movemany . ak a b c : exp
      if k > 0 then
          movemany(k - 1, a, c, b);
          moveone(a, b);
          movemany(k - 1, c, b, a)
: exp → exp → exp → exp → comm

```

The critical point is that, during the typing of the body, we turn interference control off by using the bunch

$$\begin{array}{l} \text{moveone} : \mathbf{exp} \rightarrow \mathbf{exp} \rightarrow \mathbf{comm} \\ ; \text{movemany} : \mathbf{exp} \rightarrow \mathbf{exp} \rightarrow \mathbf{exp} \rightarrow \mathbf{exp} \rightarrow \mathbf{comm} \end{array}$$

which indicates that *moveone* and *movemany* might interfere. But more globally we can turn interference control back on and, for instance, run a call to the recursive procedure in parallel with another command, as long as that command doesn't interfere with *moveone*.

$$\begin{array}{l} \text{moveone} : \mathbf{exp} \rightarrow \mathbf{exp} \rightarrow \mathbf{comm}, c : \mathbf{comm} \\ \vdash ((\mathbf{rec} \text{movemany} . \dots)7123) \parallel c : \mathbf{comm} \end{array}$$

8.3 Passivity

The language Basic SCI in Section 7.1 is in fact only a fragment of Syntactic Control of Interference, which includes typing rules for passivity [43, 45, 30]. A passive entity, such as a side-effect free expression, can safely be shared without causing interference, and a passive type is one whose elements are all passive. Bunches are compatible with the approach to passivity in the SCIR type system from [30]; we briefly indicate how this is so.

The SCIR type system uses judgements $\Pi \mid \Gamma \vdash M : A$, where the context is split into a passive zone Π and an active zone Γ . The three critical rules of the system are the permeability rules of Activation and Passification, for moving identifier across the \mid separator, and Contraction in the passive zone.

$$\begin{array}{c} \frac{\Pi \mid x : B, \Gamma \vdash M : A}{\Pi, x : B \mid \Gamma \vdash M : A} \text{Passif (where } A \text{ is passive)} \quad \frac{\Pi, x : B \mid \Gamma \vdash M : A}{\Pi \mid x : B, \Gamma \vdash M : A} \text{Activ} \\ \\ \frac{\Pi, y : B, z : B \mid \Gamma \vdash M : A}{\Pi, y : B \mid \Gamma \vdash M[y/z] : A} \text{Contr} \end{array}$$

The zonal presentation does not work well with bunches, because we would want to be able to indicate that an identifier is passive without saying that a whole bunch is.

The solution is to allow “marked assumptions” ($x^* : A$) alongside “normal assumptions” ($x : A$). Then, with an extension of marking to contexts, the three rules are as follows:

$$\begin{array}{c} \frac{\Gamma(\Delta) \vdash M : A}{\Gamma(\Delta^*) \vdash M : A} \text{Passif (where } A \text{ is passive)} \quad \frac{\Gamma(\Delta^*) \vdash M : A}{\Gamma(\Delta) \vdash M : A} \text{Activ} \\ \\ \frac{\Gamma(\Delta^*, \Delta'^*) \vdash M : A}{\Gamma(\Delta) \vdash M[i(\Delta^*)/i(\Delta'^*)] : A} \text{Contr (where } \Delta^* \cong \Delta'^*) \end{array}$$

This gives us a limited form of Contraction for “;”, in addition to the general Contraction for “,”. These rules can all be interpreted using the *bireflective subcategory* structure found in Tennent’s model [30, 15].

8.4 Remaining Limitations

The language here uses call by name as its parameter passing mechanism. The extension of the approach to call by value does not appear to raise insuperable difficulties, but the typing rules required tend to become more complex when one wants to separate out the effects caused by evaluation to a value, from those that

are “latent”. Latent effects occur only later, when using the value; for example, in $x := 1; (\lambda y. z := 2)$ the assignment to z is latent. In call by name, all effects at higher order are latent.

A more significant limitation is that it is not obvious how to incorporate higher-order store, where a reference may hold a procedure or another reference as its contents. This problem has been addressed using different methods, and with some success, by the effect systems of Gifford and Lucassen [23]. The types used in effect systems are, however, very detailed, and they seem more suitable to an intermediate language used in a compiler (where complexity can be hidden from the programmer) than in a source language. It would be worthwhile to develop a more abstract form of control over higher-order references, along the lines of SCI.

In a recent development, Walker and Morrisett have devised a fascinating system for interference control [50], which correctly handles higher-order store, and which is remarkably similar in structure to a program logic connected to BI [42, 20]. Despite the structural similarity, the relationship of their work to bunched typing is not yet clear.

