

Parameterised Multiparty Session Types ^{*}

Nobuko Yoshida, Pierre-Malo Deniérou, Andi Bejleri, and Raymond Hu

Department of Computing, Imperial College London

Abstract. For many application-level distributed protocols and parallel algorithms, the set of participants, the number of messages or the interaction structure are only known at run-time. This paper proposes a dependent type theory for multiparty sessions which can statically guarantee type-safe, deadlock-free multiparty interactions among processes whose specifications are parameterised by indices. We use the primitive recursion operator from Gödel’s System \mathcal{T} to express a wide range of communication patterns while keeping type checking decidable. We illustrate our type theory through non-trivial programming and verification examples taken from parallel algorithms and Web services usecases.

1 Introduction

As the momentum around communications-based computing grows, the need for effective frameworks to globally *coordinate* and *structure* the application-level interactions is pressing. The structures of interactions are naturally distilled as *protocols*. Each protocol describes a bare skeleton of how interactions should proceed, through e.g. sequencing, choices and repetitions. In the theory of multiparty session types [3, 5, 13], such protocols can be captured as types for interactions, and type checking can statically ensure runtime safety and fidelity to a stipulated protocol.

One of the particularly challenging aspects of protocol descriptions is the fact that many actual communication protocols are highly *parametric* in the sense that the number of participants and even the interaction structure itself are not fixed at design time. Examples include parallel algorithms such as the Fast Fourier Transform (run on any number of communication nodes depending on resource availability) and Web services such as business negotiation involving an arbitrary number of sellers and buyers. This paper introduces a robust dependent type theory which can statically ensure communication-safe, deadlock-free process interactions which follow parameterised multiparty protocols.

We illustrate the key ideas of our proposed parametric type structures through examples. Let us first consider a simple protocol where participant `Alice` sends a message of type `nat` to participant `Bob`. To develop the code for this protocol, we start by specifying the global type, which can concisely and clearly describe a high-level protocol for multiple participants [3, 13, 17], as follows (below `end` denotes protocol termination):

$$G_1 = \text{Alice} \rightarrow \text{Bob} : \langle \text{nat} \rangle . \text{end}$$

Upon agreement on G_1 as a specification for `Alice` and `Bob`, each program can be implemented separately. For type-checking, G_1 is *projected* into end-point session types:

^{*} The work is partially supported by EPSRC EP/G015635/1 and EP/F003757/1.

one from Alice’s point of view, $!(\text{Bob}, \text{nat})$ (output to Bob with nat -type), and another from Bob’s point of view, $?(\text{Alice}, \text{nat})$ (input from Alice with nat -type), against which the respective Alice and Bob programs are checked to be compliant.

The first step towards generalised type structures for multiparty sessions is to allow modular specifications of protocols using arbitrary compositions and repetitions of interaction units (this is a standard requirement in multiparty contracts [21]). Consider the type $G_2 = \text{Bob} \rightarrow \text{Carol} : \langle \text{nat} \rangle . \text{end}$. The designer may wish to compose sequentially G_1 and G_2 together to build a larger protocol:

$$G_3 = G_1; G_2 = \text{Alice} \rightarrow \text{Bob} : \langle \text{nat} \rangle . \text{Bob} \rightarrow \text{Carol} : \langle \text{nat} \rangle . \text{end}$$

We may also want to iterate the composed protocols n -times, which can be written by $\text{foreach}(i < n)\{G_1; G_2\}$, and moreover bind the number of iteration n by a dependent product to build a *family of global specifications*, as in:

$$\Pi n. \text{foreach}(i < n)\{G_1; G_2\} \tag{1}$$

Beyond enabling a variable number of exchanges between a fixed set of participants, the ability to parameterise *participant identities* can represent a wide class of the communication topologies found in the literature. For example, the use of indexed participants $W[i]$ (denoting the i -th worker) allows to specify a family of session types such that neither the number of participants nor message exchanges are known before the runtime instantiation of the parameters. The following type and diagram both describe a sequence of messages from $W[n]$ to $W[0]$ (indices decrease in our `foreach`, see § 2):

$$\Pi n. (\text{foreach}(i < n)\{W[i+1] \rightarrow W[i] : \langle \text{nat} \rangle\}) \quad \boxed{n} \rightarrow \boxed{n-1} \rightarrow \dots \rightarrow \boxed{0} \tag{2}$$

Here we face an immediate question: *what is the underlying type structure for such parametrisation, and how can we type-check each (parametric) end-point program?* The type structure should allow the projection of a parameterised global type to an end-point type *before* knowing the exact shape of the concrete topology. In (2), if $n \geq 2$, there are three distinct *communication patterns* inhabiting this specification: the initiator (send only), the $n - 1$ middle workers (receive and send), and the last worker (receive only). This is no longer the case when $n = 1$ (there is only the initiator and the last worker) or when $n = 0$ (no communication). Can we provide a decidable projection and static type-checking by which we can preserve the main properties of the session types such as progress and communication-safety in parameterised process topologies? The key technique proposed in this paper is a projection method from a dependent global type onto a *generic end-point generator* which exactly captures the interaction structures of parameterised end-points and which can represent the class of all possible end-point types.

The main contributions of this paper follow:

- *A new expressive framework to globally specify and program* a wide range of parametric communication protocols (§ 2). We achieve this result by combining dependent type theories derived from Gödel’s System \mathcal{T} [18] (for expressiveness) and indexed dependent types from [22] (for tractability to control parameters), with multiparty session types.
- *Decidable and flexible projection methods* based on a generic end-point generator and mergeability of branching types, enlarging the typability (§ 3.1).

$i ::= i \mid n \mid i \text{ op } i'$	Indices	$G ::=$	Global types
$P ::= P \wedge P \mid i \leq i'$	Propositions	$\mid p \rightarrow p' : \langle U \rangle . G$	Message
$I ::= \text{nat} \mid \{i : I \mid P\}$	Index sorts	$\mid p \rightarrow p' : \{l_k : G_k\}_{k \in K}$	Branching
$\mathcal{P} ::= \text{Alice} \mid \text{Worker} \mid \dots$	Participants	$\mid \mu x . G$	Recursion
$p ::= p[i] \mid \mathcal{P}$	Principals	$\mid \mathbf{R} G \lambda i : I . \lambda x . G'$	Primitive recursion
$S ::= \text{nat} \mid \langle G \rangle$	Value type	$\mid x$	Type variable
$U ::= S \mid T$	Payload type	$\mid G \ i$	Application
$K ::= \{n_0, \dots, n_k\}$	Finite integer set	$\mid \text{end}$	Null

$$\begin{aligned} \mathbf{R} G \lambda i : I . \lambda x . G' \ 0 &\longrightarrow G \\ \mathbf{R} G \lambda i : I . \lambda x . G' \ (n+1) &\longrightarrow G' \{n/i\} \{ \mathbf{R} G \lambda i : I . \lambda x . G' \ n/x \} \end{aligned}$$

Fig. 1. Global types and type reduction

- *A dependent typing system* that treats the full multiparty session types integrated with dependent types. The resulting static typing system allows decidable type-checking and guarantees type-safety and deadlock-freedom for well-typed processes involved in parameterised multiparty communication protocols (§ 3).
- *Applications* featuring various process topologies, including the complex butterfly network for the parallel FFT algorithm (§ 2.3, § 3.6). As far as we know, this is the first time such a complex protocol is specified by a single *type* and that its implementation can be automatically type-checked to prove communication-safety and deadlock-freedom. We also extend the calculus with a new asynchronous join primitive for session initialisation, applied to Web services use cases [19] (§ 3.6).

The complete formal definition of our system, including proofs and additional material for examples and implementations can be found in [11].

2 Types and processes for parameterised multiparty sessions

2.1 Global types

Global types allow the description of the parameterised conversations of multiparty sessions as a type signature. Our type syntax integrates three different formulations: (1) global types from [3]; (2) dependent types with primitive recursive combinators based on [18]; and (3) parameterised dependent types from a simplified Dependent ML [1, 22].

