Type-Directed Concurrency

Deepak Garg and Frank Pfenning*

Carnegie Mellon University
{dg,fp}@cs.cmu.edu

Abstract. We introduce a novel way to integrate functional and concurrent programming based on intuitionistic linear logic. The functional core arises from interpreting proof reduction as computation. The concurrent core arises from interpreting proof search as computation. The two are tightly integrated via a monad that permits both sides to share the same logical meaning for the linear connectives while preserving their different computational paradigms. For example, concurrent computation synthesizes proofs which can be evaluated as functional programs. We illustrate our design with some small examples, including an encoding of the pi-calculus.

1 Introduction

At the core of functional programming lies the beautiful Curry-Howard isomorphism which identifies intuitionistic proofs with functional programs and propositions with types. In this paradigm, computation arises from proof reduction. One of the most striking consequences is that we can write functions and reason logically about their behavior in an integrated manner.

Concurrent computation has resisted a similarly deep, elegant, and practical analysis with logical tools, despite several explorations in this direction (see, for example, [4, 3]). We believe the lack of a satisfactory Curry-Howard isomorphism is due to the limits inherent in complete proofs: they provide an analysis of constructive truth but not of the dynamics of interaction.

An alternative logical foundation for concurrency is to view computation as proof search [6]. In this paper we show that the two views of computation, via proof reduction and via proof search, are not inherently incompatible, but can coexist harmoniously in a language that combines functional and concurrent computation. We retain the strong guarantees that can be made for functional computation without unduly restricting the dynamism of concurrent computation.

In order to achieve this synthesis, we employ several advanced building blocks. The first is linearity: as has been observed [12], the evolution and communication of processes maps naturally to the single-use semantics of assumptions in linear logic. The second is dependency: we use dependent types to model communication channels and also to retain the precision of functional specifications

* This work has been partially supported by NSF Grant CCR-0306313 Efficient Logical Frameworks.

for transmitted values. The third is monads: we use monadic types to encapsulate concurrent computation, so that the linear connectives can retain the same logical meaning on the functional and concurrent side without interference. The fourth is focusing: we use it to enforce the atomicity of concurrent interactions during proof search.

The result is a tightly integrated language in which functional computation proceeds by reduction and concurrent computation proceeds by proof search. Concurrent computation thereby synthesizes proofs which can be evaluated as functional programs. We illustrate the design with some small examples, including an encoding of the π-calculus to help gauge its expressive power.

The principal contributions of this paper are conceptual and foundational. Details of the implementation of a static and dynamic semantics are left to future work, although a simple prototype [1] and experience with some prior languages [20] indicates that there are feasible solutions to implementation challenges such as type reconstruction.

There has been significant prior work in combining functional and concurrent programming. One class of languages, including Facile [11], Concurrent ML [16, 17], JOCaml [9], and Concurrent Haskell [13], adds concurrency primitives to a language with functional abstractions. While we share some ideas (such as the use of monadic encapsulation in Concurrent Haskell), the concurrent features in these languages are motivated operationally rather than logically and are only faintly reflected in the type system. Another class of languages start from a rich concurrent formalism such as the π-calculus and either add or encode some features of functional languages [15]. While operationally adequate, these encodings generally do not have a strong logical component. An interesting intermediate point is the applied π-calculus [2] where algebraic equations are added to the π-calculus. However, it is intended for reasoning about specifications rather than as a programming language.

In the remainder of the paper we present our language (which we call CLL) in three steps. First, we present the functional core which integrates linearity and a monad. Second, we present the concurrent core, which is based on proof search, and which can call upon functional computation. Third, we complete the integration with one additional construct to facilitate the other direction. We conclude with some further remarks about related and future work.

2 fCLL: Functional core of CLL

Syntax. The functional core of CLL is a first-order dependently typed linear functional language called *functional* CLL or fCLL. It is an extension of a linear lambda calculus with first-order dependent types from DML [20] and a monad similar to the concurrent logical framework [19]. The syntax of fCLL is summarized in figure 1. Types in fCLL can depend on index terms (denoted by s, t) that are divided into a number of disjoint sorts (γ). Index terms contain index variables ($i, j, k, \ldots$) and uninterpreted function symbols ($f, g, h, \ldots$). We assume the existence of a *sorting* judgment $\Sigma \vdash t \in \gamma$, where Σ is a context that mentions the sorts of all free index variables in t.

Sorts	$\gamma ::= \texttt{chan} \mid \ldots$
Index terms	$s, t ::= i \mid f(t_1, \ldots, t_n)$
Index variable contexts	$\Sigma ::= \cdot \mid \Sigma, i : \gamma$
Sorting judgment	$\Sigma \vdash t \in \gamma$
Kinds	$K ::= \texttt{Type} \mid \gamma \to K$
Types	$T ::= A \mid S$
Asynchronous types	$A, B ::= C\ t_1 \ldots t_n \mid A \& B \mid A \to B \mid A \multimap B \mid \forall i : \gamma.A(i) \mid \{S\}$
Synchronous types	$S ::= A \mid S_1 \otimes S_2 \mid \mathbf{1} \mid !A \mid \exists i : \gamma.S(i)$
Programs	$P ::= N \mid M \mid E$
Terms	$N ::= x \mid \langle N_1, N_2 \rangle \mid \pi_1\ N \mid \pi_2\ N \mid \lambda x : A.N \mid N_1\ N_2 \mid$ $\hat{\lambda} x : A.N \mid N_1\ \hat{}\ N_2 \mid \Lambda i : \gamma.N \mid N\ [t] \mid \{E\}$
Monadic-terms	$M ::= N \mid M_1 \otimes M_2 \mid \star \mid !N \mid [t, M]$
Patterns	$p ::= x \mid \star \mid p_1 \otimes p_2 \mid !x \mid [i, p]$
Expressions	$E ::= M \mid \underline{\texttt{let}}\ \{p : S\} = N\ \underline{\texttt{in}}\ E$

Fig. 1. fCLL syntax

Type constructors (denoted by C) are classified into kinds. For every fCLL program we assume the existence of an implicit signature that mentions the kinds of all type constructors used in the program. An atomic type is formed by applying a type constructor C to index terms $t_1, \ldots, t_n$. If C has kind $\gamma_1 \to \ldots \to \gamma_n \to \texttt{Type}$, we say that the atomic type $C\ t_1 \ldots\ t_n$ is well-formed in the index variable context Σ iff for $1 \leq i \leq n$, $\Sigma \vdash t_i \in \gamma_i$. In the following we assume that all atomic types in fCLL programs are well-formed.

Types in fCLL are divided into two classes - asynchronous (A, B) and synchronous (S).[1] Asynchronous types can be freely used as synchronous types. However, synchronous types must be coerced explicitly into asynchronous types using a monad $\{\ldots\}$, which is presented in a judgmental style [**?**].

Programs (P) are divided into three syntactic classes – terms (N), monadic-terms (M) and expressions (E). Under the Curry-Howard isomorphism, terms are proofs of asynchronous types whereas monadic-terms and expressions are proofs of synchronous types that end with *introduction* rules and *elimination* rules respectively. A fCLL program is called closed if it does not contain any free term variables. Closed programs may contain free index variables.

Typing. Programs in fCLL are type-checked using four contexts – a context of index variables Σ, a context of linear variables Δ, a context of unrestricted variables Γ and a context of patterns Ψ. Only the last of these contexts is ordered. There are four typing judgments in the type system. Some interesting rules from these judgments are shown in figure 2. The remaining rules for these judgments are presented in Appendix A. Type-checking for fCLL is decidable.

