

LANGUAGE PRIMITIVES AND TYPE DISCIPLINE FOR STRUCTURED COMMUNICATION-BASED PROGRAMMING

KOHEI HONDA*, VASCO T. VASCONCELOS[†], AND MAKOTO KUBO[‡]

ABSTRACT. We introduce basic language constructs and a type discipline as a foundation of structured communication-based concurrent programming. The constructs, which are easily translatable into the summation-less asynchronous π -calculus, allow programmers to organise programs as a combination of multiple flows of (possibly unbounded) reciprocal interactions in a simple and elegant way, subsuming the preceding communication primitives such as method invocation and rendez-vous. The resulting syntactic structure is exploited by a type discipline à la ML, which offers a high-level type abstraction of interactive behaviours of programs as well as guaranteeing the compatibility of interaction patterns between processes in a well-typed program. After presenting the formal semantics, the use of language constructs is illustrated through examples, and the basic syntactic results of the type discipline are established. Implementation concerns are also addressed.

1. INTRODUCTION

Recently, significance of programming practice based on communication among processes is rapidly increasing by the development of networked computing. From network protocols over the Internet to server-client systems in local area networks to distributed applications in the world wide web to interaction between mobile robots to a global banking system, the execution of complex, reciprocal communication among multiple processes becomes an important element in the achievement of the goals of applications. Many programming languages and formalisms have been proposed so far for the description of software based on communication. As programming languages, we have CSP [9], Ada [33], languages based on Actors [2], POOL [4], ABCL [39], Concurrent Smalltalk [38], or more recently Pict and other π -calculus-based languages [7, 27, 34, 13]. As formalisms, we have CCS [19], Theoretical CSP [10], π -calculus [22], and other process algebras. In another vein, we have functional programming languages augmented with communication primitives, such as CML [30], dML [25], and Concurrent Haskell [15]. In these languages and formalisms, various communication primitives have been offered (such as synchronous/asynchronous message passing, remote procedure call, method-call and rendez-vous), and the description of communication behaviour is done by combining these primitives. What we observe in these primitives is that, while they do express one-time interaction between processes, there is no construct to structure a series of reciprocal interactions between two parties as such. That is, the only way to represent a series of communications following a certain scenario (think of interactions between a file server and its client) is to describe them as a collection of distinct, unrelated interactions. In applications based on complex interactions among concurrent processes, which are appearing more and more in these days, the lack of structuring methods would result in low readability and careless bugs in final programs, apart from the case when the whole communication behaviour can be simply described as, say, a one-time remote

*Dept. of Computer Science, University of Edinburgh, UK. [†]Dept. of Computer Science, University of Lisbon, Portugal. [‡]Dept. of Computer Science, Chiba University of Commerce, Japan.

procedure call. The situation may be illustrated in comparison with the design history of imperative programming languages. In early imperative languages, programs were constructed as a bare sequence of primitives which correspond to machine operations, such as assignment and goto (consider early Fortran). As is well-known, as more programs in large size were constructed, it had become clear that such a method leads to programs without lucidity, readability or verifiability, so that the notion of structured programming was proposed in 1970's. In present days, having the language constructs for structured programming is a norm in imperative languages.

Such a comparison raises the question as to whether we can have a similar structuring method in the context of communication-based concurrent programming. Its objective is to offer a basic means to describe complex interaction behaviours with clarity and discipline at the high-level of abstraction, together with a formal basis for verification. Its key elements would, above all, consist of (1) the basic communication primitives (corresponding to assignment and arithmetic operations in the imperative setting), and (2) the structuring constructs to combine them (corresponding to “if-then-else” and “while”). Verification methodologies on their basis should then be developed.

The present paper proposes structuring primitives and the accompanying type discipline, as a basic structuring method for communication-based concurrent programming. The proposed constructs have a simple operational semantics, and various communication patterns, including those of the conventional primitives as well as those which go beyond them, are representable as their combination with clarity and rigor at the high-level of abstraction. The type discipline plays a fundamental role, guaranteeing compatibility of interaction patterns among processes via a type inference algorithm in the line of ML [23]. Concretely our proposal consists of the following key ideas.

- A basic structuring concept for communication-based programming, called *session*. A session is a chain of dyadic interactions whose collection constitutes a program. A session is designated by a private port called *channel*, through which interactions belonging to that session are performed. Channels form a distinct syntactic domain; its differentiation from the usual port names is a basic design decision we take to make the logical structure of programs explicit. Other than session, the standard structuring constructs for concurrent programming, *parallel composition*, *name hiding*, *conditional* and *recursion* are provided. In particular the combination of recursion and session allows the expression of unbounded thread of interactions as a single abstraction unit.
- Three basic communication primitives, *value passing*, *label branching*, and *delegation*. The first is the standard synchronous message passing as found in e.g. CSP or π -calculus. The second is a purified form of method invocation, deprived of value passing. The third is for passing a channel to another process, thus allowing a programmer to dynamically distribute a single session among multiple processes in a well-structured way. Together with the session structure, the combination of these primitives allows the flexible description of complex communication structures with clarity and discipline.
- A *basic type discipline* for the communication primitives, as an indispensable element of the structuring method. The typability of a program ensures two possibly communicating processes always own compatible communication patterns. For example, a procedural call has a pattern of output-input from the caller's viewpoint; then the callee should have a communication pattern of input-output. Because incompatibility of interaction patterns between processes would be one of the main reasons for bugs in communication-based programming, we believe such a type

discipline has important pragmatic significance. The derived type gives high-level abstraction of interactive behaviours of a program.

Because communication between processes over the network can be done between modules written in different programming languages, the proposed communication constructs may as well be used by embedding them in various programming languages; however for simplicity we present them as a self-contained small programming language, which has been stripped to the barest minimum necessary for explanation of its novel features. Using the language, the basic concept of the proposed constructs is illustrated through programming examples. They show how extant communication primitives such as remote procedural-call and method invocation can be concisely expressed as sessions of specific patterns (hence with specific type abstractions). They also show how sessions can represent those communication structures which do not conform to the preceding primitives but which would naturally arise in practice. This suggests that the session-based program organisation may constitute a synthesis of a number of familiar as well as new programming ideas concerning communication. We also show the proposed constructs can be simply translatable into the asynchronous polyadic π -calculus with branching [36]. This suggests the feasibility of implementation in distributed environment. Yet much remains to be studied concerning the proposed structuring method, including the efficient implementation of the constructs and the accompanying reasoning principles. See Section 6 for more discussions.

The technical content of the present paper is developed on the basis of the preceding proposal [32] due to Kaku Takeuchi and the present authors. The main contributions of the present proposal in comparison with [32] are: the generalisation of session structure by delegation and recursive sessions, which definitely enlarges the applicability of the proposed structuring method; the typing system incorporating these novel concepts; representation of conventional communication primitives by the structuring constructs; and basic programming examples which show how the constructs, in particular through the use of the above mentioned novel features, can lucidly represent complex communication structures which would not be amenable to conventional communication primitives.

In the rest of the paper, Section 2 introduces the language primitives and their operational semantics. Section 3 illustrates how the primitives allow clean representation of extant communication primitives. Section 4 shows how the primitives can represent those interaction structures beyond those of the conventional communication primitives through key programming examples. Section 5 presents the typing system and establishes the basic syntactic results. Section 6 concludes with discussions on the implementation concerns, related works and further issues.

2. SYNTAX AND OPERATIONAL SEMANTICS

2.1. Basic Concepts. The central idea in the present structuring method is a *session*. A session is a series of reciprocal interactions between two parties, possibly with branching and recursion, and serves as a unit of abstraction for describing interaction. Communications belonging to a session are done via a port specific to that session, called a *channel*. A fresh channel is generated when initiating each session, for the use in communications in the session. We use the following syntax for the initiation of a session.

$$\mathbf{request} \ a(k) \ \mathbf{in} \ P \quad \mathbf{accept} \ a(k) \ \mathbf{in} \ P \quad \text{initiation of session}$$

The **request** first requests, via a name a , the initiation of a session as well as the generation of a fresh channel, then P would use the channel for later communications. The **accept**, on the other hand, receives the request for the initiation of a session via a ,

generates a new channel k , which would be used for communications in P . In the above grammar, the parenthesis (k) and the key word **in** shows the binding and its scope. Thus, in **request** $a(k)$ **in** P , the part (k) binds the free occurrences of k in P . This convention is used uniformly throughout the present paper.

Via a channel of a session, three kinds of atomic interactions are performed: *value sending* (including name passing), *branching* and *channel passing* (or *delegation*).