The grammar of global types (G, G', \dots) is given in figure 1. *Parameterised principals* p, p', q, \dots can be indexed by one or more parameters, e.g. $\text{Worker}[5][i+1]$. Index i ranges over index variables i, j, n , naturals n or arithmetic operations. A global interaction can be a message exchange ($p \rightarrow p' : \langle U \rangle . G$), where p, p' denote the sending and receiving principals, U the payload type of the message and G the subsequent interaction. Payload types U are either value types S (which contain base type nat and session channel types $\langle G \rangle$), or *end-point types* T (which correspond to the behaviour of one of the session participants and will be explained in § 3) for delegation. Branching ($p \rightarrow p' : \{l_k : G_k\}_{k \in K}$) allows the session to follow one of the different G_k paths in




Fig. 2. Parameterised multiparty protocol on a mesh topology

the interaction (K is a ground and finite set of integers). $\mu\mathbf{x}.G$ is a recursive type where type variable \mathbf{x} is guarded in the standard way.

The interesting addition is the primitive recursion operator $\mathbf{R} G \lambda i : I. \lambda \mathbf{x}. G'$ from Gödel's System \mathcal{T} [12] whose reduction semantics is given in figure 1. Its parameters are a global type G , an index variable i with range I , a type variable for recursion \mathbf{x} and a recursion body G' .¹ When applied to an index i , its semantics corresponds to the repetition i -times of the body G' , with the index variable i value going down by one at each iteration, from $i - 1$ to 0. The final behaviour is given by G when the index reaches 0. The index sorts comprise the set of natural numbers and its restrictions by predicates (P, P', \dots) that are, in our case, conjunctions of inequalities. op represents first-order indices operators (such as $+$, $-$, $*$, \dots). We often omit I and end in our examples. Using \mathbf{R} , we define the product, composition, repetition and test operators (seen in § 1):

$$\begin{array}{l} \Pi i. G = \mathbf{R} \text{end} \lambda i. \lambda \mathbf{x}. G\{i + 1/i\} \quad \bigg| \quad \text{foreach}(i < j)\{G\} = \mathbf{R} \text{end} \lambda i. \lambda \mathbf{x}. G\{\mathbf{x}/\text{end}\} j \\ G_1; G_2 = \mathbf{R} G_2 \lambda i. \lambda \mathbf{x}. G_1\{\mathbf{x}/\text{end}\} 1 \quad \bigg| \quad \text{if } j \text{ then } G_1 \text{ else } G_2 = \mathbf{R} G_2 \lambda i. \lambda \mathbf{x}. G_1 j \end{array}$$

where we assume that \mathbf{x} is not free in G and G_1 , and that the leaves of the syntax trees of G_1 and G are end . These definitions rely on a special substitution of each end by \mathbf{x} (for example, $\text{p} \rightarrow \text{p}'\{l_1:!(\text{nat}); \text{end}, l_2:\text{end}\}\{\mathbf{x}/\text{end}\} = \text{p} \rightarrow \text{p}'\{l_1:!(\text{nat}); \mathbf{x}, l_2:\mathbf{x}\}$). The composition operator appends the execution of G_2 to G_1 ; the repetition operator above repeats G j -times²; the boolean values are integers 0 (**false**) and 1 (**true**). Similar syntactic sugar is defined for local types and processes. Note that composition and repetition do not necessarily impose sequentiality: only the order of the asynchronous messages and the possible dependency [13] between receivers and subsequent senders controls the sequentiality. A parallel version of the sequence example of (§ 1 (2)) can be written: $\Pi n. (\text{foreach}(i < n)\{W[n - i] \rightarrow W[n - i - 1]: \langle \text{nat} \rangle\})$.

Mesh example The session presented in figure 2 describes a particular protocol over a standard mesh topology [15]. In this two dimensional array of workers W , each worker receives messages from his left and top neighbours (if they exist) before sending messages to his right and bottom (if they exist). Our session takes two parameters n and m which represent the number of rows and the number of columns. Then we have two iterators that repeat $W[i + 1][j + 1] \rightarrow W[i][j + 1]: \langle \text{nat} \rangle$ and $W[i + 1][j + 1] \rightarrow W[i + 1][j]: \langle \text{nat} \rangle$

¹ We distinguish recursion and primitive recursion in order to get decidability results, see § 3.4.

² This version of **foreach** uses decreasing indices. One can write an increasing version [11].

$c ::= y \mid s[\mathbf{p}]$	Channels	$\hat{\mathbf{p}}, \hat{\mathbf{q}} ::= \hat{\mathbf{p}}[n] \mid \mathcal{P}$	Principal values
$u ::= x \mid a$	Identifiers	$m ::= (\hat{\mathbf{q}}, \hat{\mathbf{p}}, v) \mid (\hat{\mathbf{q}}, \hat{\mathbf{p}}, s[\hat{\mathbf{p}}']) \mid (\hat{\mathbf{q}}, \hat{\mathbf{p}}, l)$	Messages in transit
$v ::= a \mid n$	Values	$h ::= \epsilon \mid m \cdot h$	Queue types
$e ::= \mathbf{i} \mid v \mid x \mid s[\mathbf{p}] \mid e \text{ op } e'$	Expressions		
$P ::=$	Processes	$\mu X.P$	Recursion
$\mid \bar{u}[\mathbf{p}_0, \dots, \mathbf{p}_n](y).P$	Init	$\mathbf{0}$	Inaction
$\mid u[\mathbf{p}](y).P$	Accept	$P \mid Q$	Parallel
$\mid \bar{a}[\mathbf{p}] : s$	Request	$\mathbf{R} P \lambda i. \lambda X. Q$	Primitive recursion
$\mid c!\langle \mathbf{p}, e \rangle; P$	Value sending	X	Process variable
$\mid c?\langle \mathbf{p}, x \rangle; P$	Value reception	$(P \mathbf{i})$	Application
$\mid c \oplus \langle \mathbf{p}, l \rangle; P$	Selection	$(\nu s)P$	Session restriction
$\mid c\&\langle \mathbf{p}, \{l_k : P_k\}_{k \in K} \rangle$	Branching	$s:h$	Queues

Fig. 3. Syntax for user-defined and run-time processes

for all i and j . The communication flow goes from the top-left $W[n][m]$ and converges towards the bottom-right $W[0][0]$ in $n + m$ parallel message exchanges.

2.2 Process syntax and semantics

Syntax The syntax of expressions and processes is given in figure 3, extended from [3], adding the primitive recursion operator and a new request process. Identifiers u can be variables x or channel names a . Values v are either channels a or natural numbers n . Expressions e are built out of indices \mathbf{i} , values v , variables x , session end points (for delegation) and operations over expressions. In processes, sessions are asynchronously initiated by $\bar{u}[\mathbf{p}_0, \dots, \mathbf{p}_n](y).P$. It spawns, for each of the $\{\mathbf{p}_0, \dots, \mathbf{p}_n\}$, a request that is accepted by the participant through $u[\mathbf{p}](y).P$. Messages are sent by $c!\langle \mathbf{p}, e \rangle; P$ to the participant \mathbf{p} and received by $c?\langle \mathbf{q}, x \rangle; P$ from the participant \mathbf{q} . Selection $c \oplus \langle \mathbf{p}, l \rangle; P$, and branching $c\&\langle \mathbf{q}, \{l_k : P_k\}_{k \in K} \rangle$, allow a participant to choose a branch from those supported by another. Standard language constructs include recursive processes $\mu X.P$, restriction $(\nu s)P$ and parallel composition $P \mid Q$. The primitive recursion operator $\mathbf{R} P \lambda i. \lambda X. Q$ takes as parameters a process P , a function taking an index parameter i and a recursion variable X . A queue $s : h$ stores the asynchronous messages in transit.

An *annotated* P is the result of annotating P 's bound names and variables as in e.g. $(\nu a : \langle G \rangle)Q$ or $s?(x : \langle G \rangle)Q$ or $\mathbf{R} Q \lambda i : I. \lambda X. Q'$. We omit the annotations unless needed. We often omit $\mathbf{0}$ and the participant \mathbf{p} from the session primitives. Requests, session hiding and channel queues appear only at runtime, as explained below.