Operational Semantics. We use a call-by-value reduction semantics for fCLL. Figure 3 shows the definition of values in fCLL and some interesting reduction

[1] This terminology comes from Andreoli's work [5] on focused proofs in linear logic.

$$\Sigma ::= \cdot \mid \Sigma, i : \gamma \qquad \Delta ::= \cdot \mid \Delta, x : A$$
$$\Gamma ::= \cdot \mid \Gamma, x : A \qquad \Psi ::= \cdot \mid \Psi, p : S$$

$\boxed{\mathbf{\Sigma; \Gamma; \Delta \ \vdash \ N : A}}$

$$\frac{}{\Sigma; \Gamma; x : A \ \vdash \ x : A}\ \textsf{Hyp1} \qquad \frac{}{\Sigma; \Gamma, x : A; \cdot \ \vdash \ x : A}\ \textsf{Hyp2} \qquad \frac{\Sigma; \Gamma; \Delta \ \vdash \ E \div S}{\Sigma; \Gamma; \Delta \ \vdash \ \{E\} : \{S\}}\ \{\}\textsf{I}$$

$$\frac{\Sigma; \Gamma, x : A; \Delta \ \vdash \ N : B}{\Sigma; \Gamma; \Delta \ \vdash \ \lambda x : A.N : A \to B}\ \to\textsf{I} \qquad \frac{\Sigma; \Gamma; \Delta, x : A \ \vdash \ N : B}{\Sigma; \Gamma; \Delta \ \vdash \ \hat{\lambda} x : A.N : A \multimap B}\ \multimap\textsf{I}$$

$$\frac{\Sigma, i : \gamma; \Gamma; \Delta \ \vdash \ N : A}{\Sigma; \Gamma; \Delta \ \vdash \ \Lambda i : \gamma.N : \forall i : \gamma.A}\ \forall\textsf{I} \qquad \frac{\Sigma; \Gamma; \Delta \ \vdash \ N_1 : A_1 \qquad \Sigma; \Gamma; \Delta \ \vdash \ N_2 : A_2}{\Sigma; \Gamma; \Delta \ \vdash \ \langle N_1, N_2 \rangle : A_1 \& A_2}\ \&\textsf{I}$$

$\boxed{\mathbf{\Sigma; \Gamma; \Delta \ \vdash \ M \mathrel{\dot{\div}} S}}$

$$\frac{\Sigma; \Gamma; \cdot \ \vdash \ N : A}{\Sigma; \Gamma; \cdot \ \vdash \ !N \mathrel{\dot{\div}} !A}\ !\textsf{R} \qquad \frac{\Sigma; \Gamma; \Delta \ \vdash \ M \mathrel{\dot{\div}} S(t) \qquad \Sigma \ \vdash \ t \in \gamma}{\Sigma; \Gamma; \Delta \ \vdash \ [t, M] \mathrel{\dot{\div}} \exists i : \gamma.S(i)}\ \exists\textsf{R}$$

$$\frac{\Sigma; \Gamma; \Delta_1 \ \vdash \ M_1 \mathrel{\dot{\div}} S_1 \qquad \Sigma; \Gamma; \Delta_2 \ \vdash \ M_2 \mathrel{\dot{\div}} S_2}{\Sigma; \Gamma; \Delta_1, \Delta_2 \ \vdash \ M_1 \otimes M_2 \mathrel{\dot{\div}} S_1 \otimes S_2}\ \otimes\textsf{R}$$

$\boxed{\mathbf{\Sigma; \Gamma; \Delta \ \vdash \ E \div S}}$

$$\frac{\Sigma; \Gamma; \Delta_1 \ \vdash \ N : \{S\} \qquad \Sigma; \Gamma; \Delta_2; p : S \ \vdash \ E \div S'}{\Sigma; \Gamma; \Delta_1, \Delta_2 \ \vdash \ \underline{\texttt{let}}\ \{p : S\} = N\ \underline{\texttt{in}}\ E \div S'}\ \{\}\textsf{E}$$

$\boxed{\mathbf{\Sigma; \Gamma; \Delta; \Psi \ \vdash \ E \div S}}$

$$\frac{\Sigma; \Gamma; \Delta \ \vdash \ E \div S}{\Sigma; \Gamma; \Delta; \cdot \ \vdash \ E \div S}\ \div\div \qquad \frac{\Sigma; \Gamma, x : A; \Delta; \Psi \ \vdash \ E \div S}{\Sigma; \Gamma; \Delta; !x : !A, \Psi \ \vdash \ E \div S}\ !\textsf{L} \qquad \frac{\Sigma; \Gamma; \Delta; \Psi \ \vdash \ E \div S}{\Sigma; \Gamma; \Delta; \star : \mathbf{1}, \Psi \ \vdash \ E \div S}\ \textsf{1L}$$

$$\frac{\Sigma; \Gamma; \Delta; p_1 : S_1, p_2 : S_2, \Psi \ \vdash \ E \div S}{\Sigma; \Gamma; \Delta; p_1 \otimes p_2 : S_1 \otimes S_2, \Psi \ \vdash \ E \div S}\ \otimes\textsf{L} \qquad \frac{\Sigma, i : \gamma; \Gamma; \Delta; p : S', \Psi \ \vdash \ E \div S}{\Sigma; \Gamma; \Delta; [i, p] : \exists i : \gamma.S', \Psi \ \vdash \ E \div S}\ \exists\textsf{L}\ (i \textsf{ fresh})$$

Fig. 2. fCLL type system (selected rules)

rules. The substitution relation $P[M_V/p]$ substitutes the monadic-value M_V for a pattern p in the program P. It is defined by induction on the pattern p. $P[V/x]$ and $P[t/i]$ are the usual capture avoiding substitutions for term and index variables respectively. There are some differences from a standard call-by-value reduction semantics. In fCLL, pairs of the form $\langle N_1, N_2 \rangle$ are values because N_1 and N_2 share linear resources and cannot be evaluated simultaneously. The monad $\{E\}$ is also a value because after we extend the language in section 4, expressions have effects. Reduction of the two components of a $\otimes$ can be interleaved arbitrarily, or it may performed in parallel.

Expressions are reduced in a context of index variables Σ. This context plays no role in fCLL, but when we extend fCLL to full-CLL in section 4, the context Σ becomes computationally significant. The complete set of reduction rules for fCLL may be found in Appendix B. We state preservation and progress theorems for fCLL below. The proofs of these theorems are relatively standard. See the associated technical report [10] for more details.

Theorem 1 (Preservation for fCLL).

1. If $\Sigma; \Gamma; \Delta \ \vdash \ N : A$ and $N \rightsquigarrow N'$, then $\Sigma; \Gamma; \Delta \ \vdash \ N' : A$.

Term values $V ::= \lambda x : A.N \mid \hat{\lambda} x : A.N \mid \{E\} \mid \langle N_1, N_2 \rangle \mid \Lambda i : \gamma.N$
Monadic values $M_V ::= V \mid M_{V_1} \otimes M_{V_2} \mid \star \mid !V \mid [t, M_V]$
Expression values $E_V ::= M_V$

$\boxed{\mathbf{P[M_V/p]}}$

$$P[\star/\star] = P \qquad P[[t, M_V]/[i,p]] = (P[t/i])[M_V/p]$$
$$P[!V/!x] = P[V/x] \qquad P[M_{V_1} \otimes M_{V_2}/p_1 \otimes p_2] = (P[M_{V_1}/p_1])[M_{V_2}/p_2]$$

$\boxed{\mathbf{N \rightsquigarrow N'}}$

$$\frac{}{\pi_1\ \langle N_1, N_2\rangle \rightsquigarrow N_1} \rightsquigarrow \langle\rangle\pi_1 \qquad \frac{}{\pi_2\ \langle N_1, N_2\rangle \rightsquigarrow N_2} \rightsquigarrow \langle\rangle\pi_2 \qquad \frac{}{(\Lambda i : \gamma.N)\ [t] \rightsquigarrow N[t/i]} \rightsquigarrow \Lambda APP$$

$$\frac{}{(\lambda x : A.N)\ V \rightsquigarrow N[V/x]} \rightsquigarrow \lambda APP \qquad \frac{}{(\hat{\lambda} x : A.N)\ \hat{}\ V \rightsquigarrow N[V/x]} \rightsquigarrow \hat{\lambda} LAPP$$

$\boxed{\mathbf{M \mapsto M'}}$

$$\frac{N \rightsquigarrow N'}{N \mapsto N'} \rightsquigarrow\mapsto \qquad \frac{N \rightsquigarrow N'}{!N \mapsto !N'} \mapsto ! \qquad \frac{M \mapsto M'}{[t, M] \mapsto [t, M']} \mapsto \exists$$

$$\frac{M_1 \mapsto M_1'}{M_1 \otimes M_2 \mapsto M_1' \otimes M_2} \mapsto \otimes_1 \qquad \frac{M_2 \mapsto M_2'}{M_1 \otimes M_2 \mapsto M_1 \otimes M_2'} \mapsto \otimes_2$$

$\boxed{\mathbf{\Sigma; E \hookrightarrow \Sigma; E'}}$

$$\frac{M \mapsto M'}{\Sigma; M \hookrightarrow \Sigma; M'} \mapsto\hookrightarrow \qquad \frac{}{\Sigma; \underline{\texttt{let}}\ \{p : S\} = \{M_V\}\ \underline{\texttt{in}}\ E \hookrightarrow \Sigma; E[M_V/p]} \hookrightarrow LETRED$$

$$\frac{N \rightsquigarrow N'}{\Sigma; \underline{\texttt{let}}\ \{p : S\} = N\ \underline{\texttt{in}}\ E \hookrightarrow \Sigma; \underline{\texttt{let}}\ \{p : S\} = N'\ \underline{\texttt{in}}\ E} \hookrightarrow LET_1$$

$$\frac{\Sigma; E \hookrightarrow \Sigma; E'}{\Sigma; \underline{\texttt{let}}\ \{p : S\} = \{E\}\ \underline{\texttt{in}}\ E_1 \hookrightarrow \underline{\texttt{let}}\ \Sigma; \{p : S\} = \{E'\}\ \underline{\texttt{in}}\ E_1} \hookrightarrow LET_2$$

Fig. 3. fCLL operational semantics (selected rules)

2. If $\Sigma; \Gamma; \Delta \vdash M \mathrel{\dot{\underset{\cdot}{\sim}}} S$ and $M \rightsquigarrow M'$, then $\Sigma; \Gamma; \Delta \vdash M' : S$.
3. If $\Sigma; \Gamma; \Delta \vdash E \div S$ and $\Sigma; E \hookrightarrow \Sigma; E'$, then $\Sigma; \Gamma; \Delta \vdash E' \div S$.