$k![e_1 \cdots e_n]; P$	$k?(x_1 \cdots x_n)$ in P	data sending/receiving
$k < l; P$	$k \triangleright \{l_1 : P_1 \parallel \cdots \parallel l_n : P_n\}$	label selection/branching
throw $k[k']; P$	catch $k(k')$ in P	channel sending/receiving (delegation)

Data sending/receiving is the standard synchronous message passing. Here e_i denotes an expression such as arithmetic/boolean formulae as well as names. We assume variables $x_1..x_n$ are all distinct. We do not consider program passing for simplicity, cf. [12]. The *branching/selection* is the minimisation of method invocation in object-based programming. l_1, \dots, l_n are *labels* which are assumed to be pairwise distinct. The *channel sending/receiving*, which we often call *delegation*, passes a channel which is being used in a session to another process, thus radically changing the structure of a session. Delegation is the generalisation of the concept with the same name originally conceived in the concurrent object-oriented programming [40]. See Section 4.3 for detailed discussions. In passing we note that its treatment distinct from the usual value passing is essential for both disciplined programming and a tractable type inference.

Communication primitives, organised by sessions, are further combined by the following standard constructs in concurrent programming.

$P_1 \mid P_2$	concurrent composition
$(\nu a)P$ $(\nu k)P$	name/channel hiding
if e then P else Q	conditional
def $X_1(\tilde{x}_1 \tilde{k}_1) = P_1$ and \cdots and $X_n(\tilde{x}_n \tilde{k}_n) = P_n$ in P	recursion

We do not need sequencing since each communication primitive already accompanies one. We also use **inact**, the *inaction*, which denotes the lack of action (acting as the unit of “|”). *Hiding* declares a name/channel to be local in its scope (here P). Channel hiding may not be used for usual programming, but is needed for the operational semantics presented later. In conditional, e should be a boolean expression. In recursion, X , a process variable, would occur in $P_1 \dots P_n$ and P zero or more times. Identifiers in $\tilde{x}_i \tilde{k}_i$ should be pairwise distinct. We can use replication (or a single recursion) to achieve the same effect, but multiple recursion is preferable for well-structured programs.

This finishes the introduction of all language constructs we shall use in this paper. We give a simple example of a program.

accept $a(k)$ **in** $k![1]; k?(y)$ **in** P | **request** $a(k)$ **in** $k?(x)$ **in** $k![x + 1];$ **inact**.

The first process receives a request for a new session via a , generates k , sends 1 and receives a return value via k , while the second requests the initiation of a session via a , receives the value via the generated channel, then returns the result of adding 1 to the value. Observe the compatibility of communication patterns between two processes.

2.2. Syntax Summary. We summarise the syntax we have introduced so far. Base sets are: *names*, ranged over by a, b, \dots ; *channels*, ranged over by k, k' ; *variables*, ranged over by x, y, \dots ; *constants* (including names, integers and booleans), ranged over by c, c', \dots ; *expressions* (including constants), ranged over by e, e', \dots ; *labels*, ranged over by l, l', \dots ; and *process variables*, ranged over by X, Y, \dots . u, u', \dots denote names and

channels. Then *processes*, ranged over by P, Q, \dots , are given by the following grammar.

$P ::=$	$\text{request } a(k) \text{ in } P$ $\text{accept } a(k) \text{ in } P$ $k![\tilde{e}]; P$ $k?(\tilde{x}) \text{ in } P$ $k \triangleleft l; P$ $k \triangleright \{l_1 : P_1 \parallel \dots \parallel l_n : P_n\}$ $\text{throw } k[k']; P$ $\text{catch } k(k') \text{ in } P$ $\text{if } e \text{ then } P \text{ else } Q$ $P \mid Q$ inact $(\nu u)P$ $\text{def } D \text{ in } P$ $X[\tilde{e}\tilde{k}]$	session request session acceptance data sending data reception label selection label branching channel sending channel reception conditional branch parallel composition inaction name/channel hiding recursion process variables
$D ::=$	$X_1(\tilde{x}_1\tilde{k}_1) = P_1 \text{ and } \dots \text{ and } X_n(\tilde{x}_n\tilde{k}_n) = P_n$	declaration for recursion

The association of “|” is the weakest, others being the same. Parenthesis (\cdot) denotes binders which bind the corresponding free occurrences. The standard *simultaneous substitution* is used, writing e.g. $P[\tilde{c}/\tilde{x}]$. The sets of free names/channels/variables/process variables of say P , defined in the standard way, are respectively denoted by $\text{fn}(P)$, $\text{fc}(P)$, $\text{fv}(P)$ and $\text{fpv}(P)$. The alpha-equality is written \equiv_α . We also set $\text{fu}(P) \stackrel{\text{def}}{=} \text{fc}(P) \cup \text{fn}(P)$. Processes without free variables or free channels are called *programs*.

2.3. Operational Semantics. For the concise definition of the operational semantics of the language constructs, we introduce the *structural equality* \equiv (cf. [5, 20]), which is the smallest congruence relation on processes including the following equations.

1. $P \equiv Q$ if $P \equiv_\alpha Q$.
2. $P \mid \text{inact} \equiv P$, $P \mid Q \equiv Q \mid P$, $(P \mid Q) \mid R \equiv P \mid (Q \mid R)$.
3. $(\nu u)\text{inact} \equiv \text{inact}$, $(\nu uu)P \equiv (\nu u)P$, $(\nu uu')P \equiv (\nu u'u)P$, $(\nu u)P \mid Q \equiv (\nu u)(P \mid Q)$ if $u \notin \text{fu}(Q)$, $(\nu u)\text{def } D \text{ in } P \equiv \text{def } D \text{ in } (\nu u)P$ if $u \notin \text{fu}(D)$.
4. $(\text{def } D \text{ in } P) \mid Q \equiv \text{def } D \text{ in } (P \mid Q)$ if $\text{fpv}(D) \cap \text{fpv}(Q) = \emptyset$.
5. $\text{def } D \text{ in } (\text{def } D' \text{ in } P) \equiv \text{def } D \text{ and } D' \text{ in } P$ if $\text{fpv}(D) \cap \text{fpv}(D') = \emptyset$.

Now the operational semantics is given by the *reduction relation* \rightarrow , denoted $P \rightarrow Q$, which is the smallest relation on processes generated by the following rules.

[LINK]	$(\text{accept } a(k) \text{ in } P_1) \mid (\text{request } a(k) \text{ in } P_2) \rightarrow (\nu k)(P_1 \mid P_2)$
[COM]	$(k![\tilde{e}]; P_1) \mid (k?(\tilde{x}) \text{ in } P_2) \rightarrow P_1 \mid P_2[\tilde{c}/\tilde{x}] \quad (\tilde{e} \downarrow \tilde{c})$
[LABEL]	$(k \triangleleft l_i; P) \mid (k \triangleright \{l_1 : P_1 \parallel \dots \parallel l_n : P_n\}) \rightarrow P \mid P_i \quad (1 \leq i \leq n)$
[PASS]	$(\text{throw } k[k']; P_1) \mid (\text{catch } k(k') \text{ in } P_2) \rightarrow P_1 \mid P_2$
[IF1]	$\text{if } e \text{ then } P_1 \text{ else } P_2 \rightarrow P_1 \quad (e \downarrow \text{true})$
[IF2]	$\text{if } e \text{ then } P_1 \text{ else } P_2 \rightarrow P_2 \quad (e \downarrow \text{false})$
[DEF]	$\text{def } D \text{ in } (X[\tilde{e}\tilde{k}] \mid Q) \rightarrow \text{def } D \text{ in } (P[\tilde{c}/\tilde{x}] \mid Q) \quad (\tilde{e} \downarrow \tilde{c}, X(\tilde{x}\tilde{k}) = P \in D)$
[SCOP]	$P \rightarrow P' \Rightarrow (\nu u)P \rightarrow (\nu u)P'$
[PAR]	$P \rightarrow P' \Rightarrow P \mid Q \rightarrow P' \mid Q$
[STR]	$P \equiv P' \text{ and } P' \rightarrow Q' \text{ and } Q' \equiv Q \Rightarrow P \rightarrow Q$

Above we assume the standard *evaluation relation* \downarrow on expressions is given. We write \rightarrow^* for the reflexive, transitive closure of \rightarrow . Note how a fresh channel is generated in [LINK] rule as the result of interaction (in this way **request** $a(k)$ **in** P corresponds to the *bound output* in π -calculus[22]). In the [LABEL] rule, one of the branches is selected, discarding the remaining ones. Note we do not allow reduction under various communication prefixes, as in standard process calculi. As an example, the simple program in 2.1 has the following reduction (below and henceforth we omit trailing inactions).

$$\begin{aligned} & \text{accept } a(k) \text{ in } k![1]; k?(y) \text{ in } P \mid \text{request } a(k) \text{ in } k?(x) \text{ in } k![x+1] \\ & \rightarrow (\nu k)(k![1]; k?(y) \text{ in } P \mid k?(x) \text{ in } k![x+1]) \\ & \rightarrow (\nu k)(k?(y) \text{ in } P \mid k![x+1]) \rightarrow P[2/y]. \end{aligned}$$

Observe how interaction proceeds in a lock-step fashion. This is due to the synchronous form of the present communication primitives.