9 A Model for SCI+

In this section we describe a semantics of SCI+. Our purpose in doing this is to back up the informal interpretation of types from Section 8.2. So we will concentrate on describing the structure of the model, and how it relates back to the informal description.

9.1 Semantics of Types

We are going to use an affine model similar to the one in Section 5.2. However, as explained in [34], in order for **new** to satisfy naturality requirements of functor categories, it is necessary to allow for renaming of locations. Therefore, we use the category \mathcal{I} of finite sets and injections as the category of worlds. Also, to interpret recursion we will use domains in place of sets in the target category.

Let **Predom** denote the category of predomains (ω -complete posets and continuous maps).

For X a finite set we define

$$\begin{aligned} \llbracket \mathbf{comm} \rrbracket X &= SX \Rightarrow SX_{\perp} \\ \llbracket \mathbf{exp} \rrbracket X &= SX \Rightarrow N_{\perp} \\ \llbracket \mathbf{cell} \rrbracket X &= X_{\perp} \end{aligned}$$

Here, $SX = X \Rightarrow N$ is the set of states at world X , and N is the set of natural numbers. Variations are possible. For example, we could allow side effects in expressions, using $SX \Rightarrow (N \times SX)_{\perp}$, or we could make cells dependent on the store (this allowing conditional storage cells).

The action of each primitive type on morphisms f in \mathcal{I} is defined by renaming cells according to f and ignoring cells not in its range. The cases of **exp** and **cell** are simple:

$$\begin{aligned} \llbracket \mathbf{exp} \rrbracket f e s &= e(f; s), \\ \llbracket \mathbf{cell} \rrbracket f &= f_{\perp}. \end{aligned}$$

In the case of **comm**, when $f : X \rightarrow Y$,

$$\llbracket \mathbf{comm} \rrbracket f c s = \begin{cases} \perp & \text{if } c(f; s) = \perp \\ s' & \text{if } c(f; s) = s'' \text{ and} \\ & \forall \ell \in Y. (\ell = f\ell' \text{ implies } s'(\ell) = s''(\ell')) \text{ and} \\ & \ell \notin \text{range}(f) \text{ implies } s'(\ell) = s(\ell). \end{cases}$$

The functor category $\mathbf{Predom}^{\mathcal{I}}$ is cartesian closed, with finite products defined pointwise. The additive function type can be defined as follows.

$$(A \rightarrow B)(X) = \mathbf{Predom}^{\mathcal{I}}[A(X + -), B(X + -)].$$

This accurately reflects the informal reading, in that the presence of X in the argument type $A(X + -)$ indicates how a function $p \in (A \rightarrow B)X$ may share access to X with its argument.

This is not the standard representation of the exponent in a functor category. We are relying on the fact that any $f : X \rightarrow Z$ factors into a left injection $i : X \rightarrow X + Y$ followed by an isomorphism $j : X + Y \rightarrow Z$. Such a factorization is used to define the morphism part of $A \rightarrow B$. If $p \in (A \rightarrow B)X$ then $(A \rightarrow B)fp \in (A \rightarrow B)Z$ is defined by the formula

$$(A \rightarrow B)fpW a = B(j)(pW'(Aj^{-1}a))$$

where $i : X \rightarrow X + W'$, $j : X + W' \rightarrow W$ is an injection/isomorphism factorization of $f; i : X \rightarrow W$.

In this description of the function type the hom sets $\mathbf{Predom}^{\mathcal{I}}[A, B]$ are considered to be ordered pointwise. Also, $+$ is the evident functor on \mathcal{I} given by disjoint union of finite sets.

The multiplicative function type once again expresses disjointness of a function from its argument:

$$(A * B)X = \mathbf{Predom}^{\mathcal{I}}[A, B(X + -)].$$

To see how the semantics is working, consider the type $\mathbf{cell} * \mathbf{cell} * \mathbf{comm}$. Semantically, an element $p \in \llbracket \mathbf{cell} * \mathbf{cell} * \mathbf{comm} \rrbracket \{\}$ accepts

two worlds Y and Z ,
cells $c \in Y_{\perp}$ and $e \in Z_{\perp}$

and produces (using $\{\} + Y + Z \cong Y + Z$)

$$p[Y]c[Z]e : S(Y + Z) \Rightarrow S(Y + Z)_{\perp}.$$

It is evident from this that the arguments c and e cannot be aliases, as they live in disjoint portions of the store at world $Y + Z$.

