

Compiling First-Order Functions to Session-Typed Parallel Code

David Castro-Perez
Imperial College London
London, UK
d.castro-perez@imperial.ac.uk

Nobuko Yoshida
Imperial College London
London, UK
n.yoshida@imperial.ac.uk

Abstract

Building correct and efficient message-passing parallel programs still poses many challenges. The incorrect use of message-passing constructs can introduce deadlocks, and a bad task decomposition will not achieve good speedups. Current approaches focus either on correctness or efficiency, but limited work has been done on ensuring both. In this paper, we propose a new parallel programming framework, PAI_g, which is a first-order language with *participant annotations* that ensures *deadlock-freedom by construction*. PAI_g programs are coupled with an abstraction of their communication structure, a *global type* from the theory of *multiparty session types* (MPST). This global type serves as an output for the programmer to assess the efficiency of their achieved parallelisation. PAI_g is implemented as an EDSL in Haskell, from which we: 1. compile to low-level message-passing C code; 2. compile to sequential C code, or interpret as sequential Haskell functions; and, 3. infer the communication protocol followed by the compiled message-passing program. We use the properties of global types to perform message reordering optimisations to the compiled C code. We prove the *extensional equivalence* of the compiled code, as well as *protocol compliance*. We achieve linear speedups on a shared-memory 12-core machine, and a speedup of 16 on a 2-node, 24-core NUMA.

Keywords multiparty session types, parallelism, arrows

1 Introduction

Structured parallel programming is a technique for parallel programming that requires the use of high-level parallel constructs, rather than low-level send/receive operations [52; 62]. A popular approach to structured parallelism is the use of *algorithmic skeletons* [20; 36], i.e. higher-order functions that implement common patterns of parallelism. Programming in terms of high-level constructs rather than low-level send/receive operations is a successful way to avoid common concurrency bugs *by construction* [38]. One limitation of structured parallelism is that it restricts programmers to use a set of fixed, predefined parallel constructs. This is

problematic if a function does not match one of the available parallel constructs, or if a program needs to be ported to an architecture where some of the skeletons have not been implemented. Unlike previous structured parallelism approaches, we do not require the existence of an underlying library or implementation of common patterns of parallelism.

In this paper, we propose a structured parallel programming framework whose front-end language is a first-order language based on the *algebra of programming* [2; 3]. The algebra of programming is a mathematical framework that codifies the basic laws of algorithmics, and it has been successfully applied to e.g. *program calculation* techniques [4], *datatype-generic programming* [35], and *parallel computing* [66]. Our framework produces message-passing parallel code from program specifications written in the front-end language. The programmer controls how the program is parallelised by *annotating* the code with *participant identifiers*. To make sure that the achieved parallelisation is satisfactory, we produce as an output a formal description of the *communication protocol* achieved by a particular parallelisation. This formal description is a *global type*, introduced by Honda et al. [42] in the theory of *Multiparty Session Types* (MPST). We prove that the parallelisation, and any optimisation performed to the low-level code respects the inferred protocol. The properties of global types justify the message reordering done by our back-end. In particular, we permute send and receive operations whenever sending does not depend on the values received. This is called *asynchronous optimisation* [57], and removes unnecessary synchronisation, while remaining communication-safe.

1.1 Overview




Figure 1. Overview


Our framework has three layers: (1) Parallel Algebraic Language (PAI_g), a point-free first-order language with *participant annotations*, which describe which process is in charge of executing which part of the computation; (2) Message

111 Passing Monad (Mp), a monadic language that represents
 112 low-level message-passing parallel code, from which we gen-
 113 erate parallel C code; and (3) *global types* (from MPST), a
 114 formal description of the protocol followed by the output
 115 Mp code. Fig. 1 shows how these layers interact. PAlg, high-
 116 lighted in green, is the input to our framework; and Mp and
 117 global types (MPST), highlighted in yellow, are the outputs.
 118 We prove that the generated code behaves as prescribed by
 119 the global type, and any low-level optimisation performed on
 120 the generated code must respect the protocol. As an example,
 121 we show below a parallel mergesort. mergesort.

```

122 1 msort :: (CVal a, CAlg f) => Int -> f [a] [a]
123 2 msort n = fix n $ \ms x -> vlet (vsize x) $ \sz ->
124 3   if sz <= 1 then x
125 4   else vlet (sz / 2) $ \sz2 ->
126 5     vlet (par ms $ vtake sz2 x) $ \x1 ->
127 6     vlet (par ms $ vdrop sz2 x) $ \xr ->
128 7     app merge $ pair (sz, pair (x1, xr))
  
```

130 The return type of `msort`, $f [a] [a]$, is the type of first-order
 131 programs that take lists of *values* $[a]$, and return $[a]$. Con-
 132 straint CAlg restricts the kind of operations that are allowed
 133 in the function definition. The integer parameter to function
 134 `fix` is used for rewriting the input programs, limiting the
 135 depth of *recursion unrolling*. `par` is used to *annotate* the func-
 136 tions that we want to run at different processes, and function
 137 `app` is used to run functions at the same participant as their
 138 inputs. In case this input comes from different participants,
 139 first all values are gathered at any of them, and then the
 140 function is applied. We can instantiate f either as a sequen-
 141 tial program, as a parallel program, or as an MPST protocol.
 142 We prove that the sequential program, and output parallel
 143 programs are *extensionally equal*, and that the output parallel
 144 program *complies* with the inferred protocol. For example,
 145 interpreting `msort 1` as a parallel program produces C code
 146 that is extensionally equal to its sequential interpretation,
 147 and behaves as the following protocol:



152 This is a depth 1 divide-and-conquer, where p_1 divides the
 153 task, sends the sub-tasks to p_2 and p_3 , and combines the
 154 results. If the input is small, p_1 produces the result directly.

155 Our prototype implementation is a *tagless-final* encoding
 156 [9] in Haskell of a point-free language. Constraint CAlg is
 157 a first-order form of *arrows* [45; 61], with a syntactic sugar
 158 layer that allows us to write code closer to (point-wise) id-
 159 iomatic Haskell. The remainder of the paper focuses on the
 160 language underlying CAlg .

162 **Why Multiparty Session Types** There are both practical
 163 and theoretical advantages. On the theoretical side, the the-
 164 ory of multiparty session types ensures *deadlock-freedom* and

166 **protocol compliance.** The MPST theory guarantees that the
 167 code that we generate complies with the inferred protocol
 168 (Theorem 5.2), which greatly simplifies the proof of *exten-*
 169 *sional* equivalence (Theorem 5.3), by allowing us to focus on
 170 representative traces, instead of all possible interleavings of
 171 actions. On the practical side, we perform message reorder-
 172 ing optimisation based on the global types [57]. Moreover,
 173 an *explicit* representation of the communication protocol is
 174 a valuable output for programmers, since it can be used to
 175 assess a parallelisation. (Fig. 4).

176 1.2 Outline and Contributions

177 $\S 2$ defines the Algebraic Functional Language (Alg), a lan-
 178 guage inspired by the algebra of programming, that we use
 179 as a basis for our work; $\S 3$ proposes the Parallel Algebraic
 180 Language (PAlg), our front-end language, as an extension
 181 of Alg with participant annotations; $\S 4$ introduces a *proto-*
 182 *col inference relation* that associates PAlg expressions with
 183 MPST protocols, specified as global types. We prove that
 184 the inferred protocols are deadlock-free: i.e. every *send* has
 185 a matching *receive*. Moreover, we use the global types to
 186 justify message reordering optimisations, while preserving
 187 communication safety; $\S 5$ develops a translation scheme
 188 which *generates* message-passing code from PAlg, that we
 189 prove to preserve the extensionality of the input programs;
 190 $\S 6$ demonstrates our approach using a number of examples.
 191 We will provide as an *artifact* our working prototype im-
 192 plementation, and the examples that we used in $\S 6$, with
 193 instructions on how to replicate our experiments.

194 2 Algebraic Functional Language

195 This section describes the Algebraic Functional Language
 196 (Alg) and its combinators. In functional programming lan-
 197 guages, it is common to provide these combinators as ab-
 198 stractions defined in a base language. For example, one such
 199 combinator is the *split* function (Δ), also known as *fanout*, or
 200 ($\&\&\&$), in the *arrow* literature [45] and `Control.Arrow` Haskell
 201 package [61]. Programming in terms of these combinators,
 202 avoiding explicit mention of variables is known as *point-free*
 203 programming. Another approach is to translate code writ-
 204 ten in a *pointed* style, i.e. with explicit use of variables, to a
 205 point-free style [23; 44]. This translation can be fully auto-
 206 mated [23; 29]. In our approach, we define common point-
 207 free combinators as syntactic constructs of Alg, and require
 208 programs to be implemented in this style. Our implementa-
 209 tion provides a layer of syntactic sugar for programmers to
 210 refer to variables explicitly, as shown in `msort` in $\S 1$, but that
 211 builds internally a point-free representation.

212 2.1 Syntax

213 $F_1, F_2 ::= 1 \mid Ka \mid F_1 + F_2 \mid F_1 \times F_2$
 214 $a, b ::= 1 \mid \text{int} \mid \dots \mid a \rightarrow b \mid a + b \mid a \times b \mid Fa \mid \mu F$
 215 $e_1, e_2 ::= f \mid v \mid \text{const } e \mid \text{id} \mid e_1 \circ e_2 \mid \pi_i \mid e_1 \Delta e_2 \mid t_i \mid e_1 \nabla e_2$
 216 $\mid Fe \mid \text{in}_F \mid \text{out}_F \mid \text{rec}_F e_1 e_2$

In our syntax, f_1, f_2, \dots , capture **atomic functions**, which are functions of which we only know their types; v_1, v_2 are **values** of primitive types (e.g. integer and boolean); e_1, e_2, \dots , represent **expressions**; F_1, F_2, \dots , are **functors**; and a, b, \dots , are **types**. The syntax and semantics are standard [34; 53].

Constant, identity functions, and function composition are const , id and \circ respectively. **Products** are represented using the standard pair notation: if $x : a$ and $y : b$, then $(x, y) : a \times b$. The functions on product types are π_i and Δ , and they represent, respectively, the **projections**, and the **split** operation: $(f \Delta g)(x) = (f x, g x)$. **Coproducts** have two constructors, the **injections** ι_i , that build values of type $a + b$. The ∇ combinator is the **case** operation: $(f_1 \nabla f_2)(\iota_i x) = f_i x$. Products and coproducts can be generalised to multiple arguments: $a \times b \times c$ is isomorphic to $a \times (b \times c)$, and to $(a \times b) \times c$. We use $\prod_{i \in [1, n]} a_i$ as notation for the product of more than two types; similarly we use \sum for coproducts. The \prod notation binds tighter than any other construct. Whenever $\forall i, j \in I, a_i = a_j = a$, we use the notation $\prod_n a$ as a synonym for $\prod_{i \in [1, n]} a_i$.