Theorem 2 (Progress for fCLL).
1. If $\Sigma; \cdot; \cdot \vdash N : A$ then either $N = V$ or $N \rightsquigarrow N'$ for some N'.
2. If $\Sigma; \cdot; \cdot \vdash M \mathrel{\dot{\underset{\cdot}{\sim}}} S$ then either $M = M_V$ or $M \mapsto M'$ for some M'.
3. If $\Sigma; \cdot; \cdot \vdash E \div S$ then either $E = E_V$ or $\Sigma; E \hookrightarrow \Sigma; E'$ for some E'.

Example 1 (Fibonacci numbers). As a simple example of programming in fCLL, we describe a function for computing Fibonacci numbers. These numbers are defined inductively as follows.

$$fib(0) = fib(1) = 1 \qquad fib(n) = fib(n-1) + fib(n-2)$$

For implementing this definition as a function in fCLL, we assume that fCLL terms have been extended with integers having type `int`, named recursive functions and a conditional <u>**`if`**</u>-<u>**`then`**</u>-<u>**`else`**</u> construct. These can be added to fCLL in a straightforward manner. Figure 4 shows the fCLL function `fib` that computes the nth Fibonacci number. It has the type `int` $\rightarrow$ {!`int`}. It is possible to write this function in a manner simpler than the one presented here, but we write it this way to highlight specific features of fCLL.

```
fib = λn : int.
   if (n = 0 or n = 1) then {!1}
   else
   {
      let {!n₁} = fib (n − 1) in
      let {!n₂} = fib (n − 2) in
         !(n₁ + n₂)
   }
```

Fig. 4. The function `fib` in fCLL

The most interesting computation in `fib`, including recursive calls, occurs inside the monad. Since the monad is evaluated lazily in fCLL, computation in `fib` will actually occur only when the caller of `fib` eliminates the monad from the returned value of type $\{!\texttt{int}\}$. Syntactically, elimination of the monadic constructor can occur only in expressions at the <u>`let`</u> construct. Hence the program that calls `fib` must be an expression. Here is an example of such a top level program that prints the 5th Fibonacci number: <u>`let`</u> $\{!x\} = \texttt{fib}\ 5$ <u>`in`</u> $\texttt{print}(x)$

3 *l*CLL: Concurrent core of CLL

The concurrent core of CLL is called lCLL. It adds a layer of concurrency over the functional language fCLL. In the structure of concurrent computations lCLL is similar to the π-calculus. However it is different in other respects. First, it embeds a functional language (fCLL) directly, resulting in a direct representation of functional computation inside concurrent ones, as opposed to the use of complex encodings for doing the same in the π-calculus [18]. Second, the semantics of lCLL are directed by types, not terms. This, we believe, is a new idea that has not been explored before.

Syntax. We present lCLL as a chemical abstract machine (CHAM) [7]. lCLL programs are called configurations, denoted by $\mathcal{C}$. Figure 5 shows the syntax of lCLL configurations. Each configuration is made of four components, written $\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}$. Σ is a context of index variables, as defined in section 2. $\hat{\sigma}$ is a sorted substitution mapping index variables to index terms. $\hat{\Gamma}$ is a multiset of closed fCLL term values along with their types. $\hat{\Delta}$ is a multiset of closed fCLL programs together with their types. We require that whenever $N : A \in \hat{\Delta}$, N have the type $A[\hat{\sigma}]$, where $A[\hat{\sigma}]$ is the result of applying the substitution $\hat{\sigma}$ to the type A. Similar conditions hold for monadic-terms and expressions in $\hat{\Delta}$ and term values in $\hat{\Gamma}$. Formally, a configuration $\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}$ is said to be well-formed if it satisfies the following conditions.

1. If $(t/i : \gamma) \in \hat{\sigma}$, then $i \notin \texttt{dom}(\Sigma)$ and $\Sigma \vdash t \in \gamma$.
2. If P is a program in $\hat{\Gamma}$ or $\hat{\Delta}$, then $\texttt{fv}(P) \cap \texttt{dom}(\hat{\sigma}) = \phi$.
3. If $V : A \in \hat{\Gamma}$, then $\Sigma; \cdot; \cdot \vdash V : A[\hat{\sigma}]$.
4. If $N : A \in \hat{\Delta}$, then $\Sigma; \cdot; \cdot \vdash N : A[\hat{\sigma}]$.
5. If $M \div\!\!\!\!\sim S \in \hat{\Delta}$, then $\Sigma; \cdot; \cdot \vdash M \div\!\!\!\!\sim S[\hat{\sigma}]$.
6. If $E \div S \in \hat{\Delta}$, then $\Sigma; \cdot; \cdot \vdash E \div S[\hat{\sigma}]$.

Configurations	$\mathcal{C} ::= \Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}$
Global index names	$\Sigma ::= \cdot \mid \Sigma, i : \gamma$
Local name substitutions	$\hat{\sigma} ::= \cdot \mid \hat{\sigma}, t/i : \gamma$
Unrestricted solutions	$\hat{\Gamma} ::= \cdot \mid \hat{\Gamma}, V : A$
Linear solutions	$\hat{\Delta} ::= \cdot \mid \hat{\Delta}, N : A \mid \hat{\Delta}, M \div S \mid \hat{\Delta}, E \div S$

Fig. 5. lCLL syntax

$$\frac{N \rightsquigarrow N'}{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, N : A \twoheadrightarrow \Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, N' : A} \twoheadrightarrow\rightsquigarrow$$

$$\frac{M \mapsto M'}{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, M \div S \twoheadrightarrow \Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, M' \div S} \twoheadrightarrow\mapsto$$

$$\frac{\Sigma; E \hookrightarrow \Sigma; E'}{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, E \div S \twoheadrightarrow \Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, E' \div S} \twoheadrightarrow\hookrightarrow$$

Fig. 6. Functional rewrite rules for lCLL configurations

We assume that all our configurations are well-formed. Programs in $\hat{\Delta}$ and values in $\hat{\Gamma}$ are collectively called processes. Intuitively, we view programs in $\hat{\Delta}$ as concurrent processes that are executing simultaneously. $\hat{\Delta}$ is called a linear solution because these processes are single-use in the sense that they can neither be replicated, nor destroyed. Term values in $\hat{\Gamma}$ are viewed as irreducible processes (like functional abstractions) that are replicable. For this reason $\hat{\Gamma}$ is also called an unrestricted solution. The context Σ can be viewed as a set of global index names, that are known to have specific sorts. The domain of the substitution $\hat{\sigma}$ can be viewed as a set of local (private) index names that are created during the evaluation of the configuration. The substitution $\hat{\sigma}$ maps these local index names to index terms that depend only on the global names (see condition (1) for well-formedness above).

3.1 Semantics of *l*CLL

The semantics of lCLL are rewrite rules that allow a configuration to step to other configuration(s). The specific rules that apply to a particular configuration are determined by the *types* of processes in that configuration. In this sense, these rules are type-directed. We classify rewrite rules into three classes – functional, structural and synchronization.

Functional rules. Functional rules allow reduction of programs in the linear solution $\hat{\Delta}$. We denote them using the arrow $\twoheadrightarrow$. Figure 6 shows the functional rewrite rules for lCLL configurations. There are three rules, one for reducing programs in each of the three syntactic classes of fCLL. Reductions of different programs in $\hat{\Delta}$ can be performed in parallel. This supports the idea that programs in $\hat{\Delta}$ can be viewed as processes executing simultaneously.