3. REPRESENTING COMMUNICATION (1) EXTANT COMMUNICATION PRIMITIVES

This section discusses how structuring primitives can represent with ease the communication patterns of conventional communication primitives which have been in use in programming languages. They show how we can understand the extant communication patterns as a fixed way of combining the proposed primitives.

3.1. Call-return. The *call-return* is a widely used communication pattern in which a process calls another process, then the callee would, after some processing, returns some value to the caller. Usually the caller just waits until the reply message arrives. This concept is widely used as a basic primitive in distributed computing under the name of *remote procedure call*. We may use the following pair of notations for call-return.

$$x = \text{call } f[e_1 \cdots e_n] \text{ in } P, \quad \text{fun } f(x_1 \cdots x_n) \text{ in } P.$$

On the right, the return command **return** $[e]$ would occur in P . Assume these constructs are added to the syntax in 2.2. A simple programming example follows.

Example 3.1 (Factorial).

$$\begin{aligned} \text{Fact}(f) = & \text{fun } f(x) \text{ in} \\ & \text{if } x = 1 \text{ then } \text{return}[1] \\ & \text{else } (\nu b)(\text{Fact}[b] \mid y = \text{call } b[x-1] \text{ in } \text{return}[x * y]) \end{aligned}$$

Here and henceforth we write $X(\tilde{x}\tilde{k}) = P$, or a sequence of such equations, for the declaration part of recursion, leaving the body part implicit. This example implements the standard recursive algorithm for the factorial function. $y = \text{call } f[5] \text{ in } P$ would give its client process.

The communication patterns based on call-return are easily representable by the combination of request/accept and send/receive. We first show the mapping of the caller. Below $\llbracket \cdot \rrbracket$ denotes the inductive mapping into the structuring primitives.

$$\llbracket x = \text{call } a[\tilde{e}] \text{ in } P \rrbracket \stackrel{\text{def}}{=} \text{request } a(k) \text{ in } k![\tilde{e}]; k?(x) \text{ in } \llbracket P \rrbracket.$$

Naturally we assume k is chosen fresh. The basic scenario is that a caller first initiates the session, sends the values, and waits until it receives the answer on the same channel. On the other hand, the callee is translated as follows:

$$\llbracket \text{fun } a(\tilde{x}) \text{ in } P \rrbracket \stackrel{\text{def}}{=} \text{accept } a(k) \text{ in } k?(\tilde{x}) \text{ in } \llbracket P \rrbracket[k![\tilde{e}]/\text{return}[e]].$$

Here $[k![e]/\mathbf{return}[e]]$ denotes the syntactic substitution of $k![e]$ for each $\mathbf{return}[e]$. Observe that the original “return” operation is expressed by the “send” primitive of the structuring primitives, cf.[20]. As one example, let us see how the factorial program in Example 3.1 is translated by the above mapping.

Example 3.2 (Factorial, translation).

$$\begin{aligned} \mathbf{Fact}(f) = & \mathbf{accept} \ f(k) \ \mathbf{in} \ k?(x) \ \mathbf{in} \\ & \mathbf{if} \ x = 1 \ \mathbf{then} \ k![1]; \\ & \mathbf{else} \ (\nu b)(\mathbf{Fact}[b] \mid \mathbf{request} \ b(k') \ \mathbf{in} \ k'![x-1]; k'(y) \ \mathbf{in} \ k![x * y]) \end{aligned}$$

Notice how the usage of k' and k differentiates the two contexts of communication. If we compose the translation of factorial with that of its user, we have the reduction:

$$\mathbf{Fact}[f] \mid \llbracket y = \mathbf{call} \ f[3] \ \mathbf{in} \ P \rrbracket \rightarrow^* \mathbf{Fact}[f] \mid \llbracket P \rrbracket[6/y].$$

In this way, the semantics of the synchronous call-return is given by that of the structuring primitives via the translation. Some observations follow.

- (1) The significance of the specific notations for call-return would lie in the declaration of the assumed fixed communication pattern, which would enhance readability and help verification. At the same time, the translation retains the same level of clarity as the original code, even if it does need a few additional key strokes.
- (2) In translation, the caller and the callee in general own complementary communication patterns, i.e. input-output meets output-input. However if, for example, the return commands appear twice in the callee, the complementarity is lost. This relates to the notion of types we discuss later: indeed, non-complementary patterns are rejected by the typing system.
- (3) The translation also suggests how structuring primitives would generalise the traditional call-return structure. That is, in addition to the fixed pattern of input-output (or, correspondingly, output-input), we can have a sequence of interactions of indefinite length. For such programming examples, see Section 4.

3.2. Method Invocation. The idea of method invocation originates in object-oriented languages, where a caller calls an object by specifying a method name, while the object waits with a few methods together with the associated codes, so that, when invoked, executes the code corresponding to the method name. The call may or may not result in returning an answer. As a notation, an “object” would be written:

$$\mathbf{obj} \ a : \{l_1(\tilde{x}_1) \ \mathbf{in} \ P_1 \mid \cdots \mid l_n(\tilde{x}_n) \ \mathbf{in} \ P_n\},$$

where a gives an object identifier, l_1, \dots, l_n are labels (all pairwise distinct) with formal parameters \tilde{x}_i , and P_i gives the code corresponding to each l_i . The return action, written $\mathbf{return}[e]$, may occur in each P_i as necessary. A caller then becomes:

$$x = a.l_i[\tilde{e}] \ \mathbf{in} \ P \quad a.l_i[\tilde{e}]; P.$$

The left-hand side denotes a process which invokes the object with a method l_i together with arguments \tilde{e} , then, receiving the answer, assigns the result to x , and finally executes the continuation P . The right-hand side is a notation for the case when the invocation does not need the return value. As an example, let us program a simple cell.

Example 3.3 (Cell).

$$\mathbf{Cell}(a, c) = \mathbf{obj} \ c : \{\mathbf{read}() \ \mathbf{in} \ (\mathbf{return}[c] \mid \mathbf{Cell}[a, c]) \mid \mathbf{write}(x) \ \mathbf{in} \ \mathbf{Cell}[a, x]\}.$$

The cell $\text{Cell}[a, c]$ denotes an object which saves its value in c and has a name a . There are two methods, **read** and **write**, each with the obvious functionalities. A caller which “reads” this object would be written, for example, $x = a.\text{read}[] \text{ in } P$.

The method invocation can be represented in the structuring constructs by combining label branching and value passing. We show the translations of an object, a caller which expects the return, and a caller which does not expect the return, in this order.

$$\begin{aligned} \llbracket \text{obj } a : \{l_1(\tilde{x}_1) \text{ in } P_1 \} \cdots \{l_n(\tilde{x}_n) \text{ in } P_n \} \rrbracket &\stackrel{\text{def}}{=} \\ \text{accept } a(k) \text{ in } k \triangleright \{l_1 : k?(x_1) \text{ in } \llbracket P_1 \rrbracket \sigma \} \cdots \{l_n : k?(x_n) \text{ in } \llbracket P_n \rrbracket \sigma \} \\ \llbracket x = a.l_i[\tilde{e}] \text{ in } P \rrbracket &\stackrel{\text{def}}{=} \text{request } a(k) \text{ in } k \triangleleft l_i; k![\tilde{e}]; k?(x) \text{ in } \llbracket P \rrbracket \\ \llbracket a.l_i[\tilde{e}]; P \rrbracket &\stackrel{\text{def}}{=} \text{request } a(k) \text{ in } k \triangleleft l_i; k![\tilde{e}]; \llbracket P \rrbracket \end{aligned}$$

In each translation, k should be fresh. In the first equation, σ denotes the substitution $[k![e]/\text{return}[e]]$, replacing all occurrences of $\text{return}[e]$ by $k![e]$. Observe that a method invocation is decomposed into label branching and value passing (cf. [37]). Using this mapping, the cell is now translated into:

Example 3.4 (Cell, translation).

$$\text{Cell}(a, c) = \text{accept } a(k) \text{ in } k \triangleright \{\text{read} : k![e]; \text{Cell}[a, c] \parallel \text{write} : k?(x) \text{ in } \text{Cell}[a, x]\}.$$

Similarly, $x = a.\text{read}[] \text{ in } P$ is translated as $\text{request } c(k) \text{ in } k \triangleleft \text{read}; k?(x) \text{ in } P$, while $a.\text{write}[3]$ becomes $\text{request } a(k) \text{ in } k \triangleleft \text{write}; k![3]$. Some observations follow.