Functors are objects that take types into types, and functions to functions, such that identities and compositions are preserved. In this work, we focus on *polynomial functors* [31], which are defined inductively: I is the identity functor, and takes a type a to itself; $\text{K}b$ is the *constant functor*, and takes any type to b ; $F_1 \times F_2$ is the *product functor*, and takes a type a to $F_1 a \times F_2 a$; $F_1 + F_2$ is the *coproduct functor*, and takes a type to a coproduct type. A term $F e$ behaves as *mapping* term e to the I positions in F . For example, if $F = \text{K}a \times \text{I} \times \text{I}$, then applying $F e$ to (x, y, z) yields $(x, e y, e z)$.

Recursion is captured by combinators in , out , rec , and type μF . We use standard isorecursive types [31; 47; 53], where μF is *isomorphic* to $F \mu F$, and the isomorphism is given by the combinators in_F (**roll**) and out_F (**unroll**). For any polynomial functor F , μF , and strict functions in_F and out_F are guaranteed to exist. In our implementation, in_F is just a constructor (like inj_i). Recursion is $\text{rec}_F e_1 e_2$, and it is known as a **hylomorphism** [53]. A hylomorphism captures a divide-and-conquer algorithm, with a structure described by F , where e_1 is the *conquer* term and e_2 the *divide* term. Using hylomorphisms requires us to work in a semantic interpretation with *algebraic compactness*, i.e. in which carriers of initial F -algebras and terminal F -coalgebras coincide (or are isomorphic). Hylomorphisms and exponentials $\text{ap} : (a \rightarrow b) \times a \rightarrow b$ allow the definition of a general fix-point operator [54]. Working with hylomorphisms implies that our input programs may not terminate. We guarantee that, given a *terminating* input program, we will *not* produce a non-terminating parallelisation (Theorem 5.3).

Example 2.1 (MergeSort in Alg). Assume a type $\text{L}s$ of lists of elements of type a . Functor $T = \text{K}(\text{L}s) + \text{I} \times \text{I}$ captures the recursive structure of $\text{ms} : \text{L}s \rightarrow \text{L}s$. When splitting some $l : \text{L}s$, we may find one of the two cases described by T : an empty or singleton list, $\text{L}s$, or a list of size ≥ 2 , that

can be split in two halves $\text{L}s \times \text{L}s$. Assume that a functions $\text{spl} : \text{L}s \rightarrow T \text{L}s$, and a function $\text{mrg} : T \text{L}s \rightarrow \text{L}s$. We define $\text{ms} = \text{rec}_T \text{mrg spl}$. By the definition of rec :

$$\begin{aligned} \text{ms} &= \text{rec}_T (\text{id} \nabla \text{mrg}) \text{spl} = (\text{id} \nabla \text{mrg}) \circ T (\text{rec}_T \text{mrg spl}) \circ \text{spl} \\ &= (\text{id} \nabla \text{mrg}) \circ (\text{id} + (\text{rec}_T \text{mrg spl}) \times (\text{rec}_T \text{mrg spl})) \circ \text{spl} \\ &= (\text{id} \nabla \text{mrg} \circ (\text{ms} \times \text{ms})) \circ \text{spl} \end{aligned}$$

Function ms first applies spl . Then, if the list was empty or singleton, it returns the input unmodified. Otherwise, ms applies recursively to the first and second halves. Finally, mrg returns a pair of sorted lists.

3 Parallel Algebraic Language

In the previous section we introduced Alg, a point-free functional language. In this section, we extend this language with *participant annotations*. Annotations occur both at the type and expression levels: at the type level, annotations represent *where* the data of the respective type is; at the expression level, it represents *by whom* the computation is performed. This language extension is called PALg.

The implicit *dataflow* of the Alg (or PALg) constructs determines which interactions must take place to evaluate an annotated program. To illustrate this, we use the Cooley-Tukey Fast-Fourier Transform algorithm [21]. The Cooley-Tukey algorithm is based on the observation that an FFT of size n , fft_n can be described as the combination of two FFTs of size $n/2$. We focus its high-level structure:

$$(\text{add} @ \mathbf{p}_1 \Delta \text{sub} @ \mathbf{p}_2) \circ ((\text{fft}_{n/2} @ \mathbf{p}_3 \circ \pi_1) \Delta ((\text{exp} \circ \text{fft}_{n/2}) @ \mathbf{p}_4 \circ \pi_2))$$

Assume that the input is a pair of vectors that contain the *deinterleaved* input, i.e. elements at even positions on the left, and odd positions on the right. We first compute the fft of size $n/2$ to the even and odd elements at \mathbf{p}_3 and \mathbf{p}_4 respectively. Then, the first half of the output is produced by adding the results pairwise (at \mathbf{p}_1), and the second half by subtracting them (at \mathbf{p}_2). In order to evaluate this expression, we need to know where is the input data. This is specified by the programmer as an annotated type, which we call **interface**. Suppose that the interface specifies that the even elements are at \mathbf{p} , and the odd elements at \mathbf{p}' . The interface that represents this scenario is $(\text{vec} \times \text{vec}) @ (\mathbf{p} \times \mathbf{p}')$, i.e. an annotated pair of vectors, with the first component at \mathbf{p} , and the second component at \mathbf{p}' . By keeping track of the locations of the data, we obtain type $(\text{vec} \times \text{vec}) @ (\mathbf{p}_1 \times \mathbf{p}_2)$, which is the *output* (or *codomain*) interface the PALg expression. We also refer to the annotations (e.g $\mathbf{p}_1 \times \mathbf{p}_2$) as interfaces, whenever there is no ambiguity. We write $\text{fft}_n : (\text{vec} \times \text{vec}) @ (\mathbf{p} \times \mathbf{p}') \rightarrow (\text{vec} \times \text{vec}) @ (\mathbf{p}_1 \times \mathbf{p}_2)$ to represent the input and output interfaces of fft_n .

Consider now $e_1 @ \mathbf{p}_1 \nabla e_2 @ \mathbf{p}_2$. The output interface of this expression is either \mathbf{p}_1 or \mathbf{p}_2 , depending on whether the input is the result of applying ι_1 or ι_2 . We represent such interfaces using *unions*: $e_1 @ \mathbf{p}_1 \nabla e_2 @ \mathbf{p}_2 : (a + b) @ \mathbf{p} \rightarrow c @ (\mathbf{p}_1 \cup \mathbf{p}_2)$. Since \mathbf{p} contains a value of a sum type $a + b$, \mathbf{p} is responsible for notifying both \mathbf{p}_1 and \mathbf{p}_2 which branch needs to be taken in the

control flow. Incorrectly notifying the necessary participants will produce incorrect parallelisations that might deadlock. For example, consider the expression $e_0@p_0 \circ (e_1@p_1 \nabla e_2@p_2)$. Assuming that the input at p , p needs to notify p_0 , otherwise p_0 will be stuck. To avoid such cases, and to compute the interfaces of an expression, we define a type system for PALg.

3.1 Syntax of PALg

$$I ::= p \mid \iota_i I \mid I \times I \quad R ::= I \mid R \cup^{\vec{p}} R \quad P ::= R \rightarrow R \\ e ::= e@p \mid [p \oplus \vec{p}] \mid \text{id} \mid e \circ e \mid \pi_i \mid e \Delta e \mid \iota_i \mid e \nabla e$$

The syntax of PALg is that of Alg, extended with participant annotations (red). Note that certain Alg constructs can only occur under annotations ($e@p$), e.g: in, out and rec. This implies that recursive functions need to be annotated at a single participant. To parallelise recursive functions, they need first to be rewritten into a suitable form, and then annotate the resulting expression. At the moment, we support automatic recursion unrolling up to a user-specified depth. We provide an overview of the main syntactic constructs of PALg: **annotations**, **interfaces**, and **annotated functions**.

Annotations are ranged over by R, R', \dots . We define them in two layers, I , or simple annotations that cannot contain choices (\cup), and R . This way, we ensure that choices only occur at the topmost level. Simple annotations are: participant ids p , that identify processes; products of interfaces $I_1 \times I_2$; and tagged interfaces $\iota_i I$, that keep track of the branch of the choice that led to I . A choice $R_1 \cup^{\vec{p}} R_2$ describes a scenario that is the result of a branch in the control flow, where a value can be found at either R_1 or R_2 . Here, $\vec{p} = p_1 \cdots p_n$ are the participants whose behaviour depends on the path in the control flow. Finally, arrows P of the form $R_1 \rightarrow R_2$ represent the input/output annotations of a parallel program.

Interfaces are annotated types. They range over A, B, \dots , and are of the form $a@R$, which means that values of type a are distributed across R . We require annotated types to be *well-formed*, $\text{WF}(a@R)$, which implies that the structure of a matches that of R . We write \mathcal{I} to represent *one-hole contexts* for interfaces, with $\mathcal{I}[p]$ representing the interface that results of placing p at the hole in \mathcal{I} .

Annotated functions are ranged over by e, e' . The annotations are introduced using $e@p$, where e is an unannotated Alg expression, and p is a single participant identifier. These annotations need to be set by the programmer, but their introduction can be also automated. Additionally, we introduce the choice point annotations: $[p \oplus \vec{p}]$. This annotation specifies that p performs a choice, and notifies \vec{p} . Choice points can be introduced fully automatically by collecting all participants whose behaviour depends on the value of a sum type.

3.2 Interfaces

An interface represents a *state* in a concurrent system: the set of participants, and the types of the values that they

contain. We use mappings from participants to values to represent such states: $V := [p \mapsto v]_{p \in \mathcal{P}}$. The programmer, additionally to writing an Alg (PALg) expression, will need to provide an *input interface*, i.e. *where* is the input to the parallel program. Consider, for example, the interface $\text{int}@p_i$. Given a concurrent system with participants $p_0 \cdots p_n$, we know that p_i contains a value of type int : $[\cdots p_i \mapsto 42 \cdots]$. An interface with a product of participants $(a \times b)@(p_1 \times p_2)$ represents a state in which p_1 contains an element of type a , and p_2 an element of type b , e.g a possible state represented by $(\text{int} \times \text{vec})@(p_1 \times p_2)$ is: $[\cdots p_1 \mapsto 42 \cdots p_2 \mapsto [1, 1, 2, \dots] \cdots]$. An interface $\iota_i I$ represents the same state as interface I , but we statically know that this state was reached after an i -th injection. Then, if a participant requires the value at I , this participant will apply the necessary injections to the received values. Finally, an interface $a@(R_1 \cup^{\vec{p}} R_2)$ means that the state might be either R_1 or R_2 , and that all participants \vec{p} should be notified of the state.