Semantics The semantics is defined by the reduction relation \longrightarrow presented in figure 4. The standard definition of evaluation contexts (that allow e.g. $W[3+1]$ to be reduced to $W[4]$) is omitted. The metavariables $\hat{\mathbf{p}}, \hat{\mathbf{q}}, \dots$ range over principal values (where all indices have been evaluated). [ZeroR] and [SuccR] are standard and identical to their global type counterparts. The rule [Init] describes the initialisation of a session by its first participant $\bar{a}[\mathbf{p}_0, \dots, \mathbf{p}_n](y_0).P_0$. It spawns asynchronous requests $\bar{a}[\hat{\mathbf{p}}_k] : s$ that allow delayed acceptance by the other session participants (rule [Join]). After the connection, the participants share the private session name s , and the queue associated to s (which

$$\begin{array}{l}
\mathbf{R} \ P \ \lambda i. \lambda X. Q \ 0 \longrightarrow P \quad \text{[ZeroR]} \\
\mathbf{R} \ P \ \lambda i. \lambda X. Q \ n + 1 \longrightarrow Q\{n/i\}\{\mathbf{R} \ P \ \lambda i. \lambda X. Q \ n/X\} \quad \text{[SuccR]} \\
\bar{a}[\hat{p}_0, \dots, \hat{p}_n](y).P \longrightarrow (\nu s)(P\{s[\hat{p}_0]/y\} \mid s : \emptyset \mid \bar{a}[\hat{p}_1] : s \mid \dots \mid \bar{a}[\hat{p}_n] : s) \quad \text{[Init]} \\
\bar{a}[\hat{p}_k] : s \mid a[\hat{p}_k](y_k).P_k \longrightarrow P_k\{s[\hat{p}_k]/y_k\} \quad \text{[Join]} \\
s[\hat{p}]!\langle \hat{q}, v \rangle; P \mid s : h \longrightarrow P \mid s : h \cdot (\hat{p}, \hat{q}, v) \quad \text{[Send]} \\
s[\hat{p}] \oplus \langle \hat{q}, l \rangle; P \mid s : h \longrightarrow P \mid s : h \cdot (\hat{p}, \hat{q}, l) \quad \text{[Label]} \\
s[\hat{p}]?(\hat{q}, x); P \mid s : (\hat{q}, \hat{p}, v) \cdot h \longrightarrow P\{v/x\} \mid s : h \quad \text{[Recv]} \\
s[\hat{p}]\&z(\hat{q}, \{l_k : P_k\}_{k \in K}) \mid s : (\hat{q}, \hat{p}, l_{k_0}) \cdot h \longrightarrow P_{k_0} \mid s : h \quad (k_0 \in K) \quad \text{[Branch]}
\end{array}$$

Fig. 4. Reduction rules

is initially empty by rule [Init]). The variables y_p in each participant p are then replaced with the corresponding session channel, $s[p]$. A more verbose, but symmetric, version of [Init] (where any participant can start the session, not only p_0) could also be used [11].

The rest of the session reductions are standard [3, 13]. The output rules [Send] and [Label] push values, channels and labels into the queue of the session s . The rules [Recv] and [Branch] perform the complementary operations. Note that these operations check that the sender and receiver match. Processes are considered modulo structural equivalence, denoted by \equiv (in particular, we note $\mu X.P \equiv P\{\mu X.P/X\}$).

2.3 Processes for parameterised multiparty protocols

We give here the processes corresponding to the interactions described in § 1 and § 2.1, then introduce a parallel implementation of the Fast Fourier Transform algorithm.

Sequence from § 1 (2) The process below generates all participants using a recursor:

$$\begin{array}{l}
\text{In.}(\text{if } n = 0 \text{ then } \mathbf{0} \\
\quad \text{else } (\mathbf{R} \ (\bar{a}[\mathbf{w}[n], \dots, \mathbf{w}[0]](y).y!\langle \mathbf{w}[n-1], v \rangle; \mathbf{0} \\
\quad \quad \mid a[\mathbf{w}[0]](y).y?(\mathbf{w}[1], z); \mathbf{0} \\
\quad \quad \lambda i. \lambda X. (a[\mathbf{w}[i+1]](y).y?(\mathbf{w}[i+2], z); y!\langle \mathbf{w}[i], z \rangle; \mathbf{0} \mid X) \quad n-1)
\end{array}$$

When $n = 0$ no message is exchanged. In the other case, the recursor creates the $n-1$ workers through the main loop and finishes by spawning the initial and final ones.

As an illustration of the semantics, we show the reduction of the above process for $n = 2$. After several applications of the [SuccR] and [ZeroR] rules, we have:

$$\bar{a}[\mathbf{w}[2], \mathbf{w}[1], \mathbf{w}[0]](y).y!\langle \mathbf{w}[1], v \rangle; \mid a[\mathbf{w}[0]](y).y?(\mathbf{w}[1], z); \mid a[\mathbf{w}[1]](y).y?(\mathbf{w}[2], z); y!\langle \mathbf{w}[0], z \rangle;$$

which, with [Init], [Join], [Send], [Recv], gives:

$$\begin{array}{l}
\longrightarrow (\nu s)(s : \epsilon \mid s[\mathbf{w}[2]]!\langle \mathbf{w}[1], v \rangle; \mid \bar{a}[\mathbf{w}[1]] : s \mid \bar{a}[\mathbf{w}[0]] : s \mid \\
\quad a[\mathbf{w}[0]](y).y?(\mathbf{w}[1], z); \mid a[\mathbf{w}[1]](y).y?(\mathbf{w}[2], z); y!\langle \mathbf{w}[0], z \rangle;) \\
\longrightarrow (\nu s)(s : \epsilon \mid s[\mathbf{w}[2]]!\langle \mathbf{w}[1], v \rangle; \mid \bar{a}[\mathbf{w}[1]] : s \mid \\
\quad s[\mathbf{w}[0]]?(\mathbf{w}[1], z); \mid a[\mathbf{w}[1]](y).y?(\mathbf{w}[2], z); y!\langle \mathbf{w}[0], z \rangle;) \\
\longrightarrow^* (\nu s)(s : \emptyset \mid s[\mathbf{w}[2]]!\langle \mathbf{w}[1], v \rangle; \mid s[\mathbf{w}[0]]?(\mathbf{w}[1], z); \mid s[\mathbf{w}[1]]?(\mathbf{w}[2], z); s[\mathbf{w}[1]]!\langle \mathbf{w}[0], z \rangle;) \\
\longrightarrow^* (\nu s)(s : \emptyset \mid s[\mathbf{w}[0]]?(\mathbf{w}[1], z); \mid s[\mathbf{w}[1]]!\langle \mathbf{w}[0], v \rangle;) \\
\longrightarrow^* \equiv \mathbf{0}
\end{array}$$




Fig. 5. Fast Fourier Transform on a butterfly network topology

Mesh from figure 2 The mesh example is more complex: when n and m are bigger than 2, there are 9 distinct roles that each have a different pattern of communication. We only list processes for (1) the centre workers $W[i][j]$ ($0 < i < n$, $0 < j < m$) who are connected in all four directions, (2) the initiator $W[n][m]$ from the top-left corner. Below, $f(i, j)$ represents the expression computed at the (i, j) -th element.

$$\begin{aligned}
 P_{\text{centre}}(i, j) &= a[W[i][j]](y).y?(W[i+1][j], z_1); y?(W[i][j+1], z_2); \\
 &\quad y!(W[i-1][j], f(i-1, j)); y!(W[i][j-1], f(i, j-1)); \mathbf{0} \\
 P_{\text{start}}(n, m) &= \bar{a}[W[0][0]].W[n][m]](y).y!(W[n-1][m], f(n-1, m));
 \end{aligned}$$

FFT We describe a parallel implementation of the Fast Fourier Transform algorithm (more precisely the radix-2 variant of the Cooley-Tukey algorithm [10]).