Structural rules. Structural rules apply to those irreducible programs in $\hat{\Delta}$ that have synchronous types. These are exactly the monadic values M_V. A struc-

$$
\begin{array}{rcll}
\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, (M_{V_1} \otimes M_{V_2}) \mathrel{\hat{\div}} (S_1 \otimes S_2) & \rightharpoonup & \Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, M_{V_1} \mathrel{\hat{\div}} S_1, M_{V_2} \mathrel{\hat{\div}} S_2 & (\rightharpoonup \otimes) \\
\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, \star \mathrel{\hat{\div}} \mathbf{1} & \rightharpoonup & \Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta} & (\rightharpoonup \mathbf{1}) \\
\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, [t, M_V] \mathrel{\hat{\div}} \exists i : \gamma.S(i) & \rightharpoonup & \Sigma; \hat{\sigma}, t/i : \gamma \triangleright \hat{\Gamma} \parallel \hat{\Delta}, M_V \mathrel{\hat{\div}} S(i) & (\rightharpoonup \exists) \\
& & (i \text{ fresh}) & \\
\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, !V \mathrel{\hat{\div}} !A & \rightharpoonup & \Sigma; \hat{\sigma} \triangleright \hat{\Gamma}, V : A \parallel \hat{\Delta} & (\rightharpoonup !) \\
\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, V \mathrel{\hat{\div}} A & \rightharpoonup & \Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, V : A & (\rightharpoonup \hat{\div}) \\
\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, M_V \div S & \rightharpoonup & \Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, M_V \mathrel{\hat{\div}} S & (\rightharpoonup \div)
\end{array}
$$

Fig. 7. Structural rewrite rules for lCLL configurations

tural rule decomposes a monadic value into smaller monadic values. We denote structural rules with the arrow $\rightharpoonup$. All structural rules for rewriting lCLL configurations are shown in figure 7. Unlike most CHAMs, our structural rules are not reversible.

The rule $\rightharpoonup \otimes$ splits the monadic value $M_{V_1} \otimes M_{V_2}$ of type $S_1 \otimes S_2$ into two monadic values M_{V_1} and M_{V_2} of types S_1 and S_2 respectively. Intuitively, we can view $M_{V_1} \otimes M_{V_2}$ as a parallel composition of the processes M_{V_1} and M_{V_2}. The rule $\rightharpoonup \otimes$ splits this parallel composition into its components, allowing each component to rewrite separately.

In the rule $\rightharpoonup \exists$, there is a side condition that i must be fresh i.e. it must not occur anywhere except in $S(i)$. Some α-renaming may have to be performed to enforce this. In lCLL, the $\exists$ type acts as a local index name creator. The rule $\rightharpoonup \exists$ creates the new index name i and records the fact that i is actually bound to the index term t in the substitution $\hat{\sigma}$.

The rule $\rightharpoonup !$ moves a program of type $!A$ to the unrestricted solution, thus allowing multiple uses of this program. For this reason, the type $!A$ serves as a replication construct in lCLL. The rules $\rightharpoonup \hat{\div}$ and $\rightharpoonup \div$ change the type ascription for programs that have been coerced from one syntactic class to another.

Synchronization Rules. Synchronization rules act on values in $\hat{\Gamma}$ and $\hat{\Delta}$ having asynchronous types. These are exactly the term values V. Synchronization rules are denoted by the arrow $\longrightarrow$. Figure 8 shows the two synchronization rules. The rule $\longrightarrow \{\}$ eliminates the monadic constructor $\{\}$ from values $\{E\}$ of asynchronous type $\{S\}$.

The second rule $\longrightarrow\Longrightarrow$ performs synchronization of several term values at the same time. It uses an auxiliary judgment $\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta} \implies N : A$, which we call the sync judgment. The rules of this judgment are also shown in figure 8. The sync judgment links values in $\hat{\Gamma}$ and $\hat{\Delta}$ to form a more complex program N. We call this process synchronization. Synchronization uses values in $\hat{\Delta}$ exactly once, while those in $\hat{\Gamma}$ may be used zero or more times.

In the rule $\longrightarrow\Longrightarrow$ shown in figure 8, $\hat{\Delta}$ denotes a subset of the linear solution that participates in the synchronization. The remaining solution $\hat{\Delta}'$ is kept as is. Some backward reasoning is performed in the judgment $\Longrightarrow$ to produce the linked program N of type $\{S\}$. This is the essential point here – the result of a synchronization must be of type $\{S\}$.

Synchronization rules, $\mathbf{\Sigma}; \hat{\sigma} \triangleright \hat{\mathbf{\Gamma}} \parallel \hat{\mathbf{\Delta}} \longrightarrow \mathbf{\Sigma}; \hat{\sigma} \triangleright \hat{\mathbf{\Gamma}} \parallel \hat{\mathbf{\Delta}}'$

$$\frac{}{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, \{E\} : \{S\} \longrightarrow \Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, E \div S} \longrightarrow \{\}$$

$$\frac{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta} \Longrightarrow N : \{S\}}{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}, \hat{\Delta}' \longrightarrow \Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel N : \{S\}, \hat{\Delta}'} \longrightarrow\Longrightarrow$$

Sync judgment, $\mathbf{\Sigma}; \hat{\sigma} \triangleright \hat{\mathbf{\Gamma}} \parallel \hat{\mathbf{\Delta}} \Longrightarrow \mathbf{N} : \mathbf{A}$

$$\frac{}{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel V : A \Longrightarrow V : A} \Longrightarrow_{HYP1} \qquad \frac{}{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma}, V : A \parallel \cdot \Longrightarrow V : A} \Longrightarrow_{HYP2}$$

$$\frac{\Sigma \cup \mathtt{dom}(\hat{\sigma}) \vdash t \in \gamma \qquad \Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta} \Longrightarrow N : \forall i : \gamma.A(i)}{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta} \Longrightarrow N\ [t[\hat{\sigma}]] : A(t)} \Longrightarrow \forall$$

$$\frac{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta} \Longrightarrow N : A_1 \& A_2}{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta} \Longrightarrow \pi_1\ N : A_1} \Longrightarrow \&_1 \qquad \frac{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta} \Longrightarrow N : A_1 \& A_2}{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta} \Longrightarrow \pi_2\ N : A_2} \Longrightarrow \&_2$$

$$\frac{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}_1 \Longrightarrow N_1 : A \qquad \Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}_2 \Longrightarrow N_2 : A \multimap B}{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}_1, \hat{\Delta}_2 \Longrightarrow N_2 \ \hat{}\ N_1 : B} \Longrightarrow\multimap$$

$$\frac{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \cdot \Longrightarrow N_1 : A \qquad \Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta} \Longrightarrow N_2 : A \rightarrow B}{\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta} \Longrightarrow N_2\ N_1 : B} \Longrightarrow\rightarrow$$

Fig. 8. Synchronization rewrite rules for *l*CLL configurations

The semantic rewriting relation for *l*CLL is defined as $\rightrightarrows = \twoheadrightarrow \cup \rightharpoonup \cup \longrightarrow$. It satisfies the following type preservation theorem.

Theorem 3 (Preservation for *l*CLL). If $\mathcal{C}$ is a well-formed configuration and $\mathcal{C} \rightrightarrows \mathcal{C}'$, then $\mathcal{C}'$ is also well-formed.

Concurrent computation as proof search. Given a *l*CLL configuration $\mathcal{C} = \Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta}$, types in $\hat{\Delta}[\hat{\sigma}]$ and $\hat{\Gamma}[\hat{\sigma}]$ can be viewed as propositions that are *simultaneously true*, in a linear and unrestricted sense respectively. Using the Curry-Howard isomorphism, the corresponding programs in $\hat{\Delta}[\hat{\sigma}]$ and $\hat{\Gamma}[\hat{\sigma}]$ can be seen as specific proofs of these propositions. The sync judgment (figure 8) is actually a linear entailment judgment – if $\Sigma; \hat{\sigma} \triangleright \hat{\Gamma} \parallel \hat{\Delta} \Longrightarrow N : A$, then from the unrestricted assumptions in $\hat{\Gamma}[\hat{\sigma}]$ and linear assumptions in $\hat{\Delta}[\hat{\sigma}]$, $A[\hat{\sigma}]$ can be proved in linear logic. The term N synthesized by this judgment is a proof of the proposition $A[\hat{\sigma}]$. As a result, each use of the synchronization rule $\longrightarrow\Longrightarrow$ can be viewed as a step of proof search in linear logic that uses several known facts to conclude a new fact, together with its proof term. By the Curry-Howard isomorphism, the proof term is a well-typed program that can be functionally reduced again.

More specifically, each use of $\longrightarrow\Longrightarrow$ corresponds to a single focusing step for eliminating asynchronous constructors from a proposition that has $\{S\}$ in the head position. For a detailed description of this see [10].