Well-formedness The above examples are of well-formed interfaces: $\text{int}@p_i$, $(\text{int} \times \text{vec})@(p_1 \times p_2)$. Well-formedness ensures that interfaces represent valid states. Generally, $a@R$ is well-formed if a matches the structure of R . For example, $\text{int}@(p_1 \times p_2)$ is *ill-formed*, since a *single* integer cannot be at *two* different participants. An interface $a@(R_1 \cup R_2)$ requires that both $a@R_1$ and $a@R_2$ are well-formed. So, $(\text{vec} \times \text{vec})@((p_1 \times p_2) \cup p_3)$ is well-formed because we can have $\text{vec}@p_1$ and $\text{vec}@p_2$, or $(\text{vec} \times \text{vec})@p_3$. However, $\text{int}@((p_1 \times p_2) \cup p_3)$ is ill-formed, because $\text{int}@(p_1 \times p_2)$ is ill-formed.

3.3 Typing of Parallel Algebraic Language

We introduce a relation that associates Alg expressions with potential parallelisations PALg, and their interfaces. This relation can be seen as a type system for both Alg and PALg. As a type system for PALg, this relation provides a way to check or infer the output interface of some e . By using this relation as a type system for Alg, we can explore potential parallelisations of some input expression e . Additionally, the type system ensures that all *choice point* annotations contain every participant that depends on each particular choice.

Typing Rules A judgement of the form $\vdash e \Rightarrow e : A \rightarrow B$ means that the PALg expression e is one potential parallelisation of the Alg expression e , with domain interface A and codomain interface B . The intuition of a judgement $\vdash e \Rightarrow e : a@R_1 \rightarrow b@R_2$ is that the participants in e collectively apply computation e to the value of type a distributed across R_1 , and produce a value of type b distributed across R_2 . We sometimes omit e and write $\vdash e : A \rightarrow B$. We ensure that given any e and e such that they are typeable against interfaces $a@R_a \rightarrow b@R_b$, then e must have type $a \rightarrow b$.

Lemma 3.1. *If $e \Rightarrow e : a@R_a \rightarrow b@R_b$, then $e : a \rightarrow b$.*

The typing rules (Fig. 2) must ensure that the participants involved in a choice are notified, and that Alg expressions

are correctly expanded. Rule **CHOICE** specifies that a choice point may be introduced at any point when a participant contains a value of a sum-type. In such cases p sends the tag of the sum-type value to any other participant whose behaviour depends on it. After the choice point, the interface is $I[l_1 p] \cup^{\vec{p}} I[l_2 p]$, with the constraint that the participants in $\mathcal{I}[p]$ must be in \vec{p} . Rule **ALT** specifies that e must be the parallelisation of e , considering both A_1 and A_2 as input interfaces. The output interface is the union of B_1 and B_2 . Any participant in e must be notified of the choice $\text{pids}(e) \subseteq \vec{p}$, to make sure that they perform the interactions that correspond to the correct A_i . Rule **ALG** specifies that given any e and participant p , $e@p$ is a valid parallelisation, with output interface $b@p$. Finally, rule **EXT** is crucial for exploring potential parallelisations. It states that if e is the parallelisation of e_2 , and e_2 is extensionally equal to e_1 , then e is also a parallelisation of e_1 . The undecidability of this rule requires that the programmer specifies *rewriting strategies* both for checking and inference.

Rewriting and Annotation Strategies We use rewriting strategies when exploring potential parallelisations of functions. This is inference problem (2) below. Let $?i$ be metavariables. The two inference problems that we are interested in are: 1. Solving $\vdash e \Rightarrow e : A \rightarrow ?0$ obtains the output interface for e , with input interface A . 2. Solutions of $\vdash e \Rightarrow ?0 : A \rightarrow ?1$ are potential parallelisations of e , and their output interface. Solving (1) is straightforward. Problem (2) requires to decide how to introduce role annotations (rule **ALG**), how to perform rewritings (rule **EXT**), and where to introduce choice points (rule **CHOICE**). Introducing choice points is straightforward: we introduce them as early as possible, as soon as an input interface contains a sum-type at a participant. For introducing annotations and doing Alg rewritings, the programmer has to specify *annotation* and *rewriting* strategies. At the moment, our tool allows the developer to introduce annotations explicitly, or to select sub-expressions that will be annotated with fresh new participants. The rewriting strategies that our current implementation supports are unrollings of recursive definitions. However, our tool is extensible: the equivalences used in the rewritings are a parameter.

Example 3.2 (Mergesort). Consider the mergesort definition $\text{ms} = \text{rec}_T \text{mrg spl}$. Solutions to the inference problem $\vdash \text{ms} \Rightarrow ?0 : \text{Ls}@p_0 \rightarrow ?1$ provide the alternative parallelisations of ms . By choosing a rewriting strategy that unrolls ms once, and annotates any remaining instances of ms at fresh new participants, we produce the following PALg expression:

$$\begin{aligned}
 & \vdash (\text{id} \nabla (\text{mrg} \circ (\text{ms} \times \text{ms}))) \circ \text{spl} \\
 & \Rightarrow (\text{id} \nabla (\text{mrg}@p_1 \circ (\text{ms}@p_2 \circ \pi_1@p_1) \Delta (\text{ms}@p_3 \circ \pi_2@p_1))) \\
 & \quad \circ [p_1 \oplus p_1 p_2 p_3] \circ \text{spl}@p_1 \\
 & : \text{Ls}@p_0 \rightarrow \text{Ls}@p_1 \cup^{p_1 p_2 p_3} \text{Ls}@p_1
 \end{aligned}$$

$$\begin{array}{c}
 \text{ALG} \\
 \frac{\vdash e : a \rightarrow b}{\vdash e \Rightarrow e@p : a@I \rightarrow b@p} \\
 \\
 \text{EXT} \\
 \frac{\vdash e_2 \Rightarrow e : a@I \rightarrow B \quad e_1 =_{\text{ext}} e_2}{\vdash e_1 \Rightarrow e : a@I \rightarrow B} \\
 \\
 \text{ALT} \\
 \frac{\vdash e \Rightarrow e : A_1 \rightarrow B_1 \quad \vdash e \Rightarrow e : A_2 \rightarrow B_2 \quad A_1 \neq A_2 \quad \text{pids}(e) \subseteq \vec{r}}{\vdash e \Rightarrow e : A_1 \cup^{\vec{p}} A_2 \rightarrow B_1 \cup^{\vec{p}} B_2} \\
 \\
 \text{CHOICE} \\
 \frac{\text{pids}(\mathcal{I}[p]) \subseteq \vec{p} \quad \text{WF}(a@I[l_i p]), i \in [1, 2] \quad \vdash e \Rightarrow e : a@(\mathcal{I}[l_1 p] \cup^{\vec{p}} \mathcal{I}[l_2 p]) \rightarrow B}{\vdash e \Rightarrow e \circ [p \oplus \vec{p}] : a@I[p] \rightarrow B}
 \end{array}$$

Figure 2. Typing rules of PALg (selected)

4 Multiparty Session Types for PALg

The dataflow of the PALg constructs determine the communication protocol of the annotated expression. However, it is hard to manually check what this communication structure is. Recall the mergesort PALg expression of §3, ms , and suppose that we want to produce a parallelisation for a 32-core machine. Then, we might be interested in using a 5-unfolding of ms , so that we have ms executing concurrently on all of the cores. How do we know, for such cases, that we produced a *sensible* parallelisation? As an example, suppose we use an annotation strategy that produces the following code:

$$\begin{aligned}
 & (\text{id} \nabla (\text{mrg}@p_1 \circ (\text{ms}@p_2 \circ \pi_1@p_1) \Delta (\text{ms}@p_2 \circ \pi_2@p_1))) \\
 & \quad \circ [p_1 \oplus p_1 p_2] \circ \text{spl}@p_1 : \text{Ls}@p_0 \rightarrow \text{Ls}@p_1 \cup^{p_1 p_2} \text{Ls}@p_1
 \end{aligned}$$

Notice that this example will run correctly, and produce the expected result. However, the achieved PALg expression is *not* parallel! If we represent the *implicit dataflow* of this expression as *explicit communication*, the reason becomes apparent. We use *global types* from *multiparty session types* to provide an explicit representation of the communication structure of the program:

$$\begin{aligned}
 & p_0 \rightarrow p_1 : \text{Ls}. p_1 \rightarrow p_2 \{t_1. \text{end}; \\
 & \quad t_2. p_1 \rightarrow p_2 : \text{Ls}. p_1 \rightarrow p_2 : \text{Ls}. p_2 \rightarrow p_1 : \text{Ls} \times \text{Ls}. \text{end}\}
 \end{aligned}$$

This global type represents the following protocol: 1. participant p_0 sends a list to p_1 ; 2. p_1 sends to p_2 either t_1 or t_2 , and if the label is t_1 , the protocol ends; 3. if p_1 sent t_2 , then p_1 sends to p_2 *two* lists, in two different interactions; and 4. p_2 replies with a message to p_1 with a pair of lists. It is clear from this protocol that p_1 and p_2 are dependent on each others' messages, and that p_2 cannot perform any computation in parallel. The larger the expression is, the harder avoiding these wrong annotations will become. By changing the annotation strategy, we produce the following parallel structure, where p_2 and p_3 can operate in parallel:

$$\begin{aligned}
 & p_0 \rightarrow p_1 : \text{Ls}. p_1 \rightarrow \{p_2 p_3\} \{t_1. \text{end}; \\
 & \quad t_2. p_1 \rightarrow p_2 : \text{Ls}. p_1 \rightarrow p_3 : \text{Ls}. p_2 \rightarrow p_1 : \text{Ls}. p_3 \rightarrow p_1 : \text{Ls}. \text{end}\}
 \end{aligned}$$

This abstraction of the communication protocol of an achieved parallelisation is therefore useful as an output for the programmer. Additionally, these global types are a *contract* that

can be enforced on the generated code. We use this for proving that our back-end is correct, but also for applying low-level code optimisations (e.g. message reordering) guided by this global type, ensuring that they do not introduce any run-time error. For example, when we find in a global type $p_1 \rightarrow p_2$, $p_2 \rightarrow p_3$, we mark the send/receive actions for p_2 as point of potential optimisation. If the messages exchanged do not depend on each other, we permute them, performing first the send action, so that p_2 is not blocked by a receive action. This is known as *asynchronous optimisation* [57].