To treat recursion in this model we must effect a transformation

$$\frac{p[X] : \llbracket \Gamma \rrbracket X \times \llbracket \theta \rrbracket X \longrightarrow \llbracket \theta \rrbracket X}{\text{rec}(p) : \llbracket \Gamma \rrbracket X \longrightarrow \llbracket \theta \rrbracket X}.$$

Here, we uncurry p to obtain a map of type $\llbracket \Gamma \rrbracket X \longrightarrow (\llbracket \theta \rrbracket X \rightarrow \llbracket \theta \rrbracket X)$ and then compose on the right with the least fixed-point operator for pointed domains.

For this interpretation to exist we must have that each $\llbracket \theta \rrbracket X$ is pointed. And for it to be natural we require that each morphism part $\llbracket \theta \rrbracket f$ preserves least elements [36]. These properties are satisfied by all the types in SCI+, and are part of the identification of a subcategory of $\mathbf{Predom}^{\mathcal{I}}$ in the following section.

9.2 Bunches, Environments and Non-interference

We now give a precise treatment of bunches. Our intention in doing this is to show one example where the syntactic ambiguity resulting from the rule for coherent equivalence

$$\frac{\Gamma \vdash M : A}{\Delta \vdash M : A} \equiv (\text{where } \Delta \equiv \Gamma)$$

is dealt with by requiring that equivalent bunches be semantically equal.


The presentation in this section will be more technical than the others, but it is mainly a pulling together of known results [36, 29]. Some readers may wish to skip forward to Section 9.3, where some of the most important valuations are presented in a way that doesn't depend on the details of this section.

We work in a full subcategory of $\mathbf{Predom}^{\mathcal{I}}$, the category whose objects A have the following two properties:

1. each AX has a least element, and each $Af : AX \rightarrow AY$ preserves least elements, and
2. A preserves pullbacks.

The strictness part of the first condition is needed for naturality of the fixed-point operator, and the second condition enables a simplified description of environments. It is straightforward to verify that the meanings of primitive types satisfy these conditions. We call this category \mathcal{M} .

Pullbacks in \mathcal{I} are like those in the category of sets. In particular,



is a pullback square, where the unlabeled arrows are inclusion functions. Because of this, for pullback-preserving functors there is always a smallest world that any $a \in AX$ comes from:

- we say that a comes from $Y \subseteq X$ iff $\exists a' \in A(\text{support}(a)). a = A(\text{support}(a) \hookrightarrow X)a'$, and
- we refer to the smallest world that a comes from as the support of a , written $\text{support}(a)$.

The notion of support does not work for arbitrary functors, as there need not be a unique smallest world that a comes from. But the existence of such worlds, as guaranteed by pullback preservation, seems intuitively reasonable if we want to consider support as the set of locations that a computational entity depends on.

Given this notion of support, we can define non-interference. If $a \in AX$ and $b \in BX$, then

$$\bullet a \# b \Leftrightarrow \text{support}(a) \cap \text{support}(b) = \{\}.$$

Using this notion of non-interference we can finally describe the tensor product $*$ of functors:

$$\begin{aligned} (A * B)X &= \{(a, b) \in AX \times BX \mid a \# b\}, \\ (A * B)f(a, b) &= (Afa, Bfb). \end{aligned}$$

For this definition to work correctly, it is important that $(A * B)f$ preserves non-interference and that $(A * B)X$ is ω -complete; see [29].

Proposition 20 *\mathcal{M} is an affine doubly closed category. That is, $1, \wedge, \rightarrow$ is cartesian closed structure and $1, *, \dashv$ monoidal closed structure.*

This gives us all the structure we need to interpret the affine $\alpha\lambda$ -calculus, where we interpret a bunch by mapping “,” to $*$ and “;” to \wedge . However, a more concrete semantics of bunches is useful. For this, we first define

$$\Gamma \vdash x \# y \iff x \text{ and } y \text{ have a “,” as a common ancestor node in } \Gamma.$$

Bunches are then interpreted by relating the syntactic $\#$ to the semantic one:

$$\llbracket \Gamma \rrbracket W = \{u \in \prod_{x \in i\Gamma} \Gamma(x) \mid \Gamma \vdash x \# y \implies ux \# uy\}$$

where we are writing $\Gamma(x)$ for the type of x in Γ . This representation of environments allows us to ignore coherent equivalence, while still maintaining a relationship between “,” and $*$ and between “;” and \wedge .