Additional Signature

$\texttt{procid}$: sort
$serv : \texttt{procid}$
$\texttt{mess}$: $\texttt{procid} \rightarrow \texttt{procid} \rightarrow \texttt{Type}$
$\texttt{message}$: $\forall i : \texttt{procid}.\ \forall j : \texttt{procid}.\ \texttt{int} \rightarrow \texttt{mess}\ i\ j$
$\texttt{fetchmessage}$: $\forall i : \texttt{procid}.\ \forall j : \texttt{procid}.\ \texttt{mess}\ i\ j \multimap \{!\texttt{int}\}$
$\texttt{fetchmessage}\ [i]\ [j]\ \hat{}\ (\texttt{message}\ [i]\ [j]\ n) \ \rightsquigarrow\ \{!n\}$

Fibonacci Server

$\texttt{fibservtype = } \forall i : \texttt{procid}.\ \texttt{mess}\ i\ serv \multimap \{\texttt{mess}\ serv\ i\}$
$\texttt{fibserver: fibservtype = } \Lambda i : \texttt{procid}.\ \hat{\lambda} m : \texttt{mess}\ i\ serv.$

```
{
    let {!n} = fetchmessage [i] [serv] ^ m in
    let {!v} = fib (n) in
        (message [serv] [i] v)
}
```

Sample Execution ($\Sigma = serv : \texttt{procid},\ k : \texttt{procid}$)

$$\Sigma; \cdot \rhd \cdot \Vert\ \texttt{fibserver} : \texttt{fibservtype},\ (\texttt{message}\ [k]\ [serv]\ 6) : \texttt{mess}\ k\ serv$$
$$\longrightarrow \Sigma; \cdot \rhd \cdot \Vert\ \texttt{fibserver}\ [k]\ \hat{}\ (\texttt{message}\ [k]\ [serv]\ 6) : \{\texttt{mess}\ serv\ k\}$$
$$\twoheadrightarrow^* \Sigma; \cdot \rhd \cdot \Vert \begin{pmatrix} \{ \\ \quad \underline{\texttt{let}}\ \{!n\} = \texttt{fetchmessage}\ [k]\ [serv]\ \hat{}\ (\texttt{message}\ [k]\ [serv]\ 6)\ \underline{\texttt{in}} \\ \quad \underline{\texttt{let}}\ \{!v\} = \texttt{fib}\ (n)\ \underline{\texttt{in}}\ (\texttt{message}\ [serv]\ [k]\ v) \\ \} \end{pmatrix} : \{\texttt{mess}\ serv\ k\}$$
$$\longrightarrow \Sigma; \cdot \rhd \cdot \Vert \begin{pmatrix} \underline{\texttt{let}}\ \{!n\} = \texttt{fetchmessage}\ [k]\ [serv]\ \hat{}\ (\texttt{message}\ [k]\ [serv]\ 6)\ \underline{\texttt{in}} \\ \underline{\texttt{let}}\ \{!v\} = \texttt{fib}\ (n)\ \underline{\texttt{in}}\ (\texttt{message}\ [serv]\ [k]\ v) \end{pmatrix} \div \texttt{mess}\ serv\ k$$
$$\twoheadrightarrow^* \Sigma; \cdot \rhd \cdot \Vert\ (\underline{\texttt{let}}\ \{!v\} = \texttt{fib}\ (6)\ \underline{\texttt{in}}\ (\texttt{message}\ [serv]\ [k]\ v)) \div \texttt{mess}\ serv\ k$$
$$\twoheadrightarrow^* \Sigma; \cdot \rhd \cdot \Vert\ (\texttt{message}\ [serv]\ [k]\ 13) \div \texttt{mess}\ serv\ k$$
$$\rightharpoonup^2 \Sigma; \cdot \rhd \cdot \Vert\ (\texttt{message}\ [serv]\ [k]\ 13) : \texttt{mess}\ serv\ k$$

Fig. 9. Server for computing Fibonacci numbers in lCLL

Example 2 (Client-Server Communication). We illustrate concurrent programming in lCLL with an example of a client-server interaction. The server described here listens to client requests to compute Fibonacci numbers. Each request contains an integer n. Given a request, the server computes the nth Fibonacci number and returns this value to the client.

We model communication through asynchronous message passing. Assume that all clients and the server have unique identities, which are index names from a special sort called `procid`. The identity of the server is $serv$. A message from one process to another contains three parts – the identity of the sender, the identity of the recipient and an integer, which is the content of the message. Messages are modeled using a type constructor `mess` and a term constructor `message` having the kind and type shown in figure 9. For every pair of index terms i and j of sort `procid` and every integer n, we view the value $(\texttt{message}\ [i]\ [j]\ n)$ of type $(\texttt{mess}\ i\ j)$ as a message having content n from the process with identity i to the process with identity j. In order to extract the integer content of a message, we use the destructor `fetchmessage` that has the reduction rule $\texttt{fetchmessage}\ [i]\ [j]\ \hat{}\ (\texttt{message}\ [i]\ [j]\ n) \ \rightsquigarrow\ \{!n\}$.

Syntax

Actions	$A ::= \bar{x}y \mid x(y).P$	
External Choices	$C ::= A \mid C + C$	
Processes	$P, Q ::= C \mid !C \mid P\|P \mid \nu x.P \mid 0$	
Molecules	$m ::= P \mid \nu x.S$	
Solutions	$S ::= \phi \mid S \uplus \{m\}$	

Equations on terms and solutions

$$C_1 + (C_2 + C_3) = (C_1 + C_2) + C_3 \qquad \nu x.P = \nu y.P[y/x] \qquad (y \notin P)$$
$$C_1 + C_2 = C_2 + C_1 \qquad \nu x.S = \nu y.S[y/x] \qquad (y \notin S)$$

CHAM semantics

$$P_1|P_2 \rightleftharpoons P_1, P_2 \qquad x(y).P + C_1\ ,\ \bar{x}z + C_2 \rightarrow\ P[z/y]$$
$$0 \rightleftharpoons \qquad \nu x.P \rightleftharpoons \nu x.\{P\}$$
$$!C \rightleftharpoons !C, C \qquad (\nu x.P)|Q \rightleftharpoons \nu x.(P|Q) \qquad (x \notin Q)$$

Reduction semantics

$$P \equiv P' \ \Leftrightarrow\ P \rightleftharpoons^* P' \qquad P \rightarrow P' \ \Leftrightarrow\ P \rightleftharpoons^* \rightarrow \rightleftharpoons^* P'$$

Fig. 10. A variant of the asynchronous π-calculus

The server program called `fibserver` is shown in figure 9. It waits for a message m from any client i. Then it extracts the content n from the message, computes the nth Fibonacci number using the function `fib` defined in example 1 and returns this computed value to the client i as a message. `fibserver` has the type `fibservtype` $= \forall i :$ `procid`. `mess` $i\ serv \multimap \{$`mess` $serv\ i\}$.

A sequence of rewrite steps in lCLL using `fibserver` is shown in figure 9. The initial configuration contains `fibserver` and a message to `fibserver` containing the integer 6 from a client having identity k. For brevity, we omit the client process. The crucial rewrite in this sequence is the first one, where the synchronization rule $\longrightarrow\Longrightarrow$ is used to link the `fibserver` program with the message for it. Rewriting ends with a message containing the value of the 6th Fibonacci number (namely 13) from `fibserver` to the requesting client k.

3.2 An encoding of the π-calculus in lCLL

We describe a translation of a variant of the asynchronous π-calculus [8] to lCLL. The syntax and semantics of this variant are shown in figure 10. It extends the asynchronous π-calculus with an external choice operator + and a nil process 0. The replication operator ! is restricted to choices only.

Two translations $\ulcorner\cdot\urcorner$ and $\ulcorner\!\ulcorner\cdot\urcorner\!\urcorner$ are shown in figure 11. They map π-calculus entities to programs and types of *fCLL* respectively. We model channels as index terms of a specific sort `chan`. In order to translate $\bar{x}y$, which is an output message, we introduce a type constructor `out` and a related term constructor `output`, whose kind and type are shown in figure 11. The translations of $\bar{x}y$ to terms and types are `output` $[x]$ $[y]$ and `out` $x\ y$ respectively.

To translate $x(y).P$, we introduce a term destructor `destroyout` corresponding to the constructor `output`. Its type and reduction rule are shown in figure 11. The translation $\ulcorner x(y).P\urcorner$ waits for two inputs – the channel name y and a message m that corresponds to the translation of $\bar{x}y$. It then discards the message m and starts the process P.