4.1 Multiparty Session Types

Our global types are based on the most commonly used in the literature [22]. We start with a set of *participant identifiers*, p_1, p_2, \dots , and a set of *labels*, ι_1, ι_2, \dots . These are considered as natural numbers: participant identifiers uniquely identify an independent unit of computation, e.g. thread or process ids; and labels are tags that differentiate branches in the data/control flow. The syntax of global (G) and local (L) types in MPST is given as:

$$G ::= p_1 \rightarrow p_2 : a.G \mid p_1 \rightarrow \{p_j\}_{j \in [2, n]} : \{ \iota_i.G_i \}_{i \in I} \mid \mu X.G \mid X \mid \text{end}$$

$$L ::= p!(a).L \mid p?(a).L \mid p \& \{ \iota_i.L_i \}_{i \in I} \mid \{p_j\}_{j \in [2, n]} \oplus \{ \iota_i.L_i \}_{i \in I} \mid \mu X.L \mid X \mid \text{end}$$

Global type $p_1 \rightarrow p_2 : a.G$ denotes *data* interactions from p_1 to p_2 with value of type a ; *Branching* is represented by $p_1 \rightarrow \{p_j\}_{j \in [2, n]} : \{ \iota_i.G_i \}_{i \in I}$ with actions ι_i from p_1 to all $p_j, j \in [2, n]$. *end* represents a *termination* of the protocol. $\mu X.G$ represents a *recursive* protocol, which is *equivalent* to $[\mu X. G/X]G$. We assume recursive types are guarded.

Each participant in G represents a different participant in a parallel process. *Local session types* represent the communication actions performed by each participant, i.e. the *role* of the participant. Since each participant has a unique role, we sometimes refer to them interchangeably. The *send* type $p!(a).L$ expresses the action of sending of a value of type a to p followed by interactions specified by L . The *receive* type $p?(a).L$ is the dual, where a value with type a is received from p . The *selection* type represents the transmission to all p_j of label ι_i chosen in the set of labels ($i \in I$) followed by L_i . The *branching* type is its dual. $\text{pids}(G)/\text{pids}(L)$ denote the set of participants that occur in G/L .

Projection We use a standard definition of *projection* that uses the *full merging* operator [24; 27], which allows more well-formed global types than the original projection rules [42]. We write $G \upharpoonright p$ for the projection of G onto the role of p . We illustrate the projection with an interaction $p_0 \rightarrow p_1 : a.G$. The projection onto p_0 is $p_1!(a).(G \upharpoonright p_0)$, the projection onto p_1 is $p_0?(a).(G \upharpoonright p_1)$, and the projection onto any other role p is $G \upharpoonright p$. Projection on choices is similar, with the difference that whenever the role is not at the receiving or sending ends of the choice, the different branches must be *merged*. Two local types can be merged when they

CHOICE

$$\frac{}{\models [p \oplus \bar{p}] \Leftarrow a[b+c]@I[p] \sim p \rightarrow \{\bar{p} \setminus p\} \{ \iota_1.\text{end}; \iota_2.\text{end} \}}$$

ALT

$$\frac{\models e \Leftarrow A_1 \sim G_1 \quad \models e \Leftarrow A_2 \sim G_2}{\models e \Leftarrow A_1 \cup^{\bar{p}} A_2 \sim G_1 \cup G_2}$$

ALG

$$\vdash e : a \rightarrow b$$

$$\frac{}{\models e@p \Leftarrow a@I \sim [a@I \rightsquigarrow p]}$$

$$\begin{aligned} [a@p_1 \rightsquigarrow p_2] &= p_1 \rightarrow p_2 : a.\text{end}, \text{ if } p_1 \neq p_2; [a@p \rightsquigarrow p] = \text{end}; \\ [(a \times b)@I_a \times I_b \rightsquigarrow p] &= [a@I_a \rightsquigarrow p] \text{ ; } [b@I_b \rightsquigarrow p]; \text{ and} \\ [(a_1 + a_2)@(\iota_i I) \rightsquigarrow p] &= [a_i@I \rightsquigarrow p] \end{aligned}$$

Figure 3. Protocol Relation (selected)

are the same, or they branch on the same role, and their continuations can be merged.

We use a standard definition of *well-formedness* that states that a global type is well formed if its projection on all its roles is defined. We denote: $\text{WF}(G) = \forall p \in \text{pids}(G), \exists L, G \upharpoonright p = L$.

4.2 Protocol Relation


We introduce now the set of rules that associate a PALg expression and domain interface with their global type (Fig. 3). We extend the syntax of global types with $G_1 \cup^{\bar{p}} G_2$ to represent the *external choices*, i.e. G_i are the continuations for both branches of a previous choice that affects \bar{p} . We also extend the local types, and projection rules ($G_1 \cup^{\bar{p}} G_2$) = $G_1 \upharpoonright p \cup^{\bar{p}} G_2 \upharpoonright p$, and the notion of well-formedness. We say that an external choice is well-formed, $\text{WF}(G_1 \cup^{\bar{p}} G_2)$, if $\text{WF}(G_1)$, $\text{WF}(G_2)$, and for all $p \notin \bar{p}$, $G_1 \upharpoonright p = G_2 \upharpoonright p$. We omit the annotation of the participants involved in the choice whenever it is not needed. The relation $\models p \Leftarrow A \sim (G, B)$ specifies that the parallel code for p and input interface A will behave as global type G , and output interface B (Fig. 3). The rules are similar to the typing rules of PALg.

Example 4.1 (Mergesort Protocol). The protocol for Example 3.2 is obtained by solving:

$$\models (\text{id} \nabla (\text{mrg}@p_1 \circ (\text{ms}@p_2 \circ \pi_1@p_1) \Delta (\text{ms}@p_3 \circ \pi_2@p_1))) \circ [p_1 \oplus p_1 p_2 p_3] \circ \text{spl}@p_1 \Leftarrow \text{Ls}@p_1 \sim ?\mathbf{0}.$$

$$p_1 \rightarrow \{p_2 p_3\} \left\{ \begin{array}{l} \iota_1.\text{end}; \\ \iota_2.p_1 \rightarrow p_2 : \text{Ls}.p_1 \rightarrow p_3 : \text{Ls}.\text{end} \end{array} \right\} ; \\ (\text{end} \cup (p_2 \rightarrow p_1 : \text{Ls}.p_3 \rightarrow p_1 : \text{Ls}.\text{end}))$$

$$= p_1 \rightarrow \{p_2 p_3\} \left\{ \begin{array}{l} \iota_1.\text{end}; \\ \iota_2.p_1 \rightarrow p_2 : \text{Ls}. \\ p_1 \rightarrow p_3 : \text{Ls}. \\ p_2 \rightarrow p_1 : \text{Ls}. \\ p_3 \rightarrow p_1 : \text{Ls}.\text{end} \end{array} \right\}$$



4.3 Correctness

We guarantee that for e s.t. $\vdash e \Rightarrow e : A \rightarrow B$, with A and B well-formed, there exists a protocol G and that it is well-formed and deadlock-free.

Lemma 4.2. [Existence of Associated Global Type] For all $\text{WF}(A)$, if $\vdash e : A \rightarrow B$, then there exists G s.t. $\models e \Leftarrow A \sim G$.

Lemma 4.3. [Protocol Deadlock-Freedom] For all $\text{WF}(A)$, if $\vdash e : A \rightarrow B$ and $\models e \Leftarrow A \sim G$, then $\text{WF}(G)$.

Remark. Since the local type abstracts the behaviour of multiparty typed processes, a well-formed global type ensures the end-point processes (programs) typed by that global type are guaranteed to satisfy the properties (such as safety and deadlock-freedom) of local types [27; 43].

5 Code Generation

This section addresses the problem of generating low-level parallel code from PALg expressions. We prove that the generated code complies with its inferred protocol, which has several implications: (1) code generation does not introduce any concurrency errors, and the parallel code is therefore deadlock-free; and (2) we can prove that the generated code is extensionally equal to the input expression by considering only a representative trace, since any valid interleaving of actions must respect this protocol. The target language of our tool is an indexed monad, the *Message Passing Monad* (Mp). From Mp, we implement our low-level C backend. We implement an untyped version of Mp as a deep embedding in Haskell, and session typing on top of it. This is suitable for code generation: we only generate parallel code if the monadic actions are typeable against the respective local types. Our definition of Mp has significant differences to other embeddings of session types in Haskell, such as the Session monad by Neubauer and Thiemann [58]. First, our Mp monad is deeply embedded in Haskell, and secondly, we use type indices instead of an encoding of session types in terms of type classes. Our approach is better suited for compilation since we manipulate session types, and postpone session typing until code generation.

5.1 Message Passing Monad

Mp comprises four basic operations: send, receive, choice and branching, with a standard (asynchronous) semantics. Additionally, for composing actions that depend on the same choice, we introduce case expressions. Our definition of Mp is based on the *free monad* construction:

$$\begin{aligned} v &::= x \mid (v, v) \mid \iota_i v \mid \dots \mid e v \\ m_i &::= \text{ret } v \mid \text{send } p \ v \ m \mid \text{recv } p \ a \ f \mid \text{sel } \vec{p} \ v \ f_1 \ f_2 \\ &\quad \mid \text{brn } p \ m_1 \ m_2 \mid \text{case } f_1 \ f_2 \ f ::= \lambda x. m \end{aligned}$$

Values v are either primitive values, tagged values $\iota_i v$, pairs of values, or the result of applying an Alg expression e to a value. We use standard notation for the monadic unit (**ret**), bind (\gg) and Kleisli composition: $f_1 \gg f_2 = \lambda x. f_1 \ x \gg f_2$. The message-passing constructs are standard, except **sel**, **brn** and **case**, which are used for performing choices, and composing actions that depend on the same choice.