Figure 5(a) illustrates the recursive principle of the algorithm, called *butterfly*, where two different outputs can be computed in constant time from the results of the same two recursive calls. The complete algorithm is illustrated by the diagram from figure 5(b). It features the application of the FFT on a network of $N = 2^3$ machines on an hypercube network computing the discrete Fourier transform of vector x_0, \dots, x_7 . Each row represents a single machine at each step of the algorithm. Each edge represents a

$T ::=$	End-point types		$\mu x.T$	Recursion
$!\langle p, U \rangle; T$	Output		$\mathbf{R} T \lambda i : I. \lambda x. T'$	Primitive recursion
$?\langle p, U \rangle; T$	Input		x	Type variable
$\oplus\langle p, \{l_k : T_i\}_{k \in K} \rangle$	Selection		$T \ i$	Application
$\&\langle p, \{l_k : T_i\}_{k \in K} \rangle$	Branching		end	End

Fig. 6. End-point types

$$\begin{array}{l}
p \rightarrow p' : \langle U \rangle. G \upharpoonright q = \text{if } q=p=p' \text{ then } !\langle p, U \rangle; ?\langle p, U \rangle; G \upharpoonright q \\
\quad \text{else if } q=p \text{ then } !\langle p', U \rangle; G \upharpoonright q \\
\quad \text{else if } q=p' \text{ then } ?\langle p, U \rangle; G \upharpoonright q \\
\quad \text{else } G \upharpoonright q \\
p \rightarrow p' : \{l_k : G_k\}_{k \in K} \upharpoonright q = \text{if } q=p \text{ then } \oplus\langle p', \{l_k : G_k \upharpoonright q\}_{k \in K} \rangle \\
\quad \text{else if } q=p' \text{ then } \&\langle p, \{l_k : G_k \upharpoonright q\}_{k \in K} \rangle \\
\quad \text{else } \sqcup_{k \in K} G_k \upharpoonright q \\
\mathbf{R} G \lambda i : I. \lambda x. G' \upharpoonright q = \mathbf{R} G \upharpoonright q \lambda i : I. \lambda x. G' \upharpoonright q
\end{array}
\left|
\begin{array}{l}
(\mu t. G) \upharpoonright p = \mu t. G \upharpoonright p \\
x \upharpoonright p = x \\
(G \ i) \upharpoonright p = (G \upharpoonright p) \ i \\
\text{end} \upharpoonright p = \text{end}
\end{array}
\right.$$

Fig. 7. Projection of global types to end-point types

value sent to another machine. The dotted edges represent the particular messages that a machine sends to itself to remember a value for the next step. Each machine is successively involved in a butterfly with a machine whose number differs by only one bit. Note that the recursive partition over the value of a different bit at each step requires a particular bit-reversed ordering of the input vector: the machine number p initially receives $x_{\bar{p}}$ where \bar{p} denotes the bit-reversal of p . Figure 5(c) gives the global session type describing the interactions between 2^n machines. The first iterator is the initialisation step. Then we have an iteration over variable l for the n successive steps of the algorithm. Figure 5(d) defines the processes that each of the machines runs. Each process returns the final answer at $r_{\bar{p}}$.

3 Typing parameterised multiparty interactions

3.1 End-point types and end-point projections

The syntax of end-point types is given in figure 6. Output expresses the sending to p of a value or channel of type U , followed by the interactions T . Selection represents the transmission to p of a label l_k chosen in $\{l_k\}_{k \in K}$ followed by T_k . Input and branching are their dual counterparts. The other types are similar to their global versions.

End-point projection: a generic projection The relation between end-point types and global types is formalised by the projection relation. Since the actual participant characteristics might only be determined at runtime, we cannot straightforwardly use the definition from [3, 13]. Instead, we rely on the expressive power of the primitive recursive operator: a *generic end-point projection of G onto q* , written $G \upharpoonright q$, represents the family of all the possible end-point types that a principal q can satisfy at run-time.

The general endpoint generator is defined in figure 7 using the derived construct $\text{if } _ \text{ then } _ \text{ else } _$. The projection $p \rightarrow p' : \langle U \rangle. G \upharpoonright q$ leads to a case analysis: if the

participant q is equal to p , then the end-point type of q is an output of type U to p' ; if participant q is p' then q inputs U from p' ; else we skip the prefix. The first case corresponds to the possibility for the sender and receiver to be identical. Projecting the branching global type is similarly defined, but for the operator \sqcup explained below. For the other cases (as well as for our derived operators), the projection is homomorphic.

Mergeability and injection of branching types We first recall the example from [13], which explains that naïve branching projection leads to inconsistent end-point types.

$$w[0] \rightarrow w[1] : \{ok : w[1] \rightarrow w[2] : \langle bool \rangle, \text{quit} : w[1] \rightarrow w[2] : \langle nat \rangle\}$$

We cannot project the above type onto $w[2]$ because, while the branches behave differently, $w[0]$ makes a choice without informing $w[2]$ who thus cannot know the type of the expected value. A solution is to define projection only when the branches are identical, i.e. we change the above nat to $bool$ in our example above.

In our framework, this restriction is too strong since each branch may contain different parametric interaction patterns. To overcome this, we propose two methods called *mergeability* and *injection* of branching types. Formally, the mergeability relation \bowtie is the smallest congruence relation over end-point types such that:³ if $\forall i \in (K \cap J). T_k \bowtie T'_j$ and $\forall i \in (K \setminus J) \cup (J \setminus K). l_k \neq l_j$, then $\&\langle p, \{l_k : T_k\}_{k \in K} \rangle \bowtie \&\langle p, \{l_j : T'_j\}_{j \in J} \rangle$. When $T_1 \bowtie T_2$ is defined, we define the injection \sqcup as a partial commutative operator over two types such that $T \sqcup T = T$ for all types and that:

$$\begin{aligned} \&\langle p, \{l_k : T_i\}_{k \in K} \rangle \sqcup \&\langle p, \{l_j : T'_j\}_{j \in J} \rangle = \\ \&\langle p, \{l_k : T_k \sqcup T'_k\}_{k \in K \cap J} \cup \{l_k : T_k\}_{k \in K \setminus J} \cup \{l_j : T'_j\}_{j \in J \setminus K} \rangle \end{aligned}$$

The mergeability relation states that two types are identical up to their branching types where only branches with distinct labels are allowed to be different. By this extended typing condition, we can modify our previous global type example to add ok and $quit$ labels to notify $w[2]$. We get:

$$\begin{aligned} w[0] \rightarrow w[1] : \{ok : w[1] \rightarrow w[2] : \{ok : w[1] \rightarrow w[2] : \langle bool \rangle\}, \\ \text{quit} : w[1] \rightarrow w[2] : \{\text{quit} : w[1] \rightarrow w[2] : \langle nat \rangle\}\} \end{aligned}$$

Then $w[2]$ can have the type $\&\langle w[1], \{ok : \langle w[1], bool \rangle, \text{quit} : \langle w[1], nat \rangle\} \rangle$ which could not be obtained through the original projection rule in [3, 13]. This projection is sound up to branching subtyping (cf. Lemma 3.4).

3.2 Type system

This subsection introduces the type system. Because free indices appear both in terms (e.g. participants in session initialisation) and in types, the formal definition of what constitutes a valid term and a valid type are interdependent and both in turn require a careful definition of a valid global type.

³ The idea of mergeability is introduced informally in the tutorial paper [8].

Judgements and environments One of the main differences with previous session type systems is that session environments Δ can contain dependent *process types*. The grammar of environments, process types and kinds are given below.

$$\Delta ::= \emptyset \mid \Delta, c:T \quad \Gamma ::= \emptyset \mid \Gamma, P \mid \Gamma, u : S \mid \Gamma, i : I \mid \Gamma, X : \tau \quad \tau ::= \Delta \mid \Pi i : I. \tau$$

Δ is the *session environment* which associates channels to session types. Γ is the *standard environment* which contains predicates and which associates variables to sort types, service names to global types, indices to index sets and process variables to session types. τ is a *process type* which is either a session environment or a dependent type. We write $\Gamma, u : S$ only if $u \notin \text{dom}(\Gamma)$. We use the same convention for others.