Additional Signature

$\texttt{chan}$: sort
$\texttt{out}: \texttt{chan} \rightarrow \texttt{chan} \rightarrow \texttt{Type}$
$\texttt{output}: \forall x : \texttt{chan}.\ \forall y : \texttt{chan}.\ \texttt{out}\ x\ y$
$\texttt{destroyout}: \forall x : \texttt{chan}.\ \forall y : \texttt{chan}.\ \texttt{out}\ x\ y \multimap \{\mathbf{1}\}$
$\texttt{destroyout}\ [x]\ [y]\ \hat{}\ (\texttt{output}\ [x]\ [y]) \ \rightsquigarrow\ \{\star\}$
$c_{\texttt{chan}} : \texttt{chan}$

A/C/P	fCLL Type, $\ulcorner\!\ulcorner A/C/P \urcorner\!\urcorner$	fCLL Program, $\ulcorner A/C/P \urcorner$
$\bar{x}y$	$\texttt{out}\ x\ y$	$\texttt{output}\ [x]\ [y]$
$x(y).P$	$\forall y : \texttt{chan}.\ \texttt{out}\ x\ y \multimap \{\ulcorner\!\ulcorner P \urcorner\!\urcorner\}$	$\Lambda y : \texttt{chan}.\ \hat{\lambda} m : \texttt{out}\ x\ y.$ $\{$ $\underline{\texttt{let}}\ \{\star\} = \texttt{destroyout}\ [x]\ [y]\ \hat{}\ m$ $\underline{\texttt{in}}\ \ulcorner P \urcorner$ $\}$
$C_1 + C_2$	$\ulcorner\!\ulcorner C_1 \urcorner\!\urcorner \ \&\ \ulcorner\!\ulcorner C_2 \urcorner\!\urcorner$	$\langle \ulcorner C_1 \urcorner, \ulcorner C_2 \urcorner \rangle$
0	$\mathbf{1}$	$\star$
$!C$	$!\ \ulcorner\!\ulcorner C \urcorner\!\urcorner$	$!\ \ulcorner C \urcorner$
$P_1 \vert P_2$	$\ulcorner\!\ulcorner P_1 \urcorner\!\urcorner \otimes \ulcorner\!\ulcorner P_2 \urcorner\!\urcorner$	$\ulcorner P_1 \urcorner \otimes \ulcorner P_2 \urcorner$
$\nu x.P$	$\exists x : \texttt{chan}.\ulcorner\!\ulcorner P \urcorner\!\urcorner$	$[c_{\texttt{chan}},\ (\ulcorner P \urcorner[c_{\texttt{chan}}/x])]$

Fig. 11. Translation of the π-calculus

Translations of $C_1 + C_2$, $!C$, $P_1|P_2$ and 0 are straightforward. We translate $\nu x.P$ to the type $\exists x : \texttt{chan}.\ulcorner\!\ulcorner P \urcorner\!\urcorner$. To translate $\nu x.P$ to a program, we assume that there is an index constant $c_{\texttt{chan}}$ of sort $\texttt{chan}$. Then we translate $\nu x.P$ to $[c_{\texttt{chan}},\ (\ulcorner P \urcorner[c_{\texttt{chan}}/x])]$, which has the type $\exists x : \texttt{chan}.\ulcorner\!\ulcorner P \urcorner\!\urcorner$.

For any π-calculus process P, $\texttt{fn}(P) : \texttt{chan}; \cdot; \cdot \ \vdash\ \ulcorner P \urcorner \div \ulcorner\!\ulcorner P \urcorner\!\urcorner$. The translation of a π-calculus process P to lCLL is defined as the configuration $\langle P \rangle = \texttt{fn}(P) : \texttt{chan}; \cdot \rhd \cdot \Vert \ulcorner P \urcorner \div \ulcorner\!\ulcorner P \urcorner\!\urcorner$. Although we have not formally proved it, we believe that the following correctness result holds for this translation: $P \rightarrow^* P'$ iff there is a lCLL configuration $\mathcal{C}$ such that $\langle P \rangle \rightrightarrows^* \mathcal{C}$ and $\langle P' \rangle \rightharpoonup^* \mathcal{C}$.

4 Full-CLL: The complete language

Full-CLL is an extension of fCLL that allows lCLL's concurrent computations inside functional ones. This is done by extending fCLL expressions by a single construct – $\underline{\texttt{link}}\ E \div S\ \underline{\texttt{to}}\ G$. $G \in \{A, !A, \mathbf{1}\}$ is called a goal type. Additional syntax and semantics for this construct are shown in figure 12. Other than the $\underline{\texttt{link}}$ construct, full-CLL inherits all of fCLL's syntax, typing rules and semantics.

$\underline{\texttt{link}}\ E \div S\ \underline{\texttt{to}}\ G$ is evaluated in a context of index variables Σ as follows. First, the lCLL configuration $\mathcal{C} = \Sigma; \cdot \rhd \cdot \Vert E \div S$ is created and allowed to rewrite according to the relation $\rightrightarrows$ till it reaches a *quiescent* configuration $\mathcal{C}'$. By quiescent we mean that no rewrite rule applies to $\mathcal{C}'$ i.e. $\mathcal{C}'$ is in $\rightrightarrows$-normal form. After $\mathcal{C}'$ is obtained, the result of evaluating $\underline{\texttt{link}}\ E \div S\ \underline{\texttt{to}}\ G$ depends on the goal type G.

1. If $G = A$ and $\mathcal{C}' = \Sigma; \hat{\sigma} \rhd \hat{\Gamma} \Vert V : A$ or $\mathcal{C}' = \Sigma; \hat{\sigma} \rhd \hat{\Gamma}, V : A \Vert \cdot$, then $\underline{\texttt{link}}\ E \div S\ \underline{\texttt{to}}\ G$ evaluates to V.

Syntax

$$\text{Expressions } E ::= \ldots \mid \underline{\texttt{link}}\ E \div S\ \underline{\texttt{to}}\ G$$
$$\text{Goal Types } G ::= A \mid !A \mid \mathbf{1}$$

Typing rules

$$\frac{\Sigma; \Gamma; \Delta\ \vdash\ E \div S}{\Sigma; \Gamma; \Delta\ \vdash\ (\underline{\texttt{link}}\ E \div S\ \underline{\texttt{to}}\ G) \div G}\ \textit{LINK}$$

Operational Semantics

$$\frac{\Sigma; \cdot \rhd \cdot \parallel E \div S\ \rightrightarrows^*\ \Sigma; \hat{\sigma} \rhd \hat{\Gamma} \parallel V : A}{\Sigma; \underline{\texttt{link}}\ E \div S\ \underline{\texttt{to}}\ A\ \hookrightarrow\ \Sigma; V} \hookrightarrow \textit{LINK1} \qquad \frac{\Sigma; \cdot \rhd \cdot \parallel E \div S\ \rightrightarrows^*\ \Sigma; \hat{\sigma} \rhd \hat{\Gamma}, V : A \parallel \cdot}{\Sigma; \underline{\texttt{link}}\ E \div S\ \underline{\texttt{to}}\ A\ \hookrightarrow\ \Sigma; V} \hookrightarrow \textit{LINK2}$$

$$\frac{\Sigma; \cdot \rhd \cdot \parallel E \div S\ \rightrightarrows^*\ \Sigma; \hat{\sigma} \rhd \hat{\Gamma}, V : A \parallel \cdot}{\Sigma; \underline{\texttt{link}}\ E \div S\ \underline{\texttt{to}}\ !A\ \hookrightarrow\ \Sigma; !V} \hookrightarrow \textit{LINK3} \qquad \frac{\Sigma; \cdot \rhd \cdot \parallel E \div S\ \rightrightarrows^*\ \Sigma; \hat{\sigma} \rhd \hat{\Gamma} \parallel \cdot}{\Sigma; \underline{\texttt{link}}\ E \div S\ \underline{\texttt{to}}\ \mathbf{1}\ \hookrightarrow\ \Sigma; \star} \hookrightarrow \textit{LINK4}$$

Fig. 12. Full-CLL syntax and semantics

2. If $G = !A$ and $\mathcal{C}' = \Sigma; \hat{\sigma} \rhd \hat{\Gamma}, V : A \parallel \cdot$, then $\underline{\texttt{link}}\ E \div S\ \underline{\texttt{to}}\ G$ evaluates to $!V$.
3. If $G = \mathbf{1}$ and $\mathcal{C}' = \Sigma; \hat{\sigma} \rhd \hat{\Gamma} \parallel \cdot$, then $\underline{\texttt{link}}\ E \div S\ \underline{\texttt{to}}\ G$ evaluates to $\star$.

All these conditions are summarized in figure 12. If none of these conditions hold, evaluation of the $\underline{\texttt{link}}$ construct fails and computation deadlocks. We call this condition *link failure*. In an implementation, this condition can be easily detected and an appropriate exception can be thrown, somewhat similar to a match exception in a language like ML. Since expressions are coerced into terms through a monad, link failure never occurs during evaluation of terms and monadic-terms. As a result, full-CLL has the following progress theorem.

Theorem 4 (Progress for full-CLL).
1. If $\Sigma; \cdot; \cdot\ \vdash\ N : A$ then either $N = V$ or $N \rightsquigarrow N'$ for some N'.
2. If $\Sigma; \cdot; \cdot\ \vdash\ M \div S$ then either $M = M_V$ or $M \mapsto M'$ for some M'.
3. If $\Sigma; \cdot; \cdot\ \vdash\ E \div S$ then either $E = E_V$ or $\Sigma; E \hookrightarrow \Sigma; E'$ for some E' or reduction of $\Sigma; E$ terminates with link failure.

For all practical problems that we encountered, we found it possible to write programs in which link failure never occurs. An investigation of formal methods to ensure this property automatically is a subject of future work. fCLL's preservation theorem (theorem 1) holds for full-CLL also.