- (1) The translation is not much more complex than the original text. The specific notation however has a role of declaring the fixed communication pattern and saves key-strokes.
- (2) Here again, the translation of the caller and the callee in general results in complementary patterns. There are however two main cases where the complementarity is lost, one due to the existence/lack of the return statement (e.g. $x = a.\text{write}[3] \text{ in } Q$ above) and another due to an invocation of an object with a non-existent method. Again such inconsistency is detectable by a typing system introduced later.
- (3) The translation suggests how structuring primitives can generalise the standard method invocation. For example, an object may in turn invoke the method of the caller after being invoked. We believe that, in programming practice based on the idea of interacting objects, such reciprocal, continuous interaction would arise naturally and usefully. Such examples are discussed in the next section.

In addition to call-return and method invocation, we can similarly represent other communication patterns in use with ease, for example asynchronous call-return (which includes *rendez-vous* [33] and *future* [2]) and simple message passing. For the space sake we leave their treatment to [12]. Section 6 gives a brief discussion on the significance of specific notations for these fixed communication patterns in language design.

4. REPRESENTING COMMUNICATION (2) COMPLEX COMMUNICATION PATTERNS

This section shows how the structuring primitives can cleanly represent various complex communication patterns which go beyond those represented in conventional communication primitives.

4.1. Continuous Interactions. The traditional call-return primitive already encapsulates a sequence of communications, albeit simple, as a single abstraction unit. The key possibility of the session structure lies in that it extends this idea to arbitrarily complex communication patterns, including the case when multiple interaction sequences interleave with each other. The following shows one such example, which describes the behaviour of a banking service to the user (essentially an automatic teller machine).

Example 4.1 (ATM).

$$\begin{aligned}
\text{ATM}(a, b) = & \text{accept } a(k) \text{ in } k![id]; \\
& k \triangleright \{ \text{deposit} : \text{request } b(h) \text{ in } k?(amt) \text{ in} \\
& \quad h \triangleleft \text{deposit}; h![id, amt]; \text{ATM}[a, b] \\
& \parallel \text{withdraw} : \text{request } b(h) \text{ in } k?(amt) \text{ in} \\
& \quad h \triangleleft \text{withdraw}; h![id, amt]; \\
& \quad h \triangleright \{ \text{success} : k \triangleleft \text{dispense}; k![amt]; \text{ATM}[a, b] \\
& \quad \parallel \text{failure} : k \triangleleft \text{overdraft}; \text{ATM}[a, b] \} \\
& \parallel \text{balance} : \text{request } b(h) \text{ in } h \triangleleft \text{balance}; h?(amt) \text{ in} \\
& \quad k![amt]; \text{ATM}[a, b] \}
\end{aligned}$$

The program, after establishing a session with the user via a , first lets the user input the user code (we omit such details as verification of the code etc.), then offers three choices, **deposit**, **withdraw**, and **balance**. When the user selects one of them, the corresponding code is executed. For example, when the **withdraw** is chosen, the program lets the user enter the amount to be withdrawn, then interacts with the bank via b , asking with the user code and the amount. If the bank answers there is enough amount in the account, the money is given to the user. If not, then the **overdraft** message results. In either case the system eventually returns to the original waiting mode. Note in particular the program should communicate with the bank in the midst of interaction with the user, so three parties are involved in interaction as a whole. A user may be written as:

$$\begin{aligned}
& \text{request } a(k) \text{ in } k![myId]; k \triangleleft \text{withdraw}; k![58]; \\
& \quad k \triangleright \{ \text{dispense} : k?(amt) \text{ in } P \parallel \text{overdraft} : Q \}.
\end{aligned}$$

Here we may consider Q as a code for “exception handling” (invoked when the balance is less than expected). Notice also interactions are now truly reciprocal.

4.2. Unbounded Interactions. The previous example shows how structuring primitives can easily describe the situation which would be difficult to program in a clean way using the conventional primitives. At the same time, the code does have room for amelioration. If a user wants to look at his balance before he withdraws, he should enter his user code twice. The following refinement makes this redundancy unnecessary.

Example 4.2 (Kind ATM).

$$\begin{aligned}
\text{ATM}'(a, b) &= \text{accept } a(k) \text{ in } k![id]; \text{Actions}[a, b, id, k] \\
\text{Actions}(a, b, id, k) &= k \triangleright \{ \text{deposit} : \text{request } b(h) \text{ in } k?(amt) \text{ in} \\
&\quad h \triangleleft \text{deposit}; h![id, amt]; \text{Actions}[a, b, id, k] \\
&\quad \llbracket \text{withdraw} : \text{request } b(h) \text{ in } k?(amt) \text{ in} \\
&\quad h \triangleleft \text{withdraw}; h![id]; h![amt]; \\
&\quad h \triangleright \{ \text{success} : k \triangleleft \text{dispense}; k![amt]; \text{Actions}[a, b, id, k] \\
&\quad \quad \llbracket \text{failure} : k \triangleleft \text{overdraft}; \text{Actions}[a, b, id, k] \rrbracket \\
&\quad \llbracket \text{balance} : \text{request } b(h) \text{ in } h \triangleleft \text{balance}; h?(amt) \text{ in} \\
&\quad k![amt]; \text{Actions}[a, b, id, k] \\
&\quad \llbracket \text{quit} : \text{ATM}'[a, b] \rrbracket \}
\end{aligned}$$

As can be seen, the main difference lies in that the process is still within the same session even after recurring to the waiting mode, after processing each request: once a user establishes the session and enters the user code, she can request various services as many times as she likes. To exit from this loop, the branch `quit` is introduced.

This example shows how recursion within a session allows a flexible description of interactive behaviour which goes beyond the recursion in usual objects (where each session consists of a fixed number of interactions). It is notable that the unbounded session owns a rigorous syntactic structure so as to allow type abstraction, see Section 5. Unbounded interactions with a fixed pattern naturally arise in practice, e.g. interactions between a file server and its client. We believe recursion within a session can be effectively used for varied programming practice.

4.3. Delegation. The original idea of delegation in object-based concurrent programming [40] allows an object to delegate the processing of a request it receives to another object. Its basic purpose is *distribution of processing*, while maintaining the transparency of name space for clients of that service. Practically it can be used to enhance modularity, to express exception handling, and to increase concurrency. The following example shows how we can generalise the original notion to the present setting.

Example 4.3 (Ftp server). Below $\otimes_n P$ denotes the n -fold parallel composition.

$$\begin{aligned}
\text{Init}(pid, nis) &= (\nu b)(\text{Ftpd}[pid, b] \mid \otimes_n \text{FtpThread}[b, nis]) \\
\text{Ftpd}(pid, b) &= \text{accept } pid(s) \text{ in request } b(k) \text{ in throw } k[s]; \text{Ftpd}[pid, b] \\
\text{FtpThread}(b, nis) &= \text{accept } b(k) \text{ in catch } k(s) \text{ in } s?(userid, passwd) \text{ in} \\
&\quad \text{request } nis(j) \text{ in } j \triangleleft \text{checkUser}; j![userid, passwd]; \\
&\quad j \triangleright \{ \text{invalid} : s \triangleleft \text{sorry} \mid \text{valid} : s \triangleleft \text{welcome}; \text{Actions}[s, b] \} \\
\text{Actions}(s, b) &= s \triangleright \{ \text{get} : \dots ; s?(file) \text{ in } \dots ; \text{Actions}[s, b] \\
&\quad \llbracket \text{put} : \dots ; s?(file) \text{ in } \dots ; \text{Actions}[s, b] \\
&\quad \llbracket \text{bye} : \dots ; \text{FtpThread}[b, nis] \rrbracket \}.
\end{aligned}$$

The above code shows an outline of the code of a ftp server, which follows the behaviour of the standard ftp servers with TCP/IP protocol. Initially the program `Init` generates a server `Ftpd` and n threads with the identical behaviour (for simplicity we assume all threads share the same name b). Suppose, in this situation, a server receives a request for a service from some client, establishing the session channel s . A server then requests for another session with an idle thread (if it exists) and “throws” a channel s

to that thread, while getting ready for the next request from clients itself. It is now the thread `FtpThread` which actually processes the user's request, receiving the user name, referring to NIS, and executing various operations (note recursion within a session is used). Here the delegation is used to enable the ftp server to process multiple requests concurrently without undue delay in response. The scheme is generally applicable to a server interacting with many clients. Some observations follow.