Each monadic computation f or m has a type $m : \text{Mp } L \ a$, where a is the return type of m , and L is the type index of Mp, and it represents the local type that corresponds to the behaviour of the term m . There is almost a one to one correspondence between the terms L and the monadic actions m , so we omit the full definition. The types of the constructs that deal with choices use a new type, \uplus , that is isomorphic to sum types, but that can only be constructed and eliminated by using the corresponding monadic constructs:

$$\begin{aligned} \text{sel } \vec{p} &: a + b \rightarrow (a \rightarrow \text{Mp } L_1 \ c_1) \rightarrow (b \rightarrow \text{Mp } L_2 \ c_2) \\ &\quad \rightarrow \text{Mp } (\vec{p} \oplus \{\iota_1.L_1; \iota_2.L_2\}) \ (c_1 \uplus c_2) \\ \text{brn } p &: \text{Mp } L_1 \ a_1 \rightarrow \text{Mp } L_2 \ a_2 \\ &\quad \rightarrow \text{Mp } (p \ \& \ \{\iota_1.L_1; \iota_2.L_2\}) \ (a_1 \uplus a_2) \\ \text{case} &: (a \rightarrow \text{Mp } L_1 \ c) \rightarrow (b \rightarrow \text{Mp } L_2 \ d) \rightarrow a \uplus b \\ &\quad \rightarrow \text{Mp } (L_1 \cup L_2) \ (c \uplus d) \end{aligned}$$

These constructs ensure that the tag used to build $a \uplus b$ indeed corresponds to the correct branch of the right choice. We use **case** to compose actions that depend on a previous choice. While this treatment of \uplus leads to unnecessary code duplication, our back-end easily optimises cases where we have **case** $f \ f$ to avoid code duplication.

Parallel programs We define the basic constructs of PALg in a bottom-up way by manipulating *parallel programs*. Parallel programs are mappings from participants to their monadic action: $\mathbf{E} ::= [p_i \mapsto m_i]_{i \in I}$. If $m_i : \text{Mp } L_i \ a_i$ for all $i \in I$, then we write $[p_i \mapsto m_i]_{i \in I} : \text{Mp } [p_i \mapsto L_i]_{i \in I} \ [p_i \mapsto a_i]_{i \in I}$. The semantics of both local types and monadic actions is defined in terms of such collections of actions or local types, and shared queues of values W , or queues of types Q , e.g. $\langle \mathbf{E}, W \rangle \rightsquigarrow^\ell \langle \mathbf{E}', W' \rangle$ is a transition from \mathbf{E} to \mathbf{E}' , and shared queues W to W' with *observable action* ℓ . We prove a standard safety theorem (Theorem 5.1 below) that guarantees that if a participant does a transition with some observable action, then so does the type index.

Theorem 5.1. [Soundness] Assume $\mathbf{E} : \text{Mp } C \ A$, $m : \text{Mp } L \ a$ and $W : Q$. Suppose $\langle \mathbf{E}[r \mapsto m], W \rangle \rightsquigarrow^\ell \langle \mathbf{E}[r \mapsto m'], W' \rangle$. Then there exists $\langle C[r \mapsto L], Q \rangle \rightarrow^\ell \langle C[r \mapsto L'], Q' \rangle$ such that $W' : Q'$ and $m' : \text{Mp } L' \ a$.

Mp code generation The translation scheme for Mp code generation is done recursively on the structure of PALg expressions. It takes a PALg expression e , an interface A , and produces a mapping from all participants in e and A to their respective monadic continuations. We write $\llbracket e \rrbracket (A)$, and guarantee that $\llbracket e \rrbracket (A) : A \rightarrow \text{Mp } G \ B$, if $\models e \Leftarrow A \sim (G, B)$. This means that if e induces protocol G with interfaces $A \rightarrow B$, then the generated code behaves as G , with interfaces A and B . Code generation follows a similar structure to global type inference, and is defined by building PALg constructs as Mp parallel programs. For example, the translation of $e @ p : a @ I \rightarrow B$ requires to define the interactions from an interface I that gathers a type a at p : $(\llbracket a @ I \rightsquigarrow p \rrbracket) : a @ I \rightarrow \text{Mp } [a @ I \rightsquigarrow p] \ (a @ p)$. The definition

is analogous to that of $[a@I \rightsquigarrow p]$. The remaining of the translation is straightforward, so we skip the details.

We prove two main correctness results. We guarantee that the generated code behaves as its inferred protocol (Theorem 5.2). We also guarantee that regardless of the annotations and interfaces chosen for e , *the parallel code always produces the same result as the sequential implementation* (Theorem 5.3).

Theorem 5.2. [Protocol Conformance of the Generated Code] *If $\models e \Leftarrow A \sim G$, then $\llbracket e \rrbracket (A)$ complies with protocol G .*

Theorem 5.3. [Extensionality] *Assume $e \Rightarrow e : a@p \rightarrow b@R$ and $x : a$ initially at p . If $e x = y$, then the execution of $\llbracket e \rrbracket (p)$ also produces y , distributed across R .*

Example 5.4 (MergeSort Code Generation). We show below the code generation for `ms` (Example 3.2), with p_1 as domain interface:

```

p1 ↦ λx. sel {p2, p3} (spl x) (λx. ret x)
      (λx. send p2 (π1 x) ≫ λy. send p3 (π2 x) ≫ λ_.
      recv p2 Ls ≫ λx. recv p3 Ls ≫ λy. ret (mrg (x, y)))
p2,3 ↦ λx. brn p1 (ret x) (recv p1 Ls ≫ λx. send p1 (ms x))

```

6 Parallel Algorithms and Evaluation

We evaluate our approach using a number of parallel algorithms derived from Alg expressions, and the speedups achieved. The purpose of this is twofold: (i) showing that our approach achieves speedups for an input sequential algorithm, with naïve annotation strategies, and limited optimisations (Fig. 5), and (ii) illustrating the practical value of providing a global type that describes the parallel strategy achieved by a particular annotation strategy (Fig. 4). We run all our experiments on 2 NUMA nodes, 12 cores per node and 62GB of memory, using Intel Xeon CPU E5-2650 v4 @ 2.20GHz chips. We run our experiments first restricting the execution to a single node to avoid NUMA effects, and then on the 2 NUMA nodes.

6.1 Benchmarks

Mergesort Mergesort is the usual divide-and-conquer algorithm, using a tree-like parallel reduce.

Cooley-Tukey FFT We use a recursive Cooley-Tukey algorithm. The algorithm starts by splitting the elements of the list into those that are at even and odd positions. Then, it recursively computes the FFT of them, and finally combines the results. To generate a butterfly pattern, we use: products of size n , to store the results of the subsequent interleavings; product associativity to produce a perfect tree; and *asynchronous optimisations*.

Dot Product The dot product algorithm zips the inputs, multiplies them pairwise, and then adds them by folding the result. We use products of size n to derive a scatter-gather.

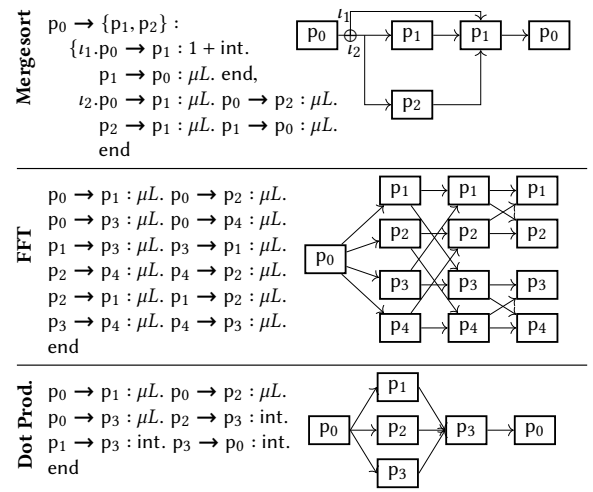


Figure 4. Benchmarks: potential parallelisations.

Additional Algorithms We implemented *scalar prod*, that recursively splits a matrix into sub-matrices, distributes them to different workers, and then multiplies their elements by a scalar, and *quicksort*, with a divide-and-conquer structure.

6.2 Evaluation

We translate Mp monadic actions to C using pthreads and shared buffers for communication, and we have a preliminary compilation of the first-order sequential terms to C. We compile the generated C code using gcc version 4.8.5. We take the average of 50 repetitions for each benchmark. Our benchmarks achieve reasonable speedups against the sequential C implementations. Fig. 5 presents the speedups against the number of participants for different input sizes, and Fig. 6 present a summary of our speedups for large inputs of size $> 10^9$. We show below an analysis of these results, by plotting the speedups against two factors: 1. the number of participants (threads) produced by a particular annotation and recursion unrolling, named K ; and 2. the input size, e.g. number of elements in the input list.

Increasing the number of threads (parameter K), increases the speedups obtained, up to a certain value that depends on the amount of available cores and the input size. For benchmarks that work better with dynamic task creation, our tool does not currently achieve good performance (e.g. quicksort). For FFT, our tool produces the usual butterfly pattern from a straightforward recursive definition, that we can achieve a speedup of 12 when running on a single shared-memory node. The rest of the examples are limited either by *Amdahl's law* (justified by their global types in Fig. 4), or by the overhead of the communication and pthread creation with respect to the cost of the computations, but still achieve speedups of up to 7 and 8 on 12 cores. We can observe that there is a slow down after creating a much larger number of participants than the ones required. This usually depends on how evenly we can distribute the data amongst workers, and whether