Proposition 21 *If $\Gamma \equiv \Delta$ then $\llbracket \Gamma \rrbracket = \llbracket \Delta \rrbracket$. Further, we have the following isomorphisms:*

$$\llbracket \Gamma, \Delta \rrbracket \cong \llbracket \Gamma \rrbracket * \llbracket \Delta \rrbracket \quad \llbracket \Gamma; \Delta \rrbracket \cong \llbracket \Gamma \rrbracket \wedge \llbracket \Delta \rrbracket$$

The point of this concrete interpretation of bunches is that it makes clear that the use of the rule for \equiv in $\alpha\lambda$ -calculus is not in any way problematic.

9.3 Selected Valuations

We will make use of a multi-map characterization of maps out of $A * B$, which enables a simple description of most of the maps in the semantics. To repeat the point made for the basic disjointness model, since we expect to have an isomorphism $\mathbf{Predom}^{\mathcal{I}}[A * B, C] \cong \mathbf{Predom}^{\mathcal{I}}[A, B \dashv C]$, we must obtain the following, no matter what $*$ is.

*Maps $p : A * B \rightarrow C$ out of a tensor are in bijection with families of functions*

$$\bar{p}[X][Y] : AX \times BY \rightarrow C(X + Y),$$

natural in X and Y .

Because of this characterization, we will expect maps out of a context Γ, Δ to be in bijection with families of maps where we use one world for Δ and another for Γ . Based on this assumption, we now give the semantics of several terms.

We begin with \parallel . First, we define a state transformation $c \parallel c'$, when c and c' are commands referring to disjoint store:

$$\frac{c : S(X) \rightarrow S(X)_{\perp} \quad c' : S(Y) \rightarrow S(Y)_{\perp}}{c \parallel c' = \lambda[s, s'] : S(X + Y). [cs, c's'] : S(X + Y) \rightarrow S(X + Y)_{\perp}}$$

where $[\cdot, \cdot] : S(X + Y) \rightarrow S(X) \times S(Y)$ is the evident isomorphism and $[cs, c's']$ is \perp if either cs or $c's'$ is. Then we can interpret the term-formation rule for \parallel as follows.

$$\frac{p[X] : \llbracket \Gamma \rrbracket X \rightarrow \llbracket \mathbf{comm} \rrbracket X \quad q[Y] : \llbracket \Gamma \rrbracket Y \rightarrow \llbracket \mathbf{comm} \rrbracket Y}{\Lambda X. \Lambda Y. \lambda\langle u, v \rangle. p[X]u \parallel q[Y]v : \llbracket \Gamma \rrbracket X \times \llbracket \Delta \rrbracket Y \rightarrow \llbracket \mathbf{comm} \rrbracket (X + Y)}.$$

Here, we have used polymorphic λ -calculus notation to talk about families of maps in what should be a clear way [33].

This semantics makes obvious that different components of \parallel act on disjoint portions of the store. In contrast, the rule for sequential composition uses the same context Γ for both commands, and so we use the transformation

$$\frac{p[X] : \llbracket \Gamma \rrbracket X \longrightarrow \llbracket \mathbf{comm} \rrbracket X \quad q[X] : \llbracket \Gamma \rrbracket X \longrightarrow \llbracket \mathbf{comm} \rrbracket X}{\Lambda X. \lambda u. (p[X]u); (q[X]u) : \llbracket \Gamma \rrbracket X \longrightarrow \llbracket \mathbf{comm} \rrbracket X}$$

where “;” is composition of partial functions. The common use of X by p and q makes clear that they access the same portion of store.

For **new** declarations we again appeal to disjointness, where the declared cell is disjoint from the store in use when a declaration begins execution. The semantic transformation is

$$\frac{p[X][Y] : \llbracket \Gamma \rrbracket X \times \llbracket \mathbf{cell} \rrbracket Y \longrightarrow \llbracket \mathbf{comm} \rrbracket X + Y}{\Lambda X. \lambda u. \lambda s. f(p[X][\{*\}] * (s \mid * \mapsto 0)) : \llbracket \Gamma \rrbracket X \longrightarrow \llbracket \mathbf{comm} \rrbracket X.}$$

where $f : S(X + \{*\})_{\perp} \rightarrow S(X)_{\perp}$ forgets the $\{*\}$ component.

Finally, for λ -abstraction a λ -bound variable abstracts over meanings defined in world that is separate from the world for other free identifiers.

$$\frac{p[X][Y] : \llbracket \Gamma \rrbracket X \times \llbracket \theta \rrbracket Y \longrightarrow \llbracket \theta \rrbracket X + Y}{\lambda u. \Lambda Y. \lambda x \in \llbracket \theta \rrbracket Y. p[X][Y]\langle u, x \rangle : \llbracket \Gamma \rrbracket X \longrightarrow \llbracket \theta * \theta' \rrbracket X.}$$

Thus, in $\lambda x. M$ the identifier x does not share storage with any other identifier free in M .