- (1) The example shows how the generalised delegation allows programmers to cleanly describe those interaction patterns which generalise the original form of delegation. Other examples of the usage of delegation abound, for example a file server with geographically distributed sites or a server with multiple services each to be processed by a different sub-server.
- (2) A key factor of the above code is that a *client does not have to be conscious of the delegation which takes place on the server's side*: that is, a client program can be written as if it is interacting with a single entity, for example as follows.

```
request pid(s) in s![myId]; s ▷ {sorry : ... || welcome : ... }
```

Observe that, between the initial `request` and the next sending operation, the `catch/throw` interaction takes place on the server's side: however the client process does not have to be conscious of the event. This shows how delegation enables distribution of computation while maintaining the transparency of the name space.

- (3) If we allow each ftp-thread to be dynamically generated, we can use parallel composition to the same effect, just as the use of “fork” to pass process resources in UNIX. While this scheme has a limitation in that we cannot send a channel to an already running process, it offers another programming method to realise flexible, dynamic communication structures. We also observe that the use of `throw/catch`, or the “fork” mentioned above, would result in complexly woven sequences of interactions, which would become more error-prone than without. In such situations, the type discipline discussed in the next section would become an indispensable tool for programming, where we can algorithmically verify if a program has coherent communication structure and, in particular, if it contains interaction errors.

5. THE TYPE DISCIPLINE

5.1. Preliminaries. The present structuring method allows the clear description of complex interaction structures beyond conventional communication primitives. The more complex the interaction becomes, however, the more difficult it would be to capture the whole interactive behaviour and to write correct programs. The type discipline we shall discuss in this section gives a simple solution to these issues at a basic level. We first introduce the basic notions concerning types, including duality on types which represents complementarity of interactions.

Definition 5.1 (Types). Given *type variables* (t, t', \dots) and *sort variables* (s, s', \dots) , *sorts* (S, S', \dots) and *types* (α, β, \dots) are defined by the following grammar.

$$\begin{aligned}
S &::= \mathbf{nat} \mid \mathbf{bool} \mid \langle \alpha, \bar{\alpha} \rangle \mid s \mid \mu s.S \\
\alpha &::= \downarrow[\tilde{S}]; \alpha \mid \downarrow[\alpha]; \beta \mid \&\{l_1: \alpha_1, \dots, l_n: \alpha_n\} \mid \mathbf{1} \mid \perp \\
&\mid \uparrow[\tilde{S}]; \alpha \mid \uparrow[\alpha]; \beta \mid \oplus\{l_1: \alpha_1, \dots, l_n: \alpha_n\} \mid t \mid \mu t.\alpha
\end{aligned}$$

where, for a type α in which \perp does not occur, we define $\bar{\alpha}$, the *co-type* of α , by:

$$\begin{aligned} \overline{\uparrow[\tilde{S}]; \alpha} &= \downarrow[\tilde{S}]; \bar{\alpha} & \overline{\oplus \{l_i : \alpha_i\}} &= \&\{l_i : \bar{\alpha}_i\} & \overline{\uparrow[\alpha]; \beta} &= \downarrow[\alpha]; \bar{\beta} & \overline{\mathbf{1}} &= \mathbf{1} \\ \overline{\downarrow[\tilde{S}]; \alpha} &= \uparrow[\tilde{S}]; \bar{\alpha} & \overline{\&\{l_i : \alpha_i\}} &= \oplus \{l_i : \bar{\alpha}_i\} & \overline{\downarrow[\alpha]; \beta} &= \uparrow[\alpha]; \bar{\beta} & \overline{t} &= t & \overline{\mu t. \alpha} &= \mu t. \bar{\alpha}. \end{aligned}$$

A *sorting* (resp. a *typing*, resp. a *basis*) is a finite partial map from names and variables to sorts (resp. from channels to types, resp. from process variables to the sequences of sorts and types). We let Γ, Γ', \dots (resp. Δ, Δ', \dots , resp. Θ, Θ', \dots) range over sortings (resp. typings, resp. bases). We regard types and sorts as representing the corresponding regular trees in the standard way [6], and consider their equality in terms of such representation.

A sort of form $\langle \alpha, \bar{\alpha} \rangle$ represents two complementary structures of interaction which are associated with a name (one denoting the behaviour starting with **accept**, another that which starts with **request**), while a type gives the abstraction of interaction to be done through a channel. The co-type of a given type denotes the complementary behaviour of the original type. Note $\bar{\bar{\alpha}} = \alpha$ holds whenever $\bar{\alpha}$ is defined. Let us give a brief description of each form of types.

- The type $\downarrow[\tilde{S}]; \alpha$ represents the behaviour of first inputting values of sorts \tilde{S} , then does the actions of type α ; $\downarrow[\alpha]; \beta$ represents the similar behaviour, which starts with the channel input (catch) instead.
- $\uparrow[\tilde{S}]; \alpha$ and $\uparrow[\alpha]; \beta$ are dual of $\downarrow[\tilde{S}]; \alpha$ and $\downarrow[\alpha]; \beta$, sending values instead of receiving.
- $\&\{l_1 : \alpha_1, \dots, l_n : \alpha_n\}$ shows the branching behaviour: it waits with n options, and behaves as type α_i if i -th action is selected (external choice). $\oplus \{l_1 : \alpha_1, \dots, l_n : \alpha_n\}$ then represents the behaviour which would select one of l_i and then behaves as α_i , according to the selected l_i (internal choice).
- $\mathbf{1}$ represents the inaction, acting as the unit of sequential composition.
- $\mu t. \alpha$ denotes a recursive behaviour, which represents the behaviour of start doing α and, when t is encountered, recurs to α again.
- Finally \perp is a specific type indicating that no further connection is possible at a given name.

The following partial algebra on the set of typings, cf. [11], plays a key role in our typing system.

Definition 5.2 (Type algebra). Typings Δ_1 and Δ_2 are *compatible*, written $\Delta_0 \asymp \Delta_1$, if $\Delta_0(k) = \bar{\Delta}_1(k)$ for all $k \in \text{dom}(\Delta_0) \cap \text{dom}(\Delta_1)$. When $\Delta_0 \asymp \Delta_1$, the *composition of Δ_0 and Δ_1* , written $\Delta_0 \circ \Delta_1$, is given as a typing such that $(\Delta_0 \circ \Delta_1)(k)$ is (1) \perp , if $k \in \text{dom}(\Delta_0) \cap \text{dom}(\Delta_1)$; (2) $\Delta_i(k)$, if $k \in \text{dom}(\Delta_i) \setminus \text{dom}(\Delta_{i+1 \bmod 2})$ for $i \in \{0, 1\}$; and (3) undefined otherwise.

Compatibility means each common channel k is associated with complementary behaviours, thus ensuring the interaction on k to run without errors. When composed, the type for k becomes \perp , preventing further connection at k (note \perp has no co-type). One can check the partial operation \circ is partially commutative and associative.

5.2. Typing System. The main sequent of our typing system has a form

$$\Theta; \Gamma \vdash P \triangleright \Delta$$

which reads: “under the environment $\Theta; \Gamma$, a process P has a typing Δ .” Sorting Γ specifies protocols at the free names of P , while typing Δ specifies P ’s behaviour at its free channels. When P is a program, Θ and Δ become both empty.

Given a typing or a sorting, say Φ , write $\Phi \cdot s : p$ for $\Phi \cup \{s : p\}$ together with the condition that $s \notin \text{dom}(\Phi)$; and $\Phi \setminus s$ for the result of taking off $s : \Phi(s)$ from Φ if $\Phi(s)$

is defined (if not we take Φ itself). Also assume given the evident inference rules for arithmetic and boolean expressions, whose sequent has the form $\Gamma \vdash e \triangleright \alpha$, enjoying the standard properties such as $\Gamma \vdash e \triangleright S$ and $e \downarrow c$ imply $\Gamma \vdash c \triangleright S$. The main definition of this section follows.