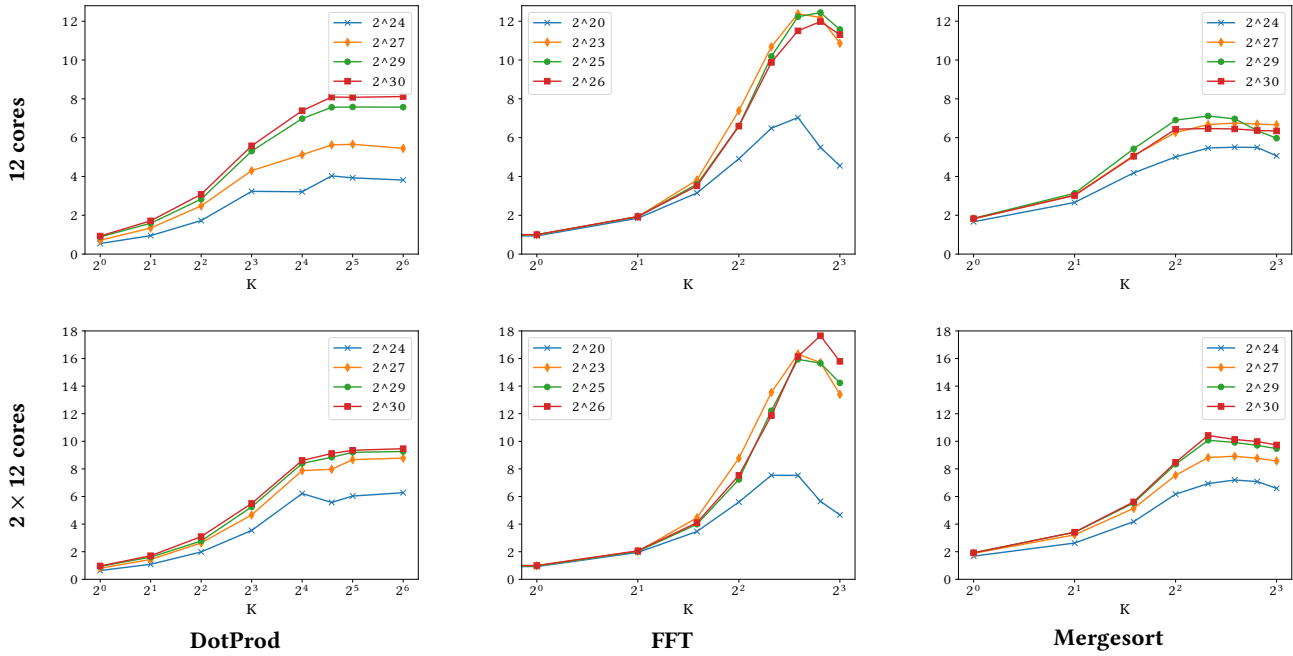


Figure 5. Benchmark speedups, run in 2 NUMA nodes with 12 cores each. The X-axis is the number of workers of the parallel program generated from a set of annotations and recursion unrolling. We show the results for 4 different input sizes.




Figure 6. Achieved speedups

the amount of workers can be evenly scheduled to different cores. We observe that we can achieve further speedups when running our benchmarks in the 2 NUMA nodes. Overall, we observe that our annotation strategies enable good speedups over the sequential implementation, with relatively little effort. Global types can be used to detect optimisation opportunities that yield efficient parallelisations, such as the Butterfly topology in Fig. 4. Without message-reordering based on the session types, FFT participant p_3 would need to wait for p_1 's message before sending its part to p_1 , i.e. p_3 's local type would be $p_1! \langle \mu L \rangle . p_1! \langle \mu L \rangle . \dots$. This means that p_3 's local computation would only become available to p_1 after

it p_1 finishes its own local computation, thus sequentialising the code. Asynchronous permutations [16; 57] allow us to *permute* such actions, and still have communication safety, i.e. $p_1! \langle \mu L \rangle . p_1? \langle \mu L \rangle . \dots$. Global types capture the structure of the parallelisation, which can in some cases be used to justify the achieved speedups. For example, we can observe that the mergesort global type contains a part that needs to happen sequentially (p_0 and the last merging point in p_1), and this will prevent us from achieving linear speedups.

7 Related Work

López et al. [50] develop a verification framework for MPI/C inspired by MPST by translating parameterised protocol specifications to protocols in VCC [19]. They focus on verification, not on code or protocol generation. Ng et al. [59; 60] use parameterised MPST [25] to generate an MPI backbone in C that encapsulates the *whole* protocol (i.e., every endpoint), and merges it with user-supplied computation kernels. Several authors (e.g. [10]) generate skeleton API from extensions of Scribble (www.scribble.org). Their approach requires the protocol to be specified beforehand, and it is not extracted from sequential code. Unlike ours, none of the above work formally defines code generation or proves its correctness.

Structured parallelism includes the use of high-level constructs in languages with implicit/data parallelism [5; 12–15; 46; 64], algorithmic skeleton APIs [1; 18; 20; 36; 48], and DSLs/APIs that compile to parallel code [8; 11; 28; 63; 69]. Besides safety, such approaches are often highly optimised.

991 However, most rely on using a fixed, predetermined range of
 992 patterns, typically by design with respect to their application
 993 domains. By contrast, our work only relies on send/receive
 994 operations, which makes it highly portable, and can be easily
 995 extended to support further parallel structures by extend-
 996 ing the annotation strategies. Optimisations for structured
 997 parallel approaches also require to study and define a set
 998 of equivalences between patterns [6; 7; 41]. In contrast, our
 999 approach does not require the definition of new sets of equiv-
 1000 alences, since these are derived from program equivalences.
 1001 *Lift* is a new language for portable parallel code genera-
 1002 tion, based on a small set of expressive parallel primitives
 1003 [40; 67; 68]. Currently, their backend focuses on generat-
 1004 ing high-performance OpenCL code, while our approach
 1005 focuses on placing computations on different participants
 1006 of a concurrent/distributed system. Both approaches could
 1007 be combined: annotations can be used to generate a high-
 1008 level message-passing layer that distributes tasks to multiple
 1009 nodes in a GPU cluster, using the global type to minimise
 1010 communication costs; then, the code at each participant can
 1011 be compiled to high-performance GPU code using *Lift*.

1012 Elliott exploits the idea of giving functional programs
 1013 multiple interpretations in different categories, and shows
 1014 examples of applications to multiple domains, including par-
 1015 allelism [29; 30]. Our approach is similar in the sense that we
 1016 allow the specifications of first-order functional programs
 1017 to have multiple different interpretations, but we focus on
 1018 generating parallel code, and provide a finer-grained con-
 1019 trol over the parallelisations by adding participant anno-
 1020 tations. There is a large body of literature in using pro-
 1021 gram equivalences to derive parallel implementations, e.g.
 1022 [17; 32; 37; 39; 49; 51; 55; 56; 65; 66]. Our framework is or-
 1023 thogonal, in that we focus on tying a low-level C back-end
 1024 with global types. Our front-end, however, supports some
 1025 basic form of rewritings, and we plan to extend it in the
 1026 future with more interesting ones from the literature.

1027 8 Conclusions and Future Work

1030 We have presented a novel approach to protocol inference
 1031 and code generation. By using this approach, we can reason
 1032 about *extensionality* of the parallel programs, and alternative
 1033 mappings of computations to *participants*. We produce the
 1034 parallel program global type, i.e. its communication protocol,
 1035 that acts as a contract for the low-level code, can be used to
 1036 pin-point potential optimisations, or assessing the suitability
 1037 of a parallelisation. This approach has several benefits:
 1038 1. our message-passing code is deadlock-free by construc-
 1039 tion, since it follows the data-flow of the program, and the
 1040 optimisations must respect the global type; 2. we prove that
 1041 our parallelisations are extensionally equivalent to the input
 1042 function. Additionally, PAlg code could be used for further
 1043 multiple purposes, such as parallel GPU/FPGA code genera-
 1044 tion, by combining our approach with other state of the art
 1045

code generation techniques. We will study this for future
 work.

Though our approach can already generate representative
 parallel protocols, our framework is extensible. E.g. we can
 extend our framework with dynamic participants to handle
 dynamic task generation [26], and we plan to use this to
 capture a wider range of communication patterns for paral-
 lel computing, such as load-balancing or work-stealing. We
 plan to study the extension of our back-end to heterogeneous
 architectures, e.g. GPU clusters, or FPGAs. Our prototype
 generates code that can achieve speedups against sequential
 implementations, the optimisations that we support are very
 basic, and our generated code can be very large. We plan to
 introduce optimisations that reduce the amount messages
 exchanged, further message reorderings guided by the global
 type, and optimisations of the size of the generated code. Fi-
 nally, we plan to study the instrumentation of global types to
 estimate statically the speedups of different parallelisations,
 and optimise communication costs.

Acknowledgements

We thank Shuhao Zhang for his contributions to the C back-
 end, described in [70]. We thank Francisco Ferreira for the
 helpful discussions in the early stages of this work. This
 work was supported in part by EPSRC projects EP/K011715/1,
 EP/K034413/1, EP/L00058X/1, EP/N027833/1, EP/N028201/1,
 and EP/T006544/1.

References

- [1] Marco Aldinucci, Marco Danelutto, Peter Kilpatrick, Massimiliano Meneghin, and Massimo Torquati. 2011. Accelerating Code on Multi-cores with FastFlow. In *Proc. of 17th International Euro-Par Conference (Euro-Par 2011) (LNCS)*, Vol. 6853. Springer, 170–181.
- [2] John Backus. 1978. Can Programming Be Liberated from the Von Neumann Style?: A Functional Style and Its Algebra of Programs. *Commun. ACM* 21, 8 (Aug. 1978), 613–641.
- [3] Richard Bird and Oege De Moor. 1996. The algebra of programming. In *NATO ASI DPD*. 167–203.
- [4] R. S. Bird. 1989. Algebraic Identities for Program Calculation. *Comput. J.* 32, 2 (01 1989), 122–126.
- [5] Guy E. Blelloch. 1996. Programming Parallel Algorithms. *Commun. ACM* 39, 3 (1996), 85–97.
- [6] Christopher Brown, Marco Danelutto, Kevin Hammond, Peter Kilpatrick, and Archibald Elliott. 2014. Cost-Directed Refactoring for Parallel Erlang Programs. *International Journal of Parallel Programming* 42, 4 (01 Aug 2014), 564–582.
- [7] Christopher Brown, Kevin Hammond, Marco Danelutto, and Peter Kilpatrick. 2012. A Language-independent Parallel Refactoring Framework. In *Proc. of the 5th Workshop on Refactoring Tools (WRT '12)*. ACM, New York, NY, USA, 54–58.
- [8] Kevin J. Brown, Arvind K. Sujeeth, HyoukJoong Lee, Tiark Rompf, Hassan Chafi, Martin Odersky, and Kunle Olukotun. 2011. A Heterogeneous Parallel Framework for Domain-Specific Languages. In *Proc. of 2011 Intl. Conf. on Parallel Architectures and Compilation Techniques (PACT'11)*. IEEE, 89–100.
- [9] Jacques Carette, Oleg Kiselyov, and Chung-chieh Shan. 2009. Finally Tagless, Partially Evaluated: Tagless Staged Interpreters for Simpler Typed Languages. *J. Funct. Program.* 19, 5 (2009), 509–543.