The semantic model described in this section accomplishes two things. First, and foremost, we claim that it achieves our basic aim, of substantiating the informal sharing reading.

Second, it shows that

Proposition 22 *SCI+ has IA and SCI as semantic sublanguages.*

To be precise, what this means is

1. the model obtained from the translation $(\cdot)^*$ from Proposition 19 is a standard functor-category model of IA [34].
2. the semantics of the SCI fragment of SCI+ is the semantics of SCI given in [29] (ignoring passivity),

The only real differences in the various interpretations are the rules in SCI or IA that were left out of SCI+, but which were shown to be derivable. The most important case is **new**: the reason it does not present a difficulty is that, even in IA, a locally declared cell doesn’t interfere with any other identifiers free in its defining block. This is why the use of “;” in the SCI+ rule for **new** is semantically sufficient to capture IA’s **new**.

10 Jumps

Jumps cause a problem broadly similar to the one with recursion in SCI. In this section we indicate how this problem can be overcome using $\alpha\lambda$.

To see the difficulty, consider a block **escape** x in M . This declares a new label which, when jumped to from within M , results in a transfer of control to the end of the block. From the point of view of continuation semantics, it binds x to the current continuation, which is a function from states to final answers that describes computation that will take place after the block is finished. This means that, if the

computation associated with the current continuation changes any storage cell then x will interfere with that cell. So, in $(\mathbf{escape} \ x \ \mathbf{in} \ M); z := 4$ the identifiers z and x interfere, if z occurs within M .

Thus, from the point of view of continuation semantics, the **escape** statement violates the requirement that distinct identifiers never interfere (unless we put rather draconian conditions on identifiers appearing in or following an **escape** block). One might attempt to use a different form of semantics to define a different notion of interference for labels. It will be simpler just to allow this interference, by arranging the typing rule so that x is set additively apart from other identifiers.

Following [36], we add a primitive type **compl** for completions (labels) and remove **comm**. We now regard **comm** as syntactic sugar for **compl** \rightarrow **compl**. The semantics of the type of completions is given using a fixed domain A of answers.

$$\llbracket \mathbf{compl} \rrbracket X = SX \Rightarrow A$$

$$\llbracket \mathbf{compl} \rrbracket f k s = k(f; s)$$

With this semantics, the language with the completion type has the same sense of non-interference as SCI+: the semantics of \rightarrow ensures that whenever we see a sequence $\alpha x \alpha y$ or $\alpha x \lambda y$ we know that x and y access different portions of the store.

The central syntactic rules are

$$\frac{\Gamma; x : \mathbf{compl} \vdash M : \mathbf{comm}}{\Gamma \vdash \mathbf{escape} \ x \ \mathbf{in} \ M : \mathbf{comm}} \quad \frac{\Gamma \vdash M : \mathbf{compl}}{\Gamma \vdash \mathbf{goto} \ M : \mathbf{comm}}$$

where the use of “;” in the rule for **escape** allows for the interference between x and identifiers in Γ . These can be given a standard continuation semantics, exactly as was done by Oles [36]. For **escape**, we effect a transformation

$$\frac{\llbracket \Gamma; x : \mathbf{compl} \rrbracket \xrightarrow{f} \llbracket \mathbf{comm} \rrbracket}{\llbracket \Gamma \rrbracket \xrightarrow{f'} \llbracket \mathbf{comm} \rrbracket}}$$

accomplished by binding x to the current continuation:

$$f'[X] u X' k s = f[X + X'](u' \mid x \mapsto k) k s$$

where

$$u \in \llbracket \Gamma \rrbracket X,$$

$$k \in \llbracket \mathbf{compl} \rrbracket (X + X'),$$

$$i : X \rightarrow X + X' \text{ is the left injection,}$$

$$u' = \llbracket \Gamma \rrbracket i u,$$

$$s \in S(X + X').$$

Here, the extra parameter X' is occurring because **comm** = **compl** \rightarrow **compl** is a procedure type, and we are using the representation of \rightarrow given in the last section.

goto is given by a map $g : \llbracket \mathbf{compl} \rrbracket \rightarrow (\llbracket \mathbf{compl} \rrbracket \rightarrow \llbracket \mathbf{compl} \rrbracket)$ from completions to commands, which ignores the current continuation:

$$g[X] k [X'] k' = \llbracket \mathbf{compl} \rrbracket i k$$

where $i : X \rightarrow X + X'$ is again the left injection.

Finally, to illustrate the effect of interference constraints we define the parallel composition of completions.