Following [22], we assume given in the typing rules two semantically defined judgements: $\Gamma \models P$ (predicate P is a consequence of Γ) and $\Gamma \models i : I$ ($i : I$ follows from the assumptions of Γ). We also inductively define well-formed types using a kind system [11]. The judgement $\Gamma \vdash U \blacktriangleright \kappa$ means type U has kind κ . Kinds include proper types for global, value, principal, end-point and process types (denoted by **Type**), and the kind of type families, written by $\Pi i : I. \kappa$. Well-formedness of a term i and P in Γ and environments is defined in the standard way [1].

3.3 Typing processes

We explain here (Figure 8) a selection of the process typing rules. Rules [TNAT] and [TVAR] are standard ($\Gamma \vdash \text{Env}$ means that Γ is well-formed). For participants, we check their typing by [TID] and [TP] in a similar way as [22] where $\Gamma \vdash \kappa$ means kinding κ is well-formed. In [TPREC], we use the abbreviation $[0..j] = \{i : \text{nat} \mid i \leq j\}$. Then we define I^- by $[0..0]^- = \emptyset$ and $[0..i]^- = [0..i-1]$. This rule deals with the changed index range within the recursor body. More precisely, we first check τ 's kind. Then we verify for the base case ($j = 0$) that P has type $\tau\{0/j\}$. Last, we check the more complex inductive case: Q should have type $\tau\{i+1/j\}$ under the environment $\Gamma, i : I^-, X : \tau\{i/j\}$ where $\tau\{i/j\}$ of X means that X satisfies the predecessor's type (induction hypothesis). [TAPP] is the elimination rule for dependent types.

Since our types include dependent types and recursors, we need a notion of type equivalence. We extend the standard method of [1, §2] with the recursor: [WF] is the main rule defining $G_1 \equiv G_2$ and relies on the existence of a common weak head normal form for the two types and [PREC] says two recursors are equated if either (1) each subgraph is equated by \equiv , or (2) they reduce to the same normal forms when applied to a finite number of indices. Rule [TEQ] allows to type processes up-to type equivalence.

[TINIT] types a session initialisation on shared channel u , binding channel y and requiring participants $\{p_0, \dots, p_n\}$. The premise verifies that the type of y is the first projection of the global type G of u and that the participants in G (denoted by $\text{pid}(G)$) can be semantically derived as $\{p_0, \dots, p_n\}$. [TACC] allows to type the p -th participant to the session initiated on u . The typing rule checks that the type of y is the p -th projection of the global type G of u and that G is fully instantiated. The kind rule ensures that G is fully instantiated (i.e. G 's kind is **Type**). [TREQ] types the process that waits for an accept from a participant: its type corresponds to the end-point projection of G .

$$\begin{array}{c}
\frac{\Gamma \vdash \text{Env}}{\Gamma \vdash n \triangleright \text{nat}} \text{[TNAT]} \quad \frac{\Gamma \vdash \kappa}{\Gamma \vdash \text{Alice} \triangleright \kappa} \text{[TID]} \quad \frac{\Gamma \vdash p \triangleright \Pi i:I.\kappa \quad \Gamma \models i:I}{\Gamma \vdash p[i] \triangleright \kappa\{i/i\}} \text{[TP]} \\
\\
\frac{\Gamma, i : I^-, X : \tau\{i/j\} \vdash Q \triangleright \tau\{i+1/j\} \quad \Gamma \vdash P \triangleright \tau\{0/j\} \quad \Gamma, j:I \vdash \tau \blacktriangleright \kappa}{\Gamma \vdash \mathbf{R} P \lambda i.\lambda X.Q \triangleright \Pi j:I.\tau} \text{[TPREC]} \\
\\
\frac{\left\{ \begin{array}{l} \Gamma \vdash G_1 \equiv G_2 \quad \Gamma \vdash G'_1 \equiv G'_2 \quad \text{or} \\ \Gamma \vdash \mathbf{R} G_1 \lambda i:I.\lambda x.G'_1 n \equiv \mathbf{R} G_2 \lambda i:I.\lambda x.G'_2 n \text{ with } \Gamma \models I = [0..m], 0 \leq n \leq m \end{array} \right.}{\Gamma \vdash \mathbf{R} G_1 \lambda i:I.\lambda x.G'_1 \equiv_{\text{wf}} \mathbf{R} G_2 \lambda i:I.\lambda x.G'_2} \text{[PREC]} \\
\\
\frac{\Gamma \vdash \text{whnf}(G_1) \equiv_{\text{wf}} \text{whnf}(G_2)}{\Gamma \vdash G_1 \equiv G_2} \text{[WF]} \quad \frac{\Gamma \vdash P \triangleright \tau \quad \Gamma \vdash \tau \equiv \tau'}{\Gamma \vdash P \triangleright \tau'} \text{[TEQ]} \quad \frac{\Gamma \vdash P \triangleright \tau \quad \Gamma \vdash \tau \leq \tau'}{\Gamma \vdash P \triangleright \tau'} \text{[TSUB]} \\
\\
\frac{\Gamma, X : \tau \vdash P \triangleright \tau}{\Gamma \vdash \mu X.P \triangleright \tau} \text{[TREC]} \quad \frac{\Gamma, X : \tau \vdash \text{Env}}{\Gamma, X : \tau \vdash X \triangleright \tau} \text{[TVAR]} \quad \frac{\Gamma \vdash P \triangleright \Pi i:I.\tau \quad \Gamma \models i \in I}{\Gamma \vdash P i \triangleright \tau\{i/i\}} \text{[TAPP]} \\
\\
\frac{\Gamma \vdash u : \langle G \rangle \quad \Gamma \vdash P \triangleright \Delta, y : G \upharpoonright p_0 \quad \Gamma \vdash u : \langle G \rangle \quad \Gamma \vdash P \triangleright \Delta, y : G \upharpoonright p}{\Gamma \vdash p_i \triangleright \text{nat} \quad \Gamma \models \text{pid}(G) = \{p_0..p_n\}} \text{[TINIT]} \quad \frac{\Gamma \vdash p \triangleright \text{nat} \quad \Gamma \models p \in \text{pid}(G)}{\Gamma \vdash u[p](y).P \triangleright \Delta} \text{[TACC]} \\
\\
\frac{\Gamma \vdash a : \langle G \rangle \quad \Gamma \vdash p \triangleright \text{nat} \quad \Gamma \models p \in \text{pid}(G)}{\Gamma \vdash \bar{a}[p] : s \triangleright s[p] : G \upharpoonright p} \text{[TREQ]} \quad \frac{\Gamma \vdash e \triangleright S \quad \Gamma \vdash P \triangleright \Delta, c : T}{\Gamma \vdash c!\langle p, e \rangle; P \triangleright \Delta, c : !\langle p, S \rangle; T} \text{[TOUT]}
\end{array}$$

Fig. 8. Process typing

Recursion [TREC], variable ([TVAR]), output ([TOUT]), input, delegation, inaction, branching/selection and the expression typing rules as well as the typing rules for queues are similar to those in [3, 13].

3.4 Properties of typing

Ensuring termination of type-checking with dependent types is not an easy task since type equivalences are often derived from term equivalences. We rely on the strong normalisation of System \mathcal{T} [12] for the termination proof.

Proposition 3.1 (Termination and Confluence) *The head relation \longrightarrow on global and end-point types (i.e. $G \longrightarrow G'$ and $T \longrightarrow T'$ for closed types in Figure 2) are strong normalising and confluent on well-formed kindings.*

The following lemma is proved by defining the weight of the equality and showing the weight of any premise of a rule is always less than the weight of the conclusion (the weight for a recursor needs to be extended to allow the inductive equality rule).