- 1101 [10] David Castro, Raymond Hu, Sung-Shik Jongmans, Nicholas Ng, and
1102 Nobuko Yoshida. 2019. Distributed Programming Using Role Parametric
1103 Session Types in Go. In *46th Symp. on Principles of Programming
1104 Languages (POPL '19)*, Vol. 3. ACM, 29:1–29:30.
- 1105 [11] Hassan Chafi, Arvind K. Sujeeth, Kevin J. Brown, HyoukJoong Lee,
1106 Anand R. Atreya, and Kunle Olukotun. 2011. A Domain-Specific
1107 Approach to Heterogeneous Parallelism. In *Proc. of the 16th Symp.
1108 on Principles and Practice of Parallel Programming (PPOPP'11)*. ACM,
1109 35–46.
- 1109 [12] Manuel M. T. Chakravarty, Roman Leshchinskiy, Simon L. Peyton
1110 Jones, Gabriele Keller, and Simon Marlow. 2007. Data Parallel Haskell:
1111 a Status Report. In *Proc. of the POPL 2007 Workshop on Declarative
1112 Aspects of Multicore Programming, (DAMP'07) (2007)*. ACM, 10–18.
- 1113 [13] Bradford L. Chamberlain, David Callahan, and Hans P. Zima. 2007.
1114 Parallel Programmability and the Chapel Language. *IJHPCA* 21, 3
1115 (2007), 291–312.
- 1116 [14] Rohit Chandra. 2001. *Parallel Programming in OpenMP*. Morgan
1117 Kaufmann.
- 1118 [15] Philippe Charles, Christian Grothoff, Vijay A. Saraswat, Christopher
1119 Donawa, Allan Kielstra, Kemal Ebcioglu, Christoph von Praun, and
1120 Vivek Sarkar. 2005. X10: an Object-Oriented Approach to Non-Uniform
1121 Cluster Computing. In *Proc. of the 20th Conf. on Object-Oriented
1122 Programming, Systems, Languages, and Applications, (OOPSLA05)*. ACM,
1123 San Diego, CA, USA, 519–538.
- 1124 [16] Tzu-chun Chen, Mariangiola Dezani-Ciancaglini, Alceste Scalas, and
1125 Nobuko Yoshida. 2017. On the Preciseness of Subtyping in Session
1126 Types. *Logical Methods in Computer Science* Volume 13, Issue 2 (June
1127 2017).
- 1128 [17] Yun-Yan Chi and Shin-Cheng Mu. 2011. Constructing List Homomorphisms
1129 from Proofs. In *Proc. APLIAS '11: Asian Symposium on
1130 Programming Languages & Systems*. 74–88.
- 1131 [18] Philipp Ciechanowicz and Herbert Kuchen. 2010. Enhancing Muesli's
1132 Data Parallel Skeletons for Multi-core Computer Architectures. In
1133 *12th Intl. Conf. on High Performance Computing and Communications,
1134 (HPCC'10)*. IEEE, 108–113.
- 1135 [19] Ernie Cohen, Markus Dahlweid, Mark A. Hillebrand, Dirk Leinenbach,
1136 Michal Moskal, Thomas Santen, Wolfram Schulte, and Stephan Tobies.
1137 2009. VCC: A Practical System for Verifying Concurrent C. In *Proc. of
1138 the 22nd Intl. Conf. on Theorem Proving in Higher Order Logics, (TPHOLS
1139 2009) (LNCS)*, Vol. 5674. Springer, 23–42.
- 1140 [20] Murray Cole. 1988. *Algorithmic skeletons: a structured approach to the
1141 management of parallel computation*. Ph.D. Dissertation. University of
1142 Edinburgh, UK.
- 1143 [21] James W Cooley and John W Tukey. 1965. An Algorithm for the
1144 Machine Calculation of Complex Fourier Series. *Mathematics of
1145 computation* 19, 90 (1965), 297–301.
- 1146 [22] Mario Coppo, Mariangiola Dezani-Ciancaglini, Luca Padovani, and
1147 Nobuko Yoshida. 2015. A Gentle Introduction to Multiparty Asynchronous
1148 Session Types. In *15th International School on Formal Methods for
1149 the Design of Computer, Communication and Software Systems: Multi-
1150 core Programming (LNCS)*, Vol. 9104. Springer, 146–178.
- 1151 [23] Alcino Cunha, Jorge Sousa Pinto, and José Proença. 2006. A Framework
1152 for Point-Free Program Transformation. In *Proc. of 17th Intl. Workshop
1153 on Implementation and Application of Functional Languages (IFL 2005)*.
1154 Springer, 1–18.
- 1155 [24] Romain Demangeon and Kohei Honda. 2012. Nested Protocols in
1156 Session Types. In *CONCUR 2012 – Concurrency Theory*. Springer, 272–
1157 286.
- 1158 [25] Pierre-Malo Deniérou, Nobuko Yoshida, Andi Bejleri, and Raymond
1159 Hu. 2012. Parameterised Multiparty Session Types. *Logical Methods
1160 in Computer Science* 8, 4 (2012).
- 1161 [26] Pierre-Malo Deniérou and Nobuko Yoshida. 2011. Dynamic Multirole
1162 Session Types (POPL '11). ACM, New York, NY, USA, 435–446.
- 1163 [27] Pierre-Malo Deniérou and Nobuko Yoshida. 2013. Multiparty Compatibility
1164 in Communicating Automata: Characterisation and Synthesis of
1165 Global Session Types. In *40th International Colloquium on Automata,
1166 Languages and Programming (LNCS)*, Vol. 7966. Springer, 174–186.
- 1167 [28] Zach DeVito, Niels Joubert, Francisco Palacios, Stephen Oakeley,
1168 Montserrat Medina, Mike Barrientos, Erich Elsen, Frank Ham, Alex
1169 Aiken, Karthik Duraisamy, Eric Darve, Juan Alonso, and Pat Hanrahan.
1170 2011. Liszt: a Domain Specific Language for Building Portable Mesh-
1171 based PDE Solvers. In *Conference on High Performance Computing
1172 Networking, Storage and Analysis, SC 2011*. ACM, 9:1–9:12.
- 1173 [29] Conal Elliott. 2017. Compiling to Categories. *Proc. ACM Program.
1174 Lang.* 1, ICFP, Article 27 (Aug. 2017), 27 pages.
- 1175 [30] Conal Elliott. 2017. Generic functional parallel algorithms: Scan and
1176 FFT. *Proc. ACM Program. Lang.* 1, ICFP, Article 48 (Sept. 2017), 24 pages.
- 1177 [31] Maarten M. Fokkinga and Erik Meijer. 1991. *Program Calculation Prop-
1178 erties of Continuous Algebras*. Number CS-R91 in Report / Department
1179 of Computer Science. CWI.
- 1180 [32] Jeremy Gibbons. 1996. Computing Downwards Accumulations on
1181 Trees Quickly. *Theoretical Computer Science* 169, 1 (1996), 67–80.
- 1182 [33] Jeremy Gibbons. 1996. The Third Homomorphism Theorem. *Journal
1183 of Functional Programming (JFP)* 6, 4 (1996), 657–665.
- 1184 [34] Jeremy Gibbons. 2002. Calculating Functional Programs. In *Algebraic
1185 and Coalgebraic Methods in the Mathematics of Program Construction*.
1186 Springer, Chapter 5, 151–203.
- 1187 [35] Jeremy Gibbons. 2007. Datatype-Generic Programming. In *Datatype-
1188 Generic Programming*. Springer, 1–71.
- 1189 [36] Horacio González-Vélez and Mario Leyton. 2010. A survey of algorithmic
1190 skeleton frameworks: high-level structured parallel programming
1191 enablers. *Softw., Pract. Exper.* 40, 12 (2010), 1135–1160.
- 1192 [37] Sergei Gorlatch. 1999. Extracting and Implementing List Homomorphisms
1193 in Parallel Program Development. *Science of Computer Program-
1194 ming* 33, 1 (1999), 1 – 27.
- 1195 [38] Sergei Gorlatch. 2004. Send-receive Considered Harmful: Myths and
1196 Realities of Message Passing. *ACM Trans. Program. Lang. Syst.* 26, 1
1197 (Jan. 2004), 47–56.
- 1198 [39] Sergei Gorlatch and Christian Lengauer. 1995. Parallelization of Divide-
1199 and-Conquer in the Bird-Meertens Formalism. *Formal Aspects of
1200 Computing* 7, 6 (1995), 663–682.
- 1201 [40] Bastian Hagedorn, Larisa Stoltzfus, Michel Steuwer, Sergei Gorlatch,
1202 and Christophe Dubach. 2018. High Performance Stencil Code Genera-
1203 tion with Lift. In *Proc. of the 2018 Intl. Symp. on Code Generation and
1204 Optimization (CGO 2018)*. ACM, New York, NY, USA, 100–112.
- 1205 [41] Kevin Hammond, Marco Aldinucci, Christopher Brown, Francesco
1206 Cesarini, Marco Danelutto, Horacio González-Vélez, Peter Kilpatrick,
1207 Rainer Keller, Michael Rossbory, and Gilad Shainer. 2013. *The Para-
1208 phrase Project: Parallel Patterns for Adaptive Heterogeneous Multicore
1209 Systems*. Springer, 218–236.
- 1210 [42] Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2008. Multiparty
1211 Asynchronous Session Types. In *Proc. of 35th Symp. on Princ. of Prog.
1212 Lang. (POPL '08)*. ACM, New York, NY, USA, 273–284.
- 1213 [43] Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2016. Multiparty
1214 Asynchronous Session Types. *J. ACM* 63, 1 (2016), 9:1–9:67.
- 1215 [44] Zhenjiang Hu, Hideya Iwasaki, and Masato Takeichi. 1996. Deriving
1216 Structural Hylomorphisms from Recursive Definitions. In *Proc. ICFP
1217 '96: ACM Int. Conf. on Functional Programming (ICFP '96)*. ACM, New
1218 York, NY, USA, 73–82.
- 1219 [45] John Hughes. 2000. Generalising monads to arrows. *Science of Com-
1220 puter Programming* 37, 1 (2000), 67 – 111.
- 1221 [46] Guy L. Steele Jr. 2005. Parallel Programming and Parallel Abstractions
1222 in Fortress. In *14th International Conference on Parallel Architecture and
1223 Compilation Techniques (PACT 2005), 17-21 September 2005, St. Louis,
1224 MO, USA*. IEEE Computer Society, 157.
- 1225 [47] Daniel J. Lehmann and Michael B. Smyth. 1981. Algebraic Specification
1226 of Data Types: A Synthetic Approach. *Mathematical Systems Theory*
1227 12, 1 (1981), 1–18.

1211 14, 1 (01 Dec 1981), 97–139.

1212 [48] Mario Leyton and José M. Piquer. 2010. Skandium: Multi-core Program-
1213 ming with Algorithmic Skeletons. In *Proceedings of the 18th Euromicro*
1214 *Conference on Parallel, Distributed and Network-based Processing, PDP*
1215 *2010, Pisa, Italy, February 17-19, 2010*. IEEE Computer Society, 289–296.