$$\frac{\Gamma \vdash M : \mathbf{compl} \quad \Delta \vdash N : \mathbf{compl}}{\Gamma, \Delta \vdash M \parallel N : \mathbf{compl}}$$

Its semantics is given by a map $par : \llbracket \mathbf{compl} \rrbracket * \llbracket \mathbf{compl} \rrbracket \longrightarrow \llbracket \mathbf{compl} \rrbracket$ and is defined similarly to the parallel composition of commands. For this, we refer again to the multi-map characterization of maps out of $*$, and define

$$par[X][X'] \langle k, k' \rangle [s, s'] = (ks \cdot k's')$$

where $(- \cdot -) : A \times A \rightarrow A$ is a function that puts together two final answers. For concreteness, we take A to be the two-point cpo $\{t\}_\perp$ and $(- \cdot -)$ to be meet. Here, we regard an answer t as indicating termination. We admit that this use of a function on answers is *ad hoc*. It does, however, enable us to show a sense in which completions typed in contexts separated by “,” do not interfere.

We have not included parallel composition for commands, because the right way to do so is not obvious. For, one of the commands in $M \parallel N$ might jump out, and ignore the current continuation. (It might be possible to use bunches to control the range of continuations; that, however, is beyond the scope of this paper.)

11 Conclusion

The $\alpha\lambda$ -calculus and BI offer a new perspective on how control over structural rules translates into control over access to resources in a computer system. As we have suggested here, the main point is the emphasis on sharing, supported by a spatial view of possible world semantics which has developed over a number of years [44, 27, 30, 33].

We began the paper by recounting an analogy between syntactic control of interference and linear logic, where both systems limit the use of Contraction. This was followed by recalling a dilemma: Although there is a formal similarity, there is also an important conceptual difference; control of Contraction in SCI is about sharing, while in LL it is primarily about duplication.

Now the reader might feel that we are splitting hairs here, as at first sight duplication versus sharing may appear to be a case of six of one versus a half dozen of the other. But the distinction is crucial in computer science. The number-of-uses explanation of linear logic calls to mind the notion of *temporary resources* in Operating Systems [8], the canonical example of which is a message produced by one process and consumed by another. The analogy with temporary resources is clear in several formal interpretations of linear logic, including the original coherence space model [16] and a concurrency reading [1]. In contrast, the sharing interpretation of $\alpha\lambda$ concerns what is often labelled a *permanent resource*. Here, permanent does not literally mean permanent, but potentially long lived; examples include files, external devices, or portions of the store. For this kind of resource it is sharing, rather than consumption, that is the prime concern.

The results of this paper give one answer to the question of whether the conceptual difference between sharing and duplication should lead to different formal structure. We described a new calculus, the $\alpha\lambda$ -calculus, which we showed differs from linear λ -calculi in several significant respects. And for each difference between the systems we were able to offer an explanation of $\alpha\lambda$'s stance by appeal to a *sharing interpretation*, where linear logic's stance can be understood in terms of a *number of uses reading*. Because different formal systems fit each of these readings

we claim that the differences are genuinely structural, and run deeper than merely having separate models of the same system.

Finally, it is worth mentioning a related resource perspective on BI, which does not mention λ -terms. Here we speak of the resources a function has access to which, when we erase λ -terms, corresponds to talking about proofs. A similar interpretation can be given on a purely logical level, where one views \multimap and \rightarrow as implications, and where the semantics is phrased in terms of truth conditions; proofs are not mentioned. This semantics of BI [31, 38], which was first advanced by Pym in 1997, is similar to the functional interpretation we derived from SCI, but genuinely different because of its declarative character: a number of interesting models have been described that make good sense from a truth-based perspective, but that have much less immediate type-theoretic significance [20, 25, 11, 10]. Incidentally, several of these models do not admit Weakening, and so correspond to the basic system of Section 3.1 rather than the affine variant used in the application to SCI.

ACKNOWLEDGEMENTS.

I am especially grateful to David Pym and John Reynolds, both for their original insights and for numerous discussions. Thanks to Martin Hofmann and Guy McCusker for pointing out glitches in an early version. This research was supported by the EPSRC and, during a one month visit to Carnegie Mellon in July, 1997, by the NSF.

References

- [1] S. Abramsky. Computational interpretations of linear logic. *Theoretical Computer Science*, 111:3–57, 1993.
- [2] A. Barber and G. D. Plotkin. Dual intuitionistic linear logic. Technical Report, LFCS, Univ of Edinburgh, October 1997.
- [3] N.D. Belnap. Display logic. *Journal of Philosophical Logic*, 11:375–414, 1982.