Definition 5.3 (Basic typing system). The typing system is defined by the axioms and rules in Figure 1, where we assume the range of Δ in [INACT] and [VAR] contains only $\mathbf{1}$ and \perp .

$$\begin{array}{c}
\text{[ACC]} \frac{\Theta; \Gamma \vdash P \triangleright \Delta \cdot k : \alpha}{\Theta; \Gamma, a : \langle \alpha, \bar{\alpha} \rangle \vdash \mathbf{accept } a(k) \text{ in } P \triangleright \Delta} \quad \text{[REQ]} \frac{\Theta; \Gamma \vdash P \triangleright \Delta \cdot k : \bar{\alpha}}{\Theta; \Gamma, a : \langle \alpha, \bar{\alpha} \rangle \vdash \mathbf{request } a(k) \text{ in } P \triangleright \Delta} \\
\text{[SEND]} \frac{\Gamma \vdash \bar{e} \triangleright \tilde{S} \quad \Theta; \Gamma \vdash P \triangleright \Delta \cdot k : \alpha}{\Theta; \Gamma \vdash k![\bar{e}]; P \triangleright \Delta \cdot k : \uparrow[\tilde{S}]; \alpha} \quad \text{[RCV]} \frac{\Theta; \Gamma \cdot \tilde{x} : \tilde{S} \vdash P \triangleright \Delta \cdot k : \alpha}{\Theta; \Gamma \vdash k?(\tilde{x}) \text{ in } P \triangleright \Delta \cdot k : \downarrow[\tilde{S}]; \alpha} \\
\text{[BR]} \frac{\Theta; \Gamma \vdash P_1 \triangleright \Delta \cdot k : \alpha_1 \quad \dots \quad \Theta; \Gamma \vdash P_n \triangleright \Delta \cdot k : \alpha_n}{\Theta; \Gamma \vdash k \triangleright \{l_1 : P_1 \parallel \dots \parallel l_n : P_n\} \triangleright \Delta \cdot k : \&\{l_1 : \alpha_1, \dots, l_n : \alpha_n\}} \\
\text{[SEL]} \frac{\Theta; \Gamma \vdash P \triangleright \Delta \cdot k : \alpha_j}{\Theta; \Gamma \vdash k \triangleleft l_j; P \triangleright \Delta \cdot k : \oplus\{l_1 : \alpha_1, \dots, l_n : \alpha_n\}} \quad (1 \leq j \leq n) \\
\text{[THR]} \frac{\Theta; \Gamma \vdash P \triangleright \Delta \cdot k : \beta}{\Theta; \Gamma \vdash \mathbf{throw } k[k']; P \triangleright \Delta \cdot k : \uparrow[\alpha]; \beta \cdot k' : \alpha} \quad \text{[CAT]} \frac{\Theta; \Gamma \vdash P \triangleright \Delta \cdot k : \beta \cdot k' : \alpha}{\Theta; \Gamma \vdash \mathbf{catch } k(k') \text{ in } P \triangleright \Delta \cdot k : \downarrow[\alpha]; \beta} \\
\text{[CONC]} \frac{\Theta; \Gamma \vdash P \triangleright \Delta \quad \Theta; \Gamma \vdash Q \triangleright \Delta'}{\Theta; \Gamma \vdash P \mid Q \triangleright \Delta \circ \Delta'} \quad (\Delta \asymp \Delta') \quad \text{[IF]} \frac{\Gamma \vdash e \triangleright \mathbf{bool} \quad \Theta; \Gamma \vdash P \triangleright \Delta \quad \Theta; \Gamma \vdash Q \triangleright \Delta}{\Theta; \Gamma \vdash \mathbf{if } e \text{ then } P \text{ else } Q \triangleright \Delta} \\
\text{[NRES]} \frac{\Theta; \Gamma \cdot a : S \vdash P \triangleright \Delta}{\Theta; \Gamma \vdash (\nu a)P \triangleright \Delta} \quad \text{[CRES]} \frac{\Theta; \Gamma \vdash P \triangleright \Delta \cdot k : \perp}{\Theta; \Gamma \vdash (\nu k)P \triangleright \Delta} \quad \text{[INACT]} \Theta; \Gamma \vdash \mathbf{inact} \triangleright \Delta \\
\text{[VAR]} \frac{\Gamma \vdash \bar{e} \triangleright \tilde{S}}{\Theta, X : \tilde{S}\bar{\alpha}; \Gamma \vdash X[\bar{e}k] \triangleright \Delta \cdot \tilde{k} : \bar{\alpha}} \quad \text{[DEF]} \frac{\Theta; \Gamma \cdot \tilde{x} : \tilde{S} \vdash P \triangleright \tilde{k} : \bar{\alpha} \quad \Theta; \Gamma \vdash Q \triangleright \Delta}{\Theta \setminus X; \Gamma \vdash \mathbf{def } X(\tilde{x}k) = P \text{ in } Q \triangleright \Delta} \quad (\Theta(X) = \tilde{S}\bar{\alpha})
\end{array}$$

FIGURE 1. The typing system

For simplicity, the rule [DEF] is restricted to single recursion, which is easily extendible to multiple recursion. If $\Theta; \Gamma \vdash P \triangleright \Delta$ is derivable in the system, we say P is *typable under $\Theta; \Gamma$ with Δ* , or simply P is *typable*. Some comments on the typing system follow.

- (1) In the typing system, the left-hand side of the turnstile is for shared names and variables (“classical realm”), while the right-hand side is for channels sharable only by two complementary parties (a variant of “linear realm”). It differs from various sorting disciplines in that a channel k is in general ill-sorted, e.g. it may carry an integer at one time and a boolean at another. In spite of this, the manipulation of linear realm by typing algebra ensures linearised usage of channels, as well as preventing interaction errors, cf. Theorem 5.4 below.
- (2) In [THR], the behaviour represented by α for channel k' is actually performed by the process which “catches” k' (note k' cannot occur free in Δ , hence neither in P , by our convention on “.”). To capture the interactions at k' as a whole, $k' : \alpha$ is

added to the linear realm. On the other hand, the rule [CAT] guarantees that the receiving side does use the channel k' as is prescribed. Reading from the conclusions to the antecedents, [THR] and [CAT] together illustrate how k' is “thrown” from the left to the right.

- (3) The simplicity of the typing rules is notable, utilising the explicit syntactic structure of session. In particular it is syntax-directed and has a principal type when we use a version of kinding [26]. It is then easy to show there is a typing algorithm à la ML, which extracts the principal type of a given process iff it is typable. It should be noted that simplicity and tractability of typing rules do *not* mean that the obtainable type information is uninteresting: the resulting type abstraction richly represents the interactive behaviour of programs, as later examples exhibit.

Below we briefly summarise the fundamental syntactic properties of the typing system. We need the following notion: a k -process is a prefixed process with subject k (such as $k![\tilde{e}]; P$ and $\text{catch } k(k') \text{ in } P$). Next, a k -redex is a pair of dual k -processes composed by $|$, i.e. either of forms $(k![\tilde{e}]; P | k?(x) \text{ in } Q)$, $(k \triangleleft l; P | k \triangleright \{l_1 : Q_1 \parallel \dots \parallel l_n : Q_n\})$, or $(\text{throw } k[k']; P | \text{catch } k(k'') \text{ in } Q)$. Then P is an *error* if $P \equiv \text{def } D \text{ in } (\nu \tilde{u})(P'|R)$ where P' is, for some k , the $|$ -composition of *either* two k -processes that do not form a k -redex, *or* three or more k -processes. We then have:

Theorem 5.4.

1. (Invariance under \equiv) $\Theta; \Gamma \vdash P \triangleright \Delta$ and $P \equiv Q$ imply $\Theta; \Gamma \vdash Q \triangleright \Delta$.
2. (Subject reduction) $\Theta; \Gamma \vdash P \triangleright \Delta$ and $P \rightarrow^* Q$ imply $\Theta; \Gamma \vdash Q \triangleright \Delta$.
3. (Lack of run-time errors) *A typable program never reduces into an error.*

The proofs are straightforward due to the syntax-directed nature of the typing rules. See [12] for details. We note that we can easily extend the typing system with ML-like polymorphism for recursion, which is useful for e.g. template processes (such as $\text{def Cell}(cv) = \dots \text{ in Cell}[a \ 42] | \text{Cell}[b \ \text{true}]$), with which all the properties in Theorem 5.4 are preserved. This and other basic extensions are discussed in [12].

5.3. Examples. We give a few examples of typing, taking programs in the preceding sections. We omit the final $\mathbf{1}$ from the type, e.g. we write $\downarrow[\alpha]$ for $\downarrow[\alpha]; \mathbf{1}$. First, the factorial in Example 3.2 is assigned, at its free name, a type $\downarrow[\mathbf{nat}]; \uparrow[\mathbf{nat}]$ (for factorial) and its dual $\uparrow[\mathbf{nat}]; \downarrow[\mathbf{nat}]$ (for its user). Next, the cell in Example 3.3 is given a type $\&\{\text{read} : \uparrow[\alpha], \text{write} : \downarrow[\alpha]\}$ (for the cell) and its dual $\oplus\{\text{read} : \downarrow[\alpha], \text{write} : \uparrow[\alpha]\}$ (for its user). The type of a cell says a cell waits with two options, and, when “read” is selected, it would send an integer and the session ends, and when “write” is selected, it would receive an integer and again the session ends: its dual says a user may do either “read” or “write”, and, according to which of them it selects, it behaves as prescribed.