Proposition 3.2 (Termination for Type-Equality Checking) *Assuming that proving the predicates $\Gamma \models P$ appearing in type equality derivations is decidable, then type-equality checking of $\Gamma \vdash G \equiv G'$ terminates. Similarly for other types.*

Proposition 3.3 (Termination for Type-Checking) *Assuming that proving the predicates $\Gamma \models P$ appearing in kinding, equality, projection and typing derivations is decidable, then type-checking of annotated process P , i.e. $\Gamma \vdash P \triangleright \emptyset$ terminates.*

Proof. (Outline) By the standard argument from indexed dependent types [1, 22], for the dependent λ -applications, we do not require equality of terms (i.e. we only need the equality of the indices by the semantic consequence judgements). Hence to eliminate the type equality rule [TEQ], we include the type equality check into [TINIT, TREQ, TACC] (between the global type and its projected session type), and the input rule (between the session type and the type annotating x). Similarly for recursive agents. Since $\alpha \equiv \beta$ (for any type α and β) terminates, these checks always terminate. \square

To ensure the termination of $\Gamma \models P$, several solutions include the restriction of predicates to linear equalities over natural numbers without multiplication (or to other decidable arithmetic subsets) or the restriction of indices to finite domains, cf. [22].

3.5 Type soundness and progress

The following lemma states that mergeability is sound with respect to the branching subtyping \leq . By this, we can safely replace the third clause $\sqcup_{k \in K} G_k \uparrow \mathbf{q}$ of the branching case from the projection definition by $\sqcap\{T \mid \forall k \in K. T \leq (G_k \uparrow \mathbf{q})\}$. This allows us to prove subject reduction by including subsumption as done in [13].

Lemma 3.4 (Soundness of mergeability) *Suppose $G_1 \uparrow \mathbf{p} \bowtie G_2 \uparrow \mathbf{p}$ and $\Gamma \vdash G_i$. Then there exists G such that $G \uparrow \mathbf{p} = \sqcap\{T \mid T \leq G_i \uparrow \mathbf{p} \ (i = 1, 2)\}$ where \sqcap denotes the maximum element with respect to \leq .*

As session environments record channel states, they evolve when communications proceed. This can be formalised by introducing a notion of session environments reduction. These rules are formalised below modulo \equiv .

- $\{s[\hat{\mathbf{p}}] : !(\hat{\mathbf{q}}, U); T, s[\hat{\mathbf{q}}] : ?(\hat{\mathbf{p}}, U); T'\} \Rightarrow \{s[\hat{\mathbf{p}}] : T, s[\hat{\mathbf{q}}] : T'\}$
- $\{s[\hat{\mathbf{p}}] : \oplus(\hat{\mathbf{q}}, \{l_k : T_k\}_{k \in K})\} \Rightarrow \{s[\hat{\mathbf{p}}] : \oplus(\hat{\mathbf{q}}, l_j); T_j\}$
- $\{s[\hat{\mathbf{p}}] : \oplus(\hat{\mathbf{q}}, l_j); T, s[\hat{\mathbf{q}}] : \&\mathcal{L}(\mathbf{p}, \{l_k : T_k\}_{k \in K})\} \Rightarrow \{s[\hat{\mathbf{p}}] : T, s[\hat{\mathbf{q}}] : T_j\}$
- $\Delta \cup \Delta'' \Rightarrow \Delta' \cup \Delta''$ if $\Delta \Rightarrow \Delta'$.

The first rule corresponds to the reception of a value or channel by the participant $\hat{\mathbf{q}}$; the second rule treats the case of the choice of label l_j while the third rule propagate these choices to the receiver (participant $\hat{\mathbf{q}}$). Using the above notion we can state type preservation under reductions as follows:

Theorem 3.5 (Subject Congruence and Reduction) *If $\Gamma \vdash P \triangleright \tau$ and $P \longrightarrow^* P'$, then $\Gamma \vdash P' \triangleright \tau'$ for some τ' such that $\tau \Rightarrow^* \tau'$.*

Note that communication safety [13, Theorem 5.5] and session fidelity [13, Corollary 5.6] are corollaries of the above theorem. A notable fact is, in the presence of the asynchronous join primitive, we can still obtain *progress* in a single multiparty session as in [13, Theorem 5.12], i.e. if a program P starts from one session, the reductions at session channels do not get a stuck. Formally we write $\Gamma \vdash^* P \triangleright \Delta$ if P is typable and

with a type derivation where the session typing in the premise and the conclusion of each prefix rule is restricted to at most a singleton. Another element which can hinder progress is when interactions at shared channels cannot proceed. We say P is *well-linked* when for each $P \longrightarrow^* Q$, whenever Q has an active prefix whose subject is a (free or bound) shared channels, then it is always reducible. The proof is similar to [13, Theorem 5.12].⁴

Theorem 3.6 (Progress) *If P is well-linked and without any element from the runtime syntax and $\Gamma \vdash^* P \triangleright \emptyset$. Then for all $P \longrightarrow^* Q$, either $Q \equiv \mathbf{0}$ or $Q \longrightarrow R$ for some R .*

3.6 Typing examples

Repetition example - § 1 (1) This example illustrates the repetition of a message pattern. Let $G(n) = \text{foreach}(i < n)\{\text{Alice} \rightarrow \text{Bob} : \langle \text{nat} \rangle. \text{Bob} \rightarrow \text{Carol} : \langle \text{nat} \rangle\}$. Following the projection from figure 7, Alice’s end-point projection of $G(n)$ is⁵:

$$G(n) \upharpoonright \text{Alice} = \mathbf{R} \text{ end } \lambda i. \lambda x. !\langle \text{Bob}, \text{nat} \rangle; x \ n$$

Let $\text{Alice}(n) = \bar{a}[\text{Alice}, \text{Bob}, \text{Carol}](y).(\mathbf{R} \ \mathbf{0} \ \lambda i. \lambda X. y! \langle \text{Bob}, e[i] \rangle; X \ n)$ and let $\Delta(n) = \{y : (G(n) \upharpoonright \text{Alice})\}$ and $\Gamma = n : \text{nat}, a : \langle G \rangle$. We can prove that $\Gamma \vdash \text{Alice}(n) \triangleright \emptyset$ from [TINIT] if we have $\Gamma \vdash \mathbf{R} \ \mathbf{0} \ \lambda i. \lambda X. y! \langle \text{Bob}, e[i] \rangle; X \ n \triangleright \Delta(n)$. This, in turn, is given by [TPREC] and [TAPP] from $\Gamma, i : I^-, X : \Delta(i) \vdash y! \langle \text{Bob}, e[i] \rangle; X \triangleright \Delta(i+1)$ and the trivial $\Gamma \vdash \mathbf{0} \triangleright y : \text{end}$. From [TVAR], we have $\Gamma, i : I^-, X : \Delta(i) \vdash X \triangleright \Delta(i)$. We conclude by [TOUT] and weak head normal form equivalence [WF] of the types $\Delta(i+1)$ and $y : !\langle \text{Bob}, \text{nat} \rangle; (\mathbf{R} \ \text{end } \lambda j. \lambda x. !\langle \text{Bob}, \text{nat} \rangle; x \ i)$. $\text{Bob}(n)$ and $\text{Carol}(n)$ can be similarly typed.

Sequence example - § 1 (2) The sequence example consists of three roles (when $n \geq 2$): the starter $W[n]$ sends the first message, the final worker $W[0]$ receives the final message and the middle workers first receive a message and then send another to the next worker. We write below the generic projection for participant $W[p]$ (left) and the end-point type that naturally types the processes (right):

$$\begin{array}{l} \mathbf{R} \ \text{end } \lambda i. \lambda x. \\ \text{if } p = W[i+1] \text{ then } !\langle W[i], \text{nat} \rangle; x \\ \text{else if } p = W[i] \text{ then } ?\langle W[i+1], \text{nat} \rangle; x \\ \text{else } x \end{array} \quad \left| \quad \begin{array}{l} \text{if } p = W[n] \text{ then } !\langle W[n-1], \text{nat} \rangle; \text{else} \\ \text{if } p = W[0] \text{ then } ?\langle W[1], \text{nat} \rangle; \text{else} \\ \text{if } p = W[i] \text{ then } ?\langle W[i+1], \text{nat} \rangle; !\langle W[i-1], \text{nat} \rangle; \end{array} \right. n$$

In order to type this example, we need to prove the equivalence of these two types. For any instantiation of p and n , the standard weak head normal form equivalence rule [WF] is sufficient. Proving the equivalence for all p and n requires either (a) to bound the domain I in which they live, and check all instantiations within this finite domain; or (b) to prove the equivalence through a meta-logic case analysis. In case (a), type checking terminates, while case (b) allows to easily prove strong properties about a protocol’s implementation.