- [4] P. N. Benton. A mixed linear and non-linear logic: proofs, terms and models. Proceedings of *Computer Science Logic '94*, Kazimierz, Poland. Springer-Verlag LNCS 933, 1995.
- [5] P. N. Benton, G. M. Bierman, V. C. V. de Paiva, and J. M. E. Hyland. Linear λ -calculus and categorical models revisited. In E. Börger et al., editors, *Proceedings of the Sixth Workshop on Computer Science Logic*, volume 702 of *Lecture Notes in Computer Science*, pages 61–84. Springer-Verlag, Berlin, 1992.
- [6] J. Berdine, P. W. O’Hearn, U. S. Reddy, and H. Thielecke. Linear continuation-passing. *Journal of Functional Programming*, to appear. Preliminary version appeared in the *2001 ACM SIGPLAN Workshop on Continuations (CW’01)*, 2002.
- [7] G. M. Bierman. What is a categorical model of intuitionistic linear logic? In *International Conference on Typed Lambda Calculi and Applications*, volume 902 of *Lecture Notes in Computer Science*. Springer, 1995.
- [8] P. Brinch Hansen. *Operating System Principles*. Prentice Hall, New Jersey, 1973. Series in Automatic Computation.

- [9] S. Brookes, M. Main, A. Melton, and M. Mislove, editors. *Mathematical Foundations of Programming Semantics, Eleventh Annual Conference*, volume 1 of *Electronic Notes in Theoretical Computer Science*, Tulane University, New Orleans, Louisiana, March 29–April 1 1995. Elsevier Science.
- [10] L. Cardelli, P. Gardner, and G. Ghelli. A spatial logic for querying graphs. In *ICALP'02*, LNCS. Springer-Verlag, 2002.
- [11] L. Cardelli and A. D. Gordon. Anytime, anywhere. Modal logics for mobile ambients. In *Conference Record of the 27th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, Boston, Massachusetts, 2000. ACM, New York.
- [12] B. J. Day. On closed categories of functors. In S. Mac Lane, editor, *Reports of the Midwest Category Seminar*, volume 137 of *Lecture Notes in Mathematics*, pages 1–38. Springer-Verlag, Berlin-New York, 1970.
- [13] J. M. Dunn. Relevant logic and entailment. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic, vol. III: Alternatives in Classical Logic*, number 166 in Synthese Library, pages 117–224. D. Reidel, Dordrecht, Holland, 1986.
- [14] F. Foltz, C. Lair, and G. M. Kelly. Algebraic categories with few monoidal biclosed structures or none. *J. Pure and Applied Algebra*, 17:171–177, 1980.
- [15] P. J. Freyd, P. W. O’Hearn, A. J. Power, R. Street, M. Takeyama, and R. D. Tennent. Bireflectivity. *Theoretical Computer Science*, 228(1-2):49–76, October 1999. Preliminary version in [9].
- [16] J.-Y. Girard. Linear logic. *Theoretical Computer Science*, pages 1–102, 1987.
- [17] E. Goubault. *The Geometry of Concurrency*. PhD thesis, 1995.
- [18] J. W. Gray. *Formal Category Theory – Adjointness for 2-Categories*, volume 391 of *Lecture Notes in Math*. Springer, 1974.
- [19] M. Hofmann. *Type Systems for Polynomial Time Complexity*. Habilitation thesis, Darmstadt. Available as Edinburgh report ECS-LFCS-99-406, 1999.
- [20] S. Ishtiaq and P. W. O’Hearn. BI as an assertion language for mutable data structures. In *28th POPL*, pages 14–26, London, January 2001.
- [21] S.S. Ishtiaq and D. J. Pym. A relevant analysis of natural deduction. *Journal of Logic and Computation*, 8(6):809–838, 1998.
- [22] Y. Lafont. The linear abstract machine. *Theoretical Computer Science*, 59:157–180, 1988.
- [23] J. M. Lucassen and D. K. Gifford. Polymorphic effect systems. In *Conference Record of the Fifteenth Annual ACM Symposium on Principles of Programming Languages*. ACM, New York, 1988.
- [24] S. Mac Lane. *Categories for the Working Mathematician*. Springer-Verlag, New York, 1971.
- [25] P. O’Hearn, J. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In L. Fribourg, editor, *CSL 2001*, pages 1–19. Springer-Verlag. LNCS 2142.

- [26] P. W. O’Hearn. Resource interpretations, bunched implications and the $\alpha\lambda$ -calculus (preliminary version). In *Typed λ -calculus and Applications*, J-Y Girard editor, LNCS 1581, 1999.
- [27] P. W. O’Hearn. *The Semantics of Non-Interference: A Natural Approach*. Ph.D. thesis, Queen’s University, Kingston, Canada, 1990.