Example 3 (Fibonacci numbers in full-CLL). Figure 13 shows a concurrent implementation of Fibonacci numbers in full-CLL. The function `fibc` uses the additional signature from example 2 and assumes that the sort `procid` contains at least three constants k_1, k_2 and k. `fibc` has the type `int` $\rightarrow$ {`!int`}. Given an input integer $n \geq 2$, `fibc` computes the nth Fibonacci number using a $\underline{\texttt{link}}$ construct that starts concurrent computation with a tensor of three processes having identities k_1, k_2 and k respectively. The first two processes recursively compute $fib(n-1)$ and $fib(n-2)$ and send these values as messages to the third process. The third process waits for these messages (m_1 and m_2), extracts their integer contents and adds them together to obtain $fib(n)$. This becomes the result of evaluation of the $\underline{\texttt{link}}$ construct.

During the evaluation of `fibc`, each of the two recursive calls can encounter a $\underline{\texttt{link}}$ construct and create a nested lCLL concurrent computation. Since the two

```
fibc = λn : int.
   if (n = 0 or n = 1) then {!1}
   else
   { link
      (
            {let {!n1} = fibc (n − 1) in (message [k1] [k] n1)}
         ⊗ {let {!n2} = fibc (n − 2) in (message [k2] [k] n2)}
         ⊗ λ̂m1 : mess k1 k. λ̂m2 : mess k2 k.
            {
               let {!x} = fetchmessage [k1] [k] ^ m1 in
               let {!y} = fetchmessage [k2] [k] ^ m2 in
                  !(x + y)
            }
      ) ÷ {mess k1 k} ⊗ {mess k2 k} ⊗ (mess k1 k ⊸ mess k2 k ⊸ {!int})
      to !int
   }
```

Fig. 13. The function `fibc` in full-CLL

recursive calls can be executed simultaneously, there may actually be more than one nested lCLL configuration at the same time. However, these configurations are distinct – processes in one configuration cannot synchronize with those in another. In general, full-CLL programs can spawn several nested concurrent computations that are completely disjoint from each other.

5 Conclusion

We have presented a language that combines functional and concurrent computation in a logically motivated manner. It requires linearity, a restricted form of dependent types, a monad, and focusing, in order to retain the desirable properties of each paradigm in their combination.

Most of the related work has already been discussed in the introduction or in the body of the paper. Perhaps most closely related to our work is CLF [19] and the programming language LolliMon [14] based on its first-order fragment. The similarity lies in the underlying logical foundation of a dependently typed linear logic with monadic encapsulation. The difference lies in the operational semantics and functional extensions. In LolliMon, all computation proceeds via proof search and it should therefore be classified as a concurrent logic programming language. While CLF and LolliMon have λ-abstraction, it is used exclusively for the representation of objects using higher-order abstract syntax, and not for functional computation. Extension by data types, recursion, and so on would therefore be impossible. Moreover, even the operational semantics of the concurrent part of LolliMon is quite different from the one proposed here, where functional terms may be synthesized and evaluated.

The most immediate item of future work is a more realistic implementation, including a more complete functional language and type reconstruction. Since concurrency in CLL is somewhat low-level, it will be important to build up libraries of common idioms. We also plan to investigate means for reasoning about the concurrent computation inside the monad in order to establish freedom from deadlock and other important program properties.

References

1. CLL implementation. Available electronically from http://www.cs.cmu.edu/˜dg.
2. M. Abadi and C. Fournet. Mobile values, new names, and secure communication. In *Proc. of POPL'01*, pages 104–115, 2001.
3. S. Abramsky. Computational interpretations of linear logic. *Theoretical Computer Science*, 111(1–2):3–57, 1993.
4. S. Abramsky, S. Gay, and R. Nagarajan. Specification structures and propositions-as-types for concurrency. In *Logics for Concurrency: Structure vs. Automata—Proc. of the VIIIth Banff Higher Order Workshop*, volume 1043 of *Lecture Notes in Computer Science*. Springer-Verlag, 1996.
5. J.-M. Andreoli. Logic programming with focusing proofs in linear logic. *Journal of Logic and Computation*, 2(3):297–347, 1992.
6. J.-M. Andreoli and R. Pareschi. Communication as fair distribution of knowledge. Technical Report ECRC-91-12, European Computer-Industry Research Centre, 1991.
7. G. Berry and G. Boudol. The chemical abstract machine. *Theoretical Computer Science*, 96:217–248, 1992.
8. G. Boudol. Asynchrony and the pi-calculus. Technical Report RR-1702, INRIA SofiaAntipolis, 1992.
9. S. Conchon and F. L. Fessant. Jocaml: Mobile agents for objective-caml. In *Proc. of ASAMA'99*. IEEE Computer Society, 1999.
10. D. Garg. CLL: A concurrent language built from logical principles. Technical Report CMU-CS-05-104, Computer Science Department, Carnegie Mellon University, January 2005.
11. A. Giacalone, P. Mishra, and S. Prasad. Facile: A symmetric integration of concurrent and functional programming. *International Journal of Parallel Programming*, 18(2):121–160, 1989.
12. J.-Y. Girard. Linear logic. In *Theoretical Computer Science*, volume 5, 1987.
13. S. P. Jones, A. Gordon, and S. Finne. Concurrent Haskell. In *Proc. of POPL'96*, 1996.
14. P. López, F. Pfenning, J. Polakow, and K. Watkins. Monadic concurrent linear logic programming. In *Proc. of PPDP'05*, 2005. To appear.
15. B. C. Pierce and D. N. Turner. Pict: a programming language based on the pi-calculus. In *Proof, language, and interaction: essays in honour of Robin Milner*, pages 455–494. MIT Press, 2000.
16. J. H. Reppy. CML: A higher-order concurrent language. In *Proc. of PLDI'91*, 1991.
17. J. H. Reppy. *Concurrent programming in ML*. Cambridge University Press, 1999.
18. D. Sangiorgi and D. Walker. *The π-calculus: A Theory of Mobile Processes*. Cambridge University Press, 2001. Chapters 15–17.
19. K. Watkins, I. Cervesato, F. Pfenning, and D. Walker. A concurrent logical framework I: Judgements and properties. Technical Report CMU-CS-02-101, Computer Science Department, Carnegie Mellon University, May 2003.
20. H. Xi and F. Pfenning. Dependent types in practical programming. In *Proc. of POPL'99*, 1999.

A Type system for fCLL

$\boxed{\mathbf{\Sigma;\Gamma;\Delta \ \vdash \ N : A}}$

$$\frac{}{\Sigma;\Gamma;x:A \ \vdash \ x:A}\ \textsf{Hyp1} \qquad \frac{}{\Sigma;\Gamma,x:A;\cdot \ \vdash \ x:A}\ \textsf{Hyp2} \qquad \frac{\Sigma;\Gamma;\Delta \ \vdash \ E \div S}{\Sigma;\Gamma;\Delta \ \vdash \ \{E\}:\{S\}}\ \{\}\textsf{I}$$

$$\frac{\Sigma;\Gamma,x:A;\Delta \ \vdash \ N:B}{\Sigma;\Gamma;\Delta \ \vdash \ \lambda x:A.N : A \to B}\ \to\textsf{I} \qquad \frac{\Sigma;\Gamma;\Delta,x:A \ \vdash \ N:B}{\Sigma;\Gamma;\Delta \ \vdash \ \hat{\lambda} x:A.N : A \multimap B}\ \multimap\textsf{I}$$

$$\frac{\Sigma,i:\gamma;\Gamma;\Delta \ \vdash \ N:A}{\Sigma;\Gamma;\Delta \ \vdash \ \Lambda i:\gamma.N : \forall i:\gamma.A}\ \forall\textsf{I} \qquad \frac{\Sigma;\Gamma;\Delta \ \vdash \ N_1:A_1 \qquad \Sigma;\Gamma;\Delta \ \vdash \ N_2:A_2}{\Sigma;\Gamma;\Delta \ \vdash \ \langle N_1,N_2\rangle : A_1\& A_2}\ \&\textsf{I}$$

$$\frac{\Sigma;\Gamma;\Delta \ \vdash \ N_1 : A \to B \qquad \Sigma;\Gamma;\cdot \ \vdash \ N_2 : A}{\Sigma;\Gamma;\Delta \ \vdash \ N_1\ N_2 : B}\ \to\textsf{E}$$

$$\frac{\Sigma;\Gamma;\Delta_1 \ \vdash \ N_1 : A \multimap B \qquad \Sigma;\Gamma;\Delta_2 \ \vdash \ N_2 : A}{\Sigma;\Gamma;\Delta_1,\Delta_2 \ \vdash \ N_1\ \hat{}\ N_2 : B}\ \multimap\textsf{E}$$

$$\frac{\Sigma;\Gamma;\Delta \ \vdash \ N : \forall i:\gamma.A(i) \qquad \Sigma \ \vdash \ t\in\gamma}{\Sigma;\Gamma;\Delta \ \vdash \ N\ [t] : A(t)}\ \forall\textsf{E}$$

$$\frac{\Sigma;\Gamma;\Delta \ \vdash \ N : A_1\& A_2}{\Sigma;\Gamma;\Delta \ \vdash \ \pi_1 N : A_1}\ \&E_1 \qquad \frac{\Sigma;\Gamma;\Delta \ \vdash \ N : A_1\& A_2}{\Sigma;\Gamma;\Delta \ \vdash \ \pi_2 N : A_2}\ \&E_2$$