⁴ We believe a stronger progress property for interleaved multiparty sessions ensured by the interaction typing in [3] can be obtained in this framework, too (since our typing system is an extension from the communication system in [3]).

⁵ For readability, in the following examples, we omit from the nested conditionals the impossible cases.

FFT example - Figure 5 We prove type-safety and deadlock-freedom for the FFT processes. Let P_{fft} be the following process:

$$\begin{aligned} & \Pi n. (\nu a)(\mathbf{R} \bar{a}[\mathbf{p}_0.. \mathbf{p}_{2^n-1}](y). P(2^n - 1, \mathbf{p}_0, x_{\bar{\mathbf{p}}_0}, y, r_{\mathbf{p}_0}) \\ & \lambda i. \lambda Y. (\bar{a}[\mathbf{p}_{i+1}](y). P(i + 1, \mathbf{p}_{i+1}, x_{\bar{\mathbf{p}}_{i+1}}, y, r_{\mathbf{p}_{i+1}}) \mid Y) 2^n - 1) \end{aligned}$$


As we reasoned above, each $P(n, \mathbf{p}, x_{\bar{\mathbf{p}}}, y, r_{\mathbf{p}})$ is straightforwardly typable by an end-point type which is equivalent with the one projected from the global type G from figure 5(c). Automatically checking the equivalence for all n is not easy though: we need to rely on the finite domain restriction using [WF,PREC]. The following theorem says once P_{fft} is applied to a natural number m , its evaluation always terminates with the answer at $r_{\mathbf{p}}$. The proof is by the progress (Theorem 3.5), noting $P_{\text{fft}} m$ is typable by a single, multiparty dependent session (except the answering channel at $r_{\mathbf{p}}$).

Theorem 3.7 (Type safety and deadlock-freedom of FFT) *For all $m, \emptyset \vdash P_{\text{fft}} m \triangleright \emptyset$; and if $P_{\text{fft}} m \longrightarrow^* Q$, then $Q \longrightarrow^* (r_0!(0, X_0) \mid \dots \mid r_{2^m-1}!(0, X_{2^m-1}))$ where the $r_p!(0, X_p)$ are the actions sending the final values X_p on external channels r_p .*

Web Service example - Figure 9 We program and type a real-world Web service use case: Quote Request (C-U-002) is the most complex scenario described in [19], the public document authored by the W3C Choreography Description Language Working Group [21]. As described in Figure 9, a buyer interacts with multiple suppliers who in turn interact with multiple manufacturers in order to obtain quotes for some goods or services. The Requirements from Section 3.1.2.2 of [19] include the ability to *reference a global description from within a global description* to support *recursive behaviour* as denoted in STEP 4(b, d): this can be achieved by parameterised multiparty session types.

We write the specification of the usecase program modularly, starting from the first steps of the informal description above. Here, Buyer stands for the buyer, $\text{Supp}[i]$ for a supplier, and $\text{Manu}[j]$ for a manufacturer. We alias manufacturers by $\text{Manu}[i][j]$ to express the fact that $\text{Manu}[j]$ is connected to $\text{Supp}[i]$ (a single $\text{Manu}[j]$ can have multiple aliases $\text{Manu}[i'][j]$, see figure 9). Then, we can write global types for each of the steps. STEP 1 is a simple *multicast* (type G_1). For STEP 2, we write first $G_2(i)$, the nested interaction loop between the i -th supplier and its manufacturers (J_i gives all $\text{Manu}[j]$ connected to $\text{Supp}[i]$). Then G_2 can describe the subsequent action within the main loop. For STEP 3, for simplicity we assume the preference is given by the (reverse) ordering of I . The first choice of G_3 corresponds to the two cases of STEP 3. In the innermost branch of G_3 , the branches ok, retryStep3 and reject correspond to STEP 4(a), (b) and (c) respectively, while the type variable \mathbf{t} models STEP 4(d). We can now compose these subprotocols together. The full global type is then $G = \Pi i. \Pi \bar{J}. (G_1 ; \mu \mathbf{t}. (G_2 ; G_3))$ where we have i suppliers, and \bar{J} gives the index sets J_i of the $\text{Manu}[j]$ s connected with each $\text{Supp}[i]$.

For the end-point projection, we focus on the suppliers' case. The projections of G_1 and G_2 are straightforward. For $G_3 \upharpoonright \text{Supp}[n]$, we use the branching injection and mergeability theory developed in § 3.1. After the relevant application of [TEQ], we can obtain the projection written in Figure 9. To tell the other suppliers whether the



1. A buyer requests a quote from a set of suppliers.
 $G_1 = \text{foreach}(i < n)\{\text{Buyer} \rightarrow \text{Supp}[i]: \langle \text{Quote} \rangle\}$
2. All suppliers receive the request to ask their respective manufacturers for a bill of material items. The suppliers interact with their manufacturers to build their quotes for the buyer.
 $G_2(i) = \text{foreach}(j : J_i)\{\text{Supp}[i] \rightarrow \text{Manu}[i][j]: \langle \text{Item} \rangle.$
 $\text{Manu}[i][j] \rightarrow \text{Supp}[i]: \langle \text{Quote} \rangle\}$
 The eventual quote is sent back to the buyer.
 $G_2 = \text{foreach}(i : I)\{G_2(i); \text{Supp}[i] \rightarrow \text{Buyer}: \langle \text{Quote} \rangle\}$

3. EITHER

- (a) The buyer agrees with one or more of the quotes and places the order(s). OR
- (b) The buyer responds to one or more of the quotes by modifying and sending them back to the relevant suppliers.

4. EITHER

- (a) The suppliers respond to a modified quote by agreeing to it and sending a confirmation message back to the buyer. OR
- (b) The supplier responds by modifying the quote and sending it back to the buyer and the buyer goes back to STEP 3. OR
- (c) The supplier responds to the buyer rejecting the modified quote. OR
- (d) The quotes from the manufacturers need to be renegotiated by the supplier. Go to STEP 2.

$$\begin{array}{l}
 G_3 = \mathbf{R} \ \mathbf{t} \ \lambda i. \lambda y. \text{Buyer} \rightarrow \text{Supp}[i] : \{ \\
 \text{ok} : \quad \text{end} \\
 \text{modify} : \text{Buyer} \rightarrow \text{Supp}[i]: \langle \text{Quote} \rangle. \\
 \quad \text{Supp}[i] \rightarrow \text{Buyer} : \{\text{ok} : \quad \text{end} \\
 \quad \quad \text{retryStep3} : y \\
 \quad \quad \text{reject} : \quad \text{end}\} \} i \\
 \end{array}
 \quad \left| \quad
 \begin{array}{l}
 G_3 \upharpoonright \text{Supp}[n] = \&\langle \text{Buyer}, \{ \\
 \text{ok} : \quad \text{end} \\
 \text{modify} : ?\langle \text{Buyer}, \text{Quote} \rangle; \oplus \langle \text{Buyer}, \{ \\
 \quad \text{ok} : \quad \text{end} \\
 \quad \text{retryStep3} : y \\
 \quad \text{reject} : \quad \text{end}\} \rangle \} \}
 \end{array}$$

Fig. 9. The Quote Request use case (C-U-002) [19] with the corresponding global types

loop is being reiterated or if it is finished, we can simply insert the following closing notification $\text{foreach}(j \in I \setminus i)\{\text{Buyer} \rightarrow \text{Supp}[j] : \{\text{close} : \}\}$ before each end, and a similar retry notification (with label `retryStep3`) before `x`. Finally, each end-point type is formed by $(G_1 \upharpoonright \text{Supp}[n]; \mu x. G_2 \upharpoonright \text{Supp}[n]; G_3 \upharpoonright \text{Supp}[n])$. While the global types look sequential, actual typed processes can asynchronously join a session and be executed in parallel (e.g., at STEP 1-2, no synchronisation is needed between $\text{Supp}[i]$).