As a more interesting example, take the “kind ATM” in Example 4.2. Consider $\text{ATM}'[ab]$ under the declaration in the example. Then a typing $a : \langle \alpha, \bar{\alpha} \rangle$, $b : \langle \beta, \bar{\beta} \rangle$ is given to the process, where we set α , which abstracts the interaction with a user, as:

$$\begin{aligned} \alpha \stackrel{\text{def}}{=} & \downarrow[\mathbf{nat}]; \mu t. \&\{\text{deposit} : \downarrow[\mathbf{nat}]; t, \\ & \text{withdraw} : \downarrow[\mathbf{nat}]; \oplus\{\text{dispense} : \uparrow[\mathbf{nat}]; t, \text{overdraft} : t\}, \\ & \text{balance} : \uparrow[\mathbf{nat}]; t, \\ & \text{quit} : \uparrow[\mathbf{nat}]\}, \end{aligned}$$

while $\overline{\beta}$, which abstracts the interaction with the banking system, is given as:

$$\begin{aligned} \overline{\beta} \stackrel{\text{def}}{=} & \oplus \{ \text{deposit} : \uparrow [\mathbf{nat} \ \mathbf{nat}], \\ & \text{withdraw} : \uparrow [\mathbf{nat} \ \mathbf{nat}]; \&\{ \text{success} : \mathbf{1}, \text{failure} : \mathbf{1} \}, \\ & \text{balance} : \uparrow [\mathbf{nat}]; \downarrow [\mathbf{nat}] \}. \end{aligned}$$

Note that the type abstraction is given separately for the user (at a) and the bank (at b), describing the behaviour of ATM' for each of its interacting parties. Also notice how the recursion in interactive behaviour is represented by recursion in types.

As a final example, the ftp server of Example 4.3 is given the following type at its principal name:

$$\downarrow [\mathbf{nat} \ \mathbf{nat}]; \oplus \{ \text{sorry} : \mathbf{1}, \text{welcome} : \mu t. \&\{ \text{get} : \dots; t, \text{put} : \dots; t, \text{bye} : \dots \} \}.$$

This example shows that the throw/catch action is abstracted away in the type with respect to the user (but not in the type with respect to the thread, which we omit) so that the user can interact without concerning himself with the delegation occurring on the other's side. In these ways, not only the type discipline offers the correctness verification of programs at a basic level, but also it gives a clean abstraction of interactive behaviours of programs, which would assist programming activities.

6. DISCUSSIONS

6.1. Implementation Concerns. In the previous sections, we have seen how the session structure enables clean description of complex communication behaviours, employing the synchronous form of interactions as its essential element. For implementation of the primitives, however, the use of asynchronous communication is essential, since the real distributed computing environments are inherently asynchronous. To study such implementation in a formal setting, which should then be applied to the realistic implementation, we consider a translation of the present language primitives into TyCO [18], a sorted summation-less polyadic asynchronous π -calculus with branching structure (which is ultimately translatable into its monadic, branch-less version). The existence of the branching structure makes TyCO an ideal intermediate language. For the space sake we cannot discuss the technical details of the translation, for which the reader may refer to [12]. We only list essential points.

- (1) The translation converts both channels and names into names. Each synchronous interaction (including the branching-selection) is translated into two asynchronous interactions, the second one acting as acknowledgement. This suggests the primitives are amenable for distributed implementation, at least at the rudimentary level.
- (2) In spite of (1) above, the resulting code is far from optimal: as a simple example, if a value is sent immediately after the request operation, clearly the value can be “piggy-backed” to the request message. A related example is the translation of the factorial in Section 3 in comparison with its standard “direct” encoding in Pict or TyCO in a continuation-passing style, cf. [28]. To find the effective, well-founded optimisation methods in this line would be a fruitful subject of study (see 6.2 below).
- (3) The encoding translates the typable programs into well-sorted TyCO codes. It is an intriguing question how we can capture, in a precise way, the well-typedness of the original code at the level of TyCO (this question was originally posed by Simon Gay for a slightly different kind of translation).

6.2. Related Works and Further Issues. In the following, comparisons with a few related works are given along with some of the significant further issues.

First, in the context of conventional concurrent programming languages, the key departure of the proposed framework lies in that the session structure allows us to form an arbitrary complex interaction pattern as a unit of abstraction, rather than being restricted to a fixed repertoire. Examples in Section 4 show how this feature results in clean description of complex communication behaviours which may not be easy to express in conventional languages. For example, if we use, for describing those examples, so-called *concurrent object-oriented languages* [40], whose programming concept based on objects and their interaction is proximate to the present one, we need to divide each series of interactions into multiple chunks of independent communications, which, together with the nesting of method invocations, would make the resulting programs hard to understand. The explicit treatment of session also enables the type-based verification of compatibility of communication patterns, which would facilitate the writing of correct communication-based programs at an elementary level.

In spite of these observations, we believe that various communication primitives in existing programming languages, such as object-invocation and RPC, would not diminish their significance even when the present primitives are incorporated. We already discussed how they would be useful for declaring fixed communication patterns, as well as for saving key strokes (in particular the primitives for simple message passing would better be given among language constructs since their description in the structuring primitives is, if easy, roundabout). Notice we can still maintain the same type discipline by regarding these constructs as standing for combination of structuring primitives, as we did in Section 3. In another vein, the communication structures which we can extract from such declaration would give useful information for performing optimisation.

Secondly, regarding channels as port names used in a specific way, we observe that the session structure is realised by name passing communication. Indeed, the encoding of various interaction structures in π -calculus is the starting point of the present series of study (cf. Section 2 of [32]). In this context, one may observe that the representation of communication structures in the present primitives in Section 3 has a close resemblance with the encoding of functions and objects in π -calculus in, respectively, [20] and [37]. Extraction of fundamental operational structures omnipresently arising in the past study, as well as the consideration of how these structures can be combined in a well-structured fashion, led to the present proposal. A couple of basic observations are possible in this context.

- (1) The well-sorted (synchronous) polyadic π -calculus without summation, which is used in most works in this field¹, is a subsystem of the language constructs as presented in this paper, by representing polyadic synchronous name sending/reception by the combinations of request/data-sending and accept/data-reception, respectively. We already noted that there is an simple encoding in the converse direction; Thus the sheer expressiveness does not change.
- (2) In spite of the above point, there is a basic difference between the π -calculus and the presented programming constructs as a tool for representing interaction structures. On the one hand, π -calculus is simple and, among others, is easy to give semantics based on interactive behaviour of processes. This makes it ideal for the basic

¹The use of (guarded or general) summation for branching, as found in [37], has the merit of leaving no “garbages” after selection in the encoding of branching, but is not vital for operational representation or for reasoning (because the mentioned garbages are strongly bisimilar to inaction). In the present setting, the introduction of explicit branching makes the encoding of branching by name passing itself unnecessary.

theoretical study, as has been already attested in the literature. In particular the semantics of the present language constructs would be studied following the basic ideas which have arisen in the study of π -calculus, including those based on labelled transition (cf. [22]) and those based on non-labelled transition (cf. [?, ?]). On the other hand, the presented structuring method gives better articulation of programs' behaviour than the bare use of π -calculus, especially when a series of reciprocal interactions between two parties is involved, as is in many complex communication-based applications. Notice this is achieved through the *constraint* on the usage of channels, formally stipulated as the well-typedness condition. This constraint indeed prohibits certain possibly interesting ways of using channels (such as allowing racing conditions for channels); but it does lead to the clarification of communication structures in programs, which enhances readability of programs and makes it possible to statically detect incompatibility in communication patterns.

With respect to the second point, it would be interesting to study how we can loosen the present constraint on channels in meaningful ways while still keeping the benefit of the present framework, referring to various possibilities as found in π -calculi, as well as considering practical concerns.

Thirdly, the field of research which is closely related to the present work is the study of various typed programming languages based on π -calculi [7, 27, 34] and of types for π -calculi (see for example [21, 8, 28, 29, 16, 17, 31, 35, 41]). Types impose various constraints on the syntax of processes, and thus restrain the processes to have only "good" interactive behaviours. Varied notions of types have been studied in the context of π -calculus and then have been applied to the setting of programming languages based on π -calculi. Comparisons with the works in this field may be done from two distinct viewpoints, language design concerns and notions of types.