$\boxed{\mathbf{\Sigma;\Gamma;\Delta \ \vdash \ M \div\!\!\!\!\!\div S}}$

$$\frac{\Sigma;\Gamma;\Delta \ \vdash \ N:A}{\Sigma;\Gamma;\Delta \ \vdash \ N \div\!\!\!\!\!\div A}\ :\div\!\!\!\!\!\div \qquad \frac{\Sigma;\Gamma;\cdot \ \vdash \ N:A}{\Sigma;\Gamma;\cdot \ \vdash \ !N \div\!\!\!\!\!\div !A}\ \textsf{!R}$$

$$\frac{}{\Sigma;\Gamma;\cdot \ \vdash \ \star \div\!\!\!\!\!\div \mathbf{1}}\ \textsf{1R} \qquad \frac{\Sigma;\Gamma;\Delta \ \vdash \ M \div\!\!\!\!\!\div S(t) \qquad \Sigma \ \vdash \ t\in\gamma}{\Sigma;\Gamma;\Delta \ \vdash \ [t,M] \div\!\!\!\!\!\div \exists i:\gamma.S(i)}\ \exists\textsf{R}$$

$$\frac{\Sigma;\Gamma;\Delta_1 \ \vdash \ M_1 \div\!\!\!\!\!\div S_1 \qquad \Sigma;\Gamma;\Delta_2 \ \vdash \ M_2 \div\!\!\!\!\!\div S_2}{\Sigma;\Gamma;\Delta_1,\Delta_2 \ \vdash \ M_1\otimes M_2 \div\!\!\!\!\!\div S_1\otimes S_2}\ \otimes\textsf{R}$$

$\boxed{\mathbf{\Sigma;\Gamma;\Delta \ \vdash \ E \div S}}$

$$\frac{\Sigma;\Gamma;\Delta \ \vdash \ M \div\!\!\!\!\!\div S}{\Sigma;\Gamma;\Delta \ \vdash \ M \div S}\ \div\!\!\!\!\!\div\,\div \qquad \frac{\Sigma;\Gamma;\Delta_1 \ \vdash \ N:\{S\} \qquad \Sigma;\Gamma;\Delta_2;p:S \ \vdash \ E\div S'}{\Sigma;\Gamma;\Delta_1,\Delta_2 \ \vdash \ \underline{\texttt{let}}\ \{p:S\}=N\ \underline{\texttt{in}}\ E\div S'}\ \{\}\textsf{E}$$

$\boxed{\mathbf{\Sigma;\Gamma;\Delta;\Psi \ \vdash \ E \div S}}$

$$\frac{\Sigma;\Gamma;\Delta \ \vdash \ E\div S}{\Sigma;\Gamma;\Delta;\cdot \ \vdash \ E\div S}\ \div\div \qquad \frac{\Sigma;\Gamma;\Delta;\Psi \ \vdash \ E\div S}{\Sigma;\Gamma;\Delta;\star:\mathbf{1},\Psi \ \vdash \ E\div S}\ \textsf{1L}$$

$$\frac{\Sigma;\Gamma;\Delta,x:A;\Psi \ \vdash \ E\div S}{\Sigma;\Gamma;\Delta;x:A,\Psi \ \vdash \ E\div S}\ \textsf{varL} \qquad \frac{\Sigma;\Gamma,x:A;\Delta;\Psi \ \vdash \ E\div S}{\Sigma;\Gamma;\Delta;!x:!A,\Psi \ \vdash \ E\div S}\ \textsf{!L}$$

$$\frac{\Sigma;\Gamma;\Delta;p_1:S_1,p_2:S_2,\Psi \ \vdash \ E\div S}{\Sigma;\Gamma;\Delta;p_1\otimes p_2:S_1\otimes S_2,\Psi \ \vdash \ E\div S}\ \otimes\textsf{L} \qquad \frac{\Sigma,i:\gamma;\Gamma;\Delta;p:S',\Psi \ \vdash \ E\div S}{\Sigma;\Gamma;\Delta;[i,p]:\exists i:\gamma.S',\Psi \ \vdash \ E\div S}\ \exists\textsf{L}\ (i\ \textsf{fresh})$$

Fig. 14. fCLL type system

B Operational Semantics for *f*CLL

$\boxed{\mathbf{N} \rightsquigarrow \mathbf{N}'}$

$$\frac{N \rightsquigarrow N'}{\pi_1\ N \rightsquigarrow \pi_1 N'} \rightsquigarrow \pi_1 \qquad \frac{N \rightsquigarrow N'}{\pi_2\ N \rightsquigarrow \pi_2 N'} \rightsquigarrow \pi_2$$

$$\frac{}{\pi_1\ \langle N_1, N_2\rangle \rightsquigarrow N_1} \rightsquigarrow \langle\rangle\pi_1 \qquad \frac{}{\pi_2\ \langle N_1, N_2\rangle \rightsquigarrow N_2} \rightsquigarrow \langle\rangle\pi_2$$

$$\frac{N_1 \rightsquigarrow N_1'}{N_1\ N_2 \rightsquigarrow N_1'\ N_2} \rightsquigarrow APP_1 \qquad \frac{N_2 \rightsquigarrow N_2'}{V\ N_2 \rightsquigarrow V\ N_2'} \rightsquigarrow APP_2 \qquad \frac{}{(\lambda x : A.N)\ V \rightsquigarrow N[V/x]} \rightsquigarrow \lambda APP$$

$$\frac{N_1 \rightsquigarrow N_1'}{N_1\ \hat{}\ N_2 \rightsquigarrow N_1'\ \hat{}\ N_2} \rightsquigarrow LAPP_1 \qquad \frac{N_2 \rightsquigarrow N_2'}{V\ \hat{}\ N_2 \rightsquigarrow V\ \hat{}\ N_2'} \rightsquigarrow LAPP_2 \qquad \frac{}{(\hat{\lambda} x : A.N)\ \hat{}\ V \rightsquigarrow N[V/x]} \rightsquigarrow \hat{\lambda} LAPP$$

$$\frac{N \rightsquigarrow N'}{N\ [t] \rightsquigarrow N'\ [t]} \rightsquigarrow \forall \qquad \frac{}{(\Lambda i : \gamma.N)\ [t] \rightsquigarrow N[t/i]} \rightsquigarrow \Lambda APP$$

$\boxed{\mathbf{M} \mapsto \mathbf{M}'}$

$$\frac{N \rightsquigarrow N'}{N \mapsto N'} \rightsquigarrow\mapsto \qquad \frac{N \rightsquigarrow N'}{!N \mapsto\ !N'} \mapsto !$$

$$\frac{M_1 \mapsto M_1'}{M_1 \otimes M_2 \mapsto M_1' \otimes M_2} \mapsto \otimes_1 \qquad \frac{M_2 \mapsto M_2'}{M_1 \otimes M_2 \mapsto M_1 \otimes M_2'} \mapsto \otimes_2$$

$$\frac{M \mapsto M'}{[t, M] \mapsto [t, M']} \mapsto \exists$$

$\boxed{\mathbf{\Sigma}; \mathbf{E} \hookrightarrow \mathbf{\Sigma}; \mathbf{E}'}$

$$\frac{M \mapsto M'}{\Sigma; M \hookrightarrow \Sigma; M'} \mapsto\hookrightarrow \qquad \frac{}{\Sigma; \underline{\texttt{let}}\ \{p : S\} = \{M_V\}\ \underline{\texttt{in}}\ E \hookrightarrow \Sigma; E[M_V/p]} \hookrightarrow LETRED$$

$$\frac{N \rightsquigarrow N'}{\Sigma; \underline{\texttt{let}}\ \{p : S\} = N\ \underline{\texttt{in}}\ E \hookrightarrow \Sigma; \underline{\texttt{let}}\ \{p : S\} = N'\ \underline{\texttt{in}}\ E} \hookrightarrow LET_1$$

$$\frac{\Sigma; E \hookrightarrow \Sigma; E'}{\Sigma; \underline{\texttt{let}}\ \{p : S\} = \{E\}\ \underline{\texttt{in}}\ E_1 \hookrightarrow \Sigma; \underline{\texttt{let}}\ \{p : S\} = \{E'\}\ \underline{\texttt{in}}\ E_1} \hookrightarrow LET_2$$

Fig. 15. *f*CLL operational semantics