We have explored the impact of parameterised type structures for communications through implementations of the above use case and of several parallel algorithms in Java with session types [14], including the Jacobi method (with sequence and mesh topologies) and the FFT. We observe (1) a clear coordination of the communication behaviour of each party with the construction of the whole multiparty protocol, thus reducing the programming errors and ensuring deadlock-freedom; and (2) a performance benefit against the original binary session version, reducing the overhead of multiple binary session establishments (see [11]). Full implementation and integration of our theory into [4, 14] is on-going work.

4 Related work

Dependent types The first use of primitive recursive functionals for dependent types is in Nelson’s \mathcal{T}^π [18] for the λ -calculus, which is a finite representation of \mathcal{T}^∞ by Tait and Martin L of [16, 20]. \mathcal{T}^π can type functions previously untypable in ML, and the finite representability of dependent types makes it possible to have a type-reconstruction algorithm. We also use the ideas from the DML’s dependent typing system in [1, 22] where type dependency is only allowed for index sorts, so that type-checking can be reduced to a constraint-solving problem over indices. Our design choice to combine both systems gives (1) the simplest formulation of sequences of global and end-point types and processes described by the primitive recursor; (2) a precise specification for parameters appearing in the participants based on index sorts; and (3) a clear integration with the full session types and general recursion, whilst ensuring decidability of type-checking (if the constraint-solving problem is decidable). From the basis of these works, our type equivalence does not have to rely on behavioural equivalence between processes, but only on the strongly normalising *types* represented by recursors. None of these works investigate families of global specifications using dependent types.

Types and contracts for multiparty interactions The work [7] presented an *executable global processes* for Web interactions based on the binary session types. Our work provides flexible, programmable global descriptions as *types*, offering a progress for parameterised multiparty session, which is not ensured in [7]. Recent formalisms for typing multiparty interactions include [6, 9]. These works treat different aspects of dynamic session structures. *Contracts* [9] can type more processes than session types, thanks to the flexibility of process syntax for describing protocols. However, typable processes themselves in [9] may not always satisfy the properties of session types such as progress: it is proved later by checking whether the type meets a certain form. Hence proving progress with contracts effectively requires an exploration of all possible paths (interleaving, choices) of a protocol. The most complex example of [9, § 3] (a group key agreement protocol from [2]), which is typed as π -processes with delegations, can be specified and articulated by a single parameterised global session type as:

$$\begin{aligned} \Pi n : I. (\text{foreach}(i < n) \{ \mathbb{W}[n - i] \rightarrow \mathbb{W}[n - i + 1] : \langle \text{nat} \rangle \}; \\ \text{foreach}(i < n) \{ \mathbb{W}[n - i] \rightarrow \mathbb{W}[n + 1] : \langle \text{nat} \rangle. \mathbb{W}[n + 1] \rightarrow \mathbb{W}[n - i] : \langle \text{nat} \rangle \}) \end{aligned}$$

Once the end-point process conforms to this specification, we can automatically guarantee communication safety and progress.

Conversation Calculus [6] supports the dynamic joining and leaving of participants. Though the formalism in § 2.2 can operationally capture such dynamic features, the aim of the present work is *not* the type-abstraction of dynamic interaction patterns. Our purpose is to capture, in a single type description, a family of protocols over arbitrary number of participants, to be instantiated at runtime. The parameterisation gives freedom not possible with previous session types: once typed, a parametric process is ensured that its arbitrary well-typed instantiations, in terms of both topologies and process behaviours, satisfy the safety and progress properties of typed processes. Parameterisation, composition and repetition are common idioms in parallel algorithms and choreographic/conversational interactions, all of which are uniformly treatable in our dependent type theory. Here types offer a rigorous structuring principle which can

economically abstract rich interaction structures, including parameterised ones.

Acknowledgements. We thank the reviewers, Lasse Nielsen and Dimitris Mostrous for their useful comments for the paper and Kohei Honda for discussions.

References

1. D. Aspinall and M. Hofmann. *Advanced Topics in Types and Programming Languages*, chapter Dependent Types. MIT, 2005.
2. G. Ateniese, M. Steiner, and G. Tsudik. Authenticated group key agreement and friends. In *CCS '98: Proceedings of the 5th ACM conference on Computer and communications security*, pages 17–26, New York, NY, USA, 1998. ACM.
3. L. Bettini et al. Global progress in dynamically interfered multiparty sessions. In *CONCUR'08*, volume 5201 of *LNCS*, pages 418–433, 2008.
4. K. Bhargavan, R. Corin, P.-M. Deniérou, C. Fournet, and J. Leifer. Cryptographic protocol synthesis and verification for multiparty sessions. In *CSF*, pages 124–140, 2009.
5. E. Bonelli and A. Compagnoni. Multipoint session types for a distributed calculus. In *TGC'07*, volume 4912 of *LNCS*, pages 240–256, 2008.
6. L. Caires and H. T. Vieira. Conversation types. In *ESOP*, volume 5502 of *LNCS*, pages 285–300. Springer, 2009.
7. M. Carbone, K. Honda, and N. Yoshida. Structured communication-centred programming for web services. In *ESOP'07*, volume 4421 of *LNCS*, pages 2–17, 2007.
8. M. Carbone, N. Yoshida, and K. Honda. Asynchronous session types: Exceptions and multiparty interactions. In *SFM'09*, volume 5569 of *LNCS*, pages 187–212. Springer, 2009.
9. G. Castagna and L. Padovani. Contracts for mobile processes. In *CONCUR 2009*, number 5710 in *LNCS*, pages 211–228, 2009.
10. J. W. Cooley and J. W. Tukey. An algorithm for the machine calculation of complex fourier series. *Mathematics of Computation*, 19(90):297–301, 1965.
11. Online Appendix. <http://www.doc.ic.ac.uk/~yoshida/dependent/>.
12. J.-Y. Girard, Y. Lafont, and P. Taylor. *Proofs and Types*, volume 7 of *Cambridge Tracts in Theoretical Computer Science*. CUP, 1989.
13. K. Honda, N. Yoshida, and M. Carbone. Multiparty Asynchronous Session Types. In *POPL'08*, pages 273–284. ACM, 2008.
14. R. Hu, N. Yoshida, and K. Honda. Session-Based Distributed Programming in Java. In *ECOOP'08*, volume 5142 of *LNCS*, pages 516–541, 2008.
15. F. T. Leighton. *Introduction to parallel algorithms and architectures: arrays, trees, hypercubes*. Morgan Kaufmann, 1991.
16. P. Martin-Löf. Infinite terms and a system of natural deduction. In *Compositio Mathematica*, pages 93–103. Wolters-Noordhoff, 1972.
17. D. Mostrous, N. Yoshida, and K. Honda. Global principal typing in partially commutative asynchronous sessions. In *ESOP'09*, volume 5502 of *LNCS*, pages 316–332. Springer, 2009.
18. N. Nelson. Primitive recursive functionals with dependent types. In *MFPS*, volume 598 of *LNCS*, pages 125–143, 1991.
19. Web Services Choreography Requirements (No. 11). <http://www.w3.org/TR/ws-chor-reqs>.
20. W. W. Tait. Infinitely long terms of transfinite type. In *Formal Systems and Recursive Functions*, pages 177–185. North Holland, 1965.
21. Web Services Choreography Working Group. Choreography Description Language. <http://www.w3.org/2002/ws/chor/>.
22. H. Xi and F. Pfenning. Dependent types in practical programming. In *POPL*, pages 214–227, 1999.