- [28] P. W. O’Hearn. Linear logic and interference control. In D. H. Pitt et al., editors, *Category Theory and Computer Science*, volume 530 of *Lecture Notes in Computer Science*, pages 74–93, Paris, France, September 1991. Springer-Verlag, Berlin.
- [29] P. W. O’Hearn. A model for syntactic control of interference. *Mathematical Structures in Computer Science*, 3(4):435–465, 1993.
- [30] P. W. O’Hearn, A. J. Power, M. Takeyama, and R. D. Tennent. Syntactic control of interference revisited. *Theoretical Computer Science*, 228(1-2):211–252, October 1999. Preliminary version in [9] and in [35], vol 2.
- [31] P. W. O’Hearn and D. J. Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2):215–244, June 99.
- [32] P. W. O’Hearn and U. S. Reddy. Objects, interference and the Yoneda embedding. *Theoretical Computer Science*, 228(1-2):253–282, October 1999. Preliminary version in [9].
- [33] P. W. O’Hearn and J. C. Reynolds. From Algol to polymorphic linear lambda-calculus. *J. ACM*, 47(1):167–223, January 2000.
- [34] P. W. O’Hearn and R. D. Tennent. Semantics of local variables. In M. P. Fourman, P. T. Johnstone, and A. M. Pitts, editors, *Applications of Categories in Computer Science*, volume 177 of *London Mathematical Society Lecture Note Series*, pages 217–238. Cambridge University Press, Cambridge, England, 1992.
- [35] P. W. O’Hearn and R. D. Tennent, editors. *Algol-like Languages*. Two volumes, Birkhauser, Boston, 1997.
- [36] F. J. Oles. *A Category-Theoretic Approach to the Semantics of Programming Languages*. Ph.D. thesis, Syracuse University, Syracuse, N.Y., 1982.
- [37] J. Polakow and F. Pfenning. Natural deduction for intuitionistic non-commutative linear logic. In *Typed λ -calculus and Applications*, J-Y Girard editor, LNCS 1581., 1999.
- [38] D. J. Pym. *The Semantics and Proof Theory of the Logic of Bunched Implications*. Kluwer Academic Publishers, Boston/Dordrecht/London, 2002. To appear.
- [39] S. Read. *Relevant Logic: A Philosophical Examination of Inference*. Basil Blackwell, 1987.
- [40] U. S. Reddy. A linear logic model of state. Manuscript, Univ of Illinois, 1993.
- [41] C. Retoré. Pomset logic: A non-commutative extension of classical linear logic. In *Computer Science Logic*, Paderborn, 1995.
- [42] J. C. Reynolds. Intuitionistic reasoning about shared mutable data structure. To appear in the *Proceedings of the Symposium in Celebration of the Work of C.A.R. Hoare*, 2000.

- [43] J. C. Reynolds. Syntactic control of interference. In *Conference Record of the Fifth Annual ACM Symposium on Principles of Programming Languages*, pages 39–46, Tucson, Arizona, January 1978. ACM, New York. Also in [35], vol 1.
- [44] J. C. Reynolds. The essence of Algol. In J. W. de Bakker and J. C. van Vliet, editors, *Algorithmic Languages*, pages 345–372, Amsterdam, October 1981. North-Holland, Amsterdam. Also in [35], vol 1, pages 67–88.
- [45] J. C. Reynolds. Syntactic control of interference, part 2. In G. Ausiello, M. Dezani-Ciancaglini, and S. Ronchi Della Rocca, editors, *Automata, Languages and Programming, 16th International Colloquium*, volume 372 of *Lecture Notes in Computer Science*, pages 704–722, Stresa, Italy, July 1989. Springer-Verlag, Berlin.
- [46] P. Ruet and F. Fages. Concurrent constraint programming and noncommutative logic. In *Computer Science Logic'97*, LNCS. Springer-Verlag, 1998.
- [47] P. Schroeder-Heister. Structural frameworks, substructural logics and the role of elimination inferences. In G. Huet and G. Plotkin, editors, *Logical Frameworks*, pages 385–403. Cambridge University Press, 1991.
- [48] R. A. G. Seely. Linear logic, *-autonomous categories and cofree coalgebras. In J. W. Gray and A. Scedrov, editors, *Categories in Computer Science and Logic*, volume 92 of *Contemporary Mathematics*, pages 371–382, Providence, Rhode Island, 1989. American Mathematical Society.
- [49] P. Wadler. Is there a use for linear logic? In *ACM/IFIP Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, 1991. Proceedings of the 1991 Conference.
- [50] D. Walker and G. Morrisett. Alias types for recursive data structures. *Workshop on Types in Compilation*, Montreal, September 2000.