- (1) From the viewpoint of language design, Pict and other π -calculus-based languages use the primitives of (polyadic) asynchronous π -calculus or its variants as the basic language constructs, and build further layers of abstraction on their basis. The present proposal differs in that it incorporates the session-based structure as a fundamental stratum for programming, rather than relying on chains of direct name passing for describing communication behaviour. While the former is operationally translatable into the latter as discussed in 6.1, the very translation of, say, the programming examples in the preceding sections would reveal the significance of the session structure for abstraction concerns. In particular, any such translation should use multiple names for actions belonging to a single session, which damages the clarity and readability of the program. We note that we are far from claiming that the proposed framework would form the sole stratum for high-level communication-based programming: abstraction concerns for distributed computing are so diverse that any single framework cannot meet all purposes. However we do believe that the proposed constructs (with possible variations) would offer a basic and useful building block for communication-based programming, especially when concerned communication behaviours tend to become complex.
- (2) From the viewpoint of type disciplines, one notable point of the present notion of types would be that well-typed programs in the present typing system are in general ill-sorted (in the sense of [21]) when they are regarded as π -terms, since the same channel k may be used for carrying values of different sorts at different occasions. This is in a sharp contrast with most type disciplines for π -calculus and π -calculus-based programming languages in literature. An exception is the typing system for monadic π -calculus presented in [41], where, as in the present setting, the

incorporation of sequencing information in types allows certain ill-sorted processes to be well-typed. Apart from a notable difference in motivation, a main technical difference is that the rules in [41] are not syntax-directed so that the type inference algorithm may not be as straightforward as in the present setting; at the same time, the type discipline in [41] guarantees a much stronger behavioural property. We also note that a typing system for ill-sorted processes is also studied in [17] in the context of polymorphism. In another context, we observe that one notable feature of the present typing system is the linear usage of channels it imposes on programs. In this context, the preceding study on linearity in π -calculus such as [11, 16] offers the clarification of the deterministic character of interaction at channels (notice interaction at names is still non-deterministic in general). Also [THR] and [CAT] rules have some resemblance to the rules for linear name communication presented in [11, 16]. Regarding these and other works on types for π -calculus, we note that various cases of redundant codes in translation mentioned in 6.1 above often concern certain fixed ways of using names in processes, which would be amenable to type-based analysis.

Finally one of the most important topics which we could not touch in the present paper is the development of reasoning principles based on the proposed structuring method. Notice the type discipline in Section 5 already gives one simple, though useful, example. However it is yet to be studied whether reasoning methods on deeper properties of programs (cf. [1, 4, 10, 14, 19, 24]), which are both mathematically well-founded and are applicable to conspicuous practical situations, can be developed or not. We wish to touch on this topic in our future study.

Acknowledgements. We sincerely thank three reviewers for their constructive comments. Discussions with Simon Gay, Luís Lopes, Benjamin Pierce, Kaku Takeuchi and Nobuko Yoshida have been beneficial, for which we are grateful. We also thank the last two for offering criticisms on an earlier version of the paper. Vasco Vasconcelos is partially supported by Project PRAXIS/2/2.1/MAT/46/94.

REFERENCES

- [1] S. Abramsky, S. Gay, and R. Nagarajan, Interaction Categories and Foundations of Typed Concurrent Computing. *Deductive Program Design*, Springer-Verlag, 1995.
- [2] G. Agha. *Actors: A Model of Concurrent Computation in Distributed Systems*. MIT Press, 1986.
- [3] P. America. Operational Semantics of a Parallel Object-Oriented Language. *POPL'86*, ACM Press, 1986.
- [4] P. America and F. de Boer. Reasoning about Dynamically Evolving Process Structures. *Formal Aspects of Computing*, 94:269–316, 1994.
- [5] G. Berry and G. Boudol. The chemical abstract machine. *TCS*, 96:217–248, 1992.
- [6] B. Courcelle. Fundamental properties of infinite trees. *TCS*, 25:95–169, 1983.
- [7] C. Fournet and G. Gonthier. The reflexive chemical abstract machine and the join-calculus. *POPL'96*, pp. 372–385, ACM Press, 1996.
- [8] Gay, S., A Sort Inference Algorithm for the Polyadic π -Calculus. *POPL'93*, ACM Press, 1993.
- [9] C.A.R. Hoare. Communicating sequential processes. *CACM*, 21(8):667–677, 1978.
- [10] C.A.R. Hoare. *Communicating Sequential Processes*. Prentice Hall, 1995.
- [11] K. Honda. Composing processes. *POPL'96*, pp. 344–357, ACM Press, 1996.
- [12] K. Honda, V. Vasconcelos, and M. Kubo. Language primitives and type disciplines for structured communication-based programming. Full version of this paper, in preparation.
- [13] C.B. Jones. *Process-Algebraic Foundations for an Object-Based Design Notation*. UMCS-93-10-1, Computer Science Department, Manchester University, 1993.
- [14] C.B. Jones. Accommodating Interference in the Formal Design of Concurrent Object-Based Programs. *Formal Methods in System Design*, 8(2): 105-122, March 1996
- [15] S. P. Jones, A. Gordon, and S. Finne. Concurrent Haskell. *POPL'96*, pp. 295–308, ACM Press, 1996.

- [16] N. Kobayashi, B. Pierce, and D. Turner. Linearity and the pi-calculus. *POPL'96*, pp. 358–371, ACM Press, 1996.
- [17] X. Liu and D. Walker. A polymorphic type system for the polyadic pi-calculus. *CONCUR'95*, LNCS, Springer, 1995.
- [18] L. Lopes, F. Silva, and V. Vasconcelos. *A framework for compiling object calculi*. In preparation.
- [19] R. Milner. *Communication and Concurrency*. C.A.R. Hoare Series Editor. Prentice Hall, 1989.
- [20] R. Milner. Functions as processes. *MSCS*, 2(2):119–141, 1992.
- [21] R. Milner, Polyadic π -Calculus: a tutorial. *Logic and Algebra of Specification*, Springer-Verlag, 1992.
- [22] R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes, Parts I and II. *Journal of Information and Computation*, 100:1–77, September 1992.
- [23] R. Milner and M. Tofte. *The Definition of Standard ML*. MIT Press, 1991.
- [24] H. R. Nielson and F. Nielson. Higher-order concurrent programs with finite communication topology. *POPL '94*. ACM Press, 1994.
- [25] A. Ogori and K. Kato. Semantics for communication primitives in a polymorphic language. *POPL 93*. ACM Press, 1993.
- [26] A. Ogori. A compilation method for ML-style polymorphic record calculi. *POPL 92*, pp. 154–165. ACM Press, 1992.
- [27] B. Pierce and D. Turner. Pict: A programming language based on the pi-calculus. CSCI Technical Report 476, Indiana University, March 1997.
- [28] B. Pierce, and D. Sangiorgi, Typing and subtyping for mobile processes. *LICS'93*, IEEE, pp. 187–215, 1993.
- [29] B. Pierce and D. Sangiorgi, Behavioral Equivalence in the Polymorphic Pi-Calculus. *POPL 97*, ACM Press, 1997.
- [30] J-H. Reppy. CML: A higher-order concurrent language. *PLDI 91*, pp. 293–305. ACM Press, 1991.
- [31] M. Müller, J. Niehren and G. Smolka, *Typed Concurrent Programming with Logic Variables*, technical report, Programming Systems Lab, Universität des Saarlandes, January 1998.
- [32] K. Takeuchi, K. Honda, and M. Kubo. An interaction-based programming language and its typing system. *PARLE'94*, volume 817 of LNCS, pp. 398–413. Springer-Verlag, July 1994.
- [33] US Government Printing Office, Washington DC. *The Ada Programming Language*, 1983.
- [34] V. Vasconcelos. Typed concurrent objects. *ECOOP'94*, volume 821 of LNCS, pp. 100–117. Springer-Verlag, 1994.
- [35] V. Vasconcelos and K. Honda, Principal Typing Scheme for Polyadic π -Calculus. *CONCUR'93*, Volume 715 of LNCS, pp.524-538, Springer-Verlag, 1993.
- [36] V. Vasconcelos and M. Tokoro. A typing system for a calculus of objects. *1st ISOTAS*, volume 742 of LNCS, pp. 460–474. Springer-Verlag, November 1993.
- [37] D. Walker, Objects in the π -calculus. *Information and Computation*, Vol. 116, pp.253–271, 1995.
- [38] Y. Yokote and M. Tokoro. Concurrent programming in ConcurrentSmalltalk. In [40].
- [39] A. Yonezawa, editor. *ABCL, an Object-Oriented Concurrent System*. MIT Press, 1990.
- [40] A. Yonezawa and M. Tokoro, editors. *Object-Oriented Concurrent Programming*. MIT Press, 1987.
- [41] N. Yoshida, Graph Types for Monadic Mobile Processes. *FST/TCS'16*, volume 1180 of LNCS, pp. 371–386, Springer-Verlag, 1996. The full version as LFCS Technical Report, ECS-LFCS-96-350, 1996.