1216 [49] Yu Liu, Zhenjiang Hu, and Kiminori Matsuzaki. 2011. Towards Sys-
1217 tematic Parallel Programming over Mapreduce. In *Proc. Euro-Par 2011:*
1218 *European Conference on Parallelism*. 39–50.

1219 [50] Hugo A. López, Eduardo R. B. Marques, Francisco Martins, Nicholas
1220 Ng, César Santos, Vasco Thudichum Vasconcelos, and Nobuko Yoshida.
1221 2015. Protocol-Based Verification of Message-Passing Parallel Pro-
1222 grams. In *Proc. of the 2015 Intl. Conf. on Object-Oriented Programming,*
1223 *Systems, Languages, and Applications (OOPSLA'15)*. ACM, 280–298.

1224 [51] Frédéric Loulergue, Wadoud Bousdira, and Julien Tesson. 2017. Calcu-
1225 lating Parallel Programs in Coq Using List Homomorphisms. *International*
1226 *Journal of Parallel Programming* 45, 2 (01 Apr 2017), 300–319.

1227 [52] Michael McCool, Arch Robison, and James Reinders. 2012. *Structured*
1228 *parallel programming: patterns for efficient computation*. Elsevier.

1229 [53] Erik Meijer, Maarten Fokkinga, and Ross Paterson. 1991. Functional
1230 Programming with Bananas, Lenses, Envelopes and Barbed Wire.
1231 In *Functional Programming Languages and Computer Architecture*.
1232 Springer, 124–144.

1233 [54] Erik Meijer and Graham Hutton. 1995. Bananas in Space: Extending
1234 Fold and Unfold to Exponential Types. In *Proceedings of the Seventh*
1235 *International Conference on Functional Programming Languages and*
1236 *Computer Architecture (FPCA '95)*. ACM, New York, NY, USA, 324–333.

1237 [55] Akimasa Morihata. 2013. A Short Cut to Parallelization Theorems. In
1238 *Proc. ICFP 2013: 18th Int. Conf. on Functional Programming*. 245–256.

1239 [56] Akimasa Morihata and Kiminori Matsuzaki. 2010. Automatic Paral-
1240 lelization of Recursive Functions using Quantifier Elimination. In *Proc.*
1241 *FLOPS '10: Functional and Logic Programming*. 321–336.

1242 [57] Dimitris Mostrous and Nobuko Yoshida. 2009. Session-Based Commu-
1243 nication Optimisation for Higher-Order Mobile Processes. In *Proc. of*
1244 *the 9th Intl. Conf on Typed Lambda Calculi and Applications (LNCS)*,
1245 Vol. 5608. Springer, 203–218.

1246 [58] Matthias Neubauer and Peter Thiemann. 2004. An Implementation
1247 of Session Types. In *Practical Aspects of Declarative Languages, 6th*
1248 *International Symposium, PADL 2004, Dallas, TX, USA, June 18-19, 2004,*
1249 *Proceedings*. 56–70.

1250 [59] Nicholas Ng, José Gabriel de Figueiredo Coutinho, and Nobuko Yoshida.
1251 2015. Protocols by Default - Safe MPI Code Generation Based on
1252 Session Types. In *Proc. of the 24th Intl. Conf. on Compiler Construction*
1253 *(LNCS)*, Vol. 9031. Springer, 212–232.

1254 [60] Nicholas Ng and Nobuko Yoshida. 2015. Pabble: parameterised Scribble.
1255 *Service Oriented Computing and Applications* 9, 3-4 (2015), 269–284.

1256 [61] Ross Paterson. 2018. arrows: Arrow classes and transformers. <http://hackage.haskell.org/package/arrows-0.4.4.2>.

1257 [62] Fethi A. Rabhi and Sergei Gorbach (Eds.). 2003. *Patterns and Skeletons*
1258 *for Parallel and Distributed Computing*. Springer.

1259 [63] Jonathan Ragan-Kelley, Connelly Barnes, Andrew Adams, Sylvain
1260 Paris, Frédo Durand, and Saman P. Amarasinghe. 2013. Halide: a
1261 Language and Compiler for Optimizing Parallelism, Locality, and
1262 Recomputation in Image Processing Pipelines. In *Proc. of Conf. on*
1263 *Programming Language Design and Implementation, PLDI '13*. ACM,
1264 519–530.

1265 [64] James Reinders. 2007. *Intel threading building blocks - outfitting C++*
1266 *for multi-core processor parallelism*. O'Reilly.

1267 [65] Rodrigo C. O. Rocha, Luís Fabricio Wanderley Góes, and Fernando
1268 Magno Quintão Pereira. 2016. An Algebraic Framework for Paral-
1269 lelizing Recurrence in Functional Programming. In *Proc. of the 20th*
1270 *Brazilian Symposium on Programming Languages, SBLP 2016 (LNCS)*,
1271 Vol. 9889. Springer, 140–155.

1272 [66] David B Skillicorn. 1993. *The Bird-Meertens Formalism as a Parallel*
1273 *Model*. Springer.

1266 [67] Michel Steuwer, Christian Fensch, Sam Lindley, and Christophe
1267 Dubach. 2015. Generating Performance Portable Code Using Rewrite
1268 Rules. In *Proc ICFP 2015: 20th ACM Conf. on Functional Prog. Lang. and*
1269 *Comp. Arch.* 205–217.

1270 [68] Michel Steuwer, Toomas Rempel, and Christophe Dubach. 2017.
1271 Lift: a functional data-parallel IR for high-performance GPU code
1272 generation. In *Proceedings of the 2017 International Symposium on Code*
1273 *Generation and Optimization, CGO 2017, Austin, TX, USA, February 4-8,*
1274 *2017*. ACM, 74–85.

1275 [69] Arvind K. Sujeeth, HyoukJoong Lee, Kevin J. Brown, Tiark Rompf,
1276 Hassan Chafi, Michael Wu, Anand R. Atreya, Martin Odersky, and
1277 Kunle Olukotun. 2011. OptiML: An Implicitly Parallel Domain-Specific
1278 Language for Machine Learning. In *Proc. of the 28th Intl. Conf. on*
1279 *Machine Learning (ICML'11)*. Omnipress, 609–616.

1280 [70] Shuhao Zhang. 2019. *Session Arrows: A Session-Type Based Framework*
1281 *For Parallel Code Generation*. Master's thesis. Imperial College London.

A Further Definitions

A.1 Algebraic Functional Language

$$\begin{aligned}
 F_1, F_2 &::= \mid \mid Ka \mid F_1 + F_2 \mid F_1 \times F_2 \\
 a, b &::= \mid \mid \text{int} \mid \dots \mid a \rightarrow b \mid a + b \mid a \times b \mid F a \mid \mu F \\
 e_1, e_2 &::= f \mid v \mid \text{const } e \mid \text{id} \mid e_1 \circ e_2 \mid \pi_i \mid e_1 \Delta e_2 \mid \iota_i \mid e_1 \nabla e_2 \\
 &\mid F e \mid \text{in}_F \mid \text{out}_F \mid \text{rec}_F e_1 e_2 \\
 \\
 &\frac{f : a \rightarrow b \in \Gamma}{\vdash f : a \rightarrow b} \quad \frac{\vdash e : a}{\vdash \text{const } e : b \rightarrow a} \quad \frac{}{\vdash \text{id} : a \rightarrow a} \\
 \\
 &\frac{}{\vdash \text{in}_F : F \mu F \rightarrow \mu F} \quad \frac{}{\vdash \text{out}_F : \mu F \rightarrow F \mu F} \\
 \\
 &\frac{\vdash e_1 : b \rightarrow c \quad \vdash e_2 : a \rightarrow b}{\vdash e_1 \circ e_2 : a \rightarrow c} \quad \frac{i \in [1, 2]}{\vdash \pi_i : a_1 \times a_2 \rightarrow a_i} \\
 \\
 &\frac{\vdash e_1 : a \rightarrow b \quad \vdash e_2 : a \rightarrow c}{\vdash e_1 \Delta e_2 : a \rightarrow b \times c} \quad \frac{i \in [1, 2]}{\vdash \iota_i : a_i \rightarrow a_1 + a_2} \\
 \\
 &\frac{\vdash e_1 : a \rightarrow c \quad \vdash e_2 : b \rightarrow c}{\vdash e_1 \nabla e_2 : a + b \rightarrow c} \quad \frac{\vdash e : a \rightarrow b}{\vdash F e : F a \rightarrow F b} \\
 \\
 &\frac{\vdash e_1 : F b \rightarrow b \quad \vdash e_2 : a \rightarrow F a}{\vdash \text{rec}_F e_1 e_2 : a \rightarrow b}
 \end{aligned}$$

Figure 7. Syntax and types of Alg.

A.1.1 Properties of Alg Constructs

Alg constructs are characterised by well-known properties. Sum and product functions, are uniquely determined by their universal properties. Composition and identity must satisfy the associativity and cancellation properties. These basic properties are summarised in Fig. 9. Functors preserve identities and composition, and rec satisfy the *hylomorphism laws* (Fig. 9b). The laws of hylomorphisms can be used to perform some common program optimisations. For example, the well-known *deforestation* transformation can be

	Constant, Identity and Composition
1321	$\text{const } e = \lambda x. e \quad \text{id} = \lambda x. x \quad e_1 \circ e_2 = \lambda x. e_1 (e_2 x)$
1322	Products
1323	$\pi_i = \lambda(x_1, x_2). x_i, i \in [1, 2] \quad e_1 \Delta e_2 = \lambda x. (e_1 x, e_2 x)$
1324	$e_1 \times e_2 = (e_1 \circ \pi_1) \Delta (e_2 \circ \pi_2)$
1325	Coproducts
1326	$\iota_i = \lambda x. \text{inj}_i x \quad e_1 \nabla e_2 = \lambda(\text{inj}_i x). e_i x$
1327	$e_1 + e_2 = (\iota_1 \circ e_1) \nabla (\iota_2 \circ e_2)$
1328	Functors
1329	$\vdash a = a \quad \text{Ka } b = a \quad (F_1 \dagger F_2) a = F_1 a \dagger F_2 a, \dagger \in \{+, \times\}$
1330	$\vdash e = e \quad \text{Ka } e = \text{id} \quad (F_1 \dagger F_2) e = F_1 e \dagger F_2 e$
1331	Recursion
1332	$\text{in}_F = \lambda x. \text{in}_F x \quad \text{out}_F = \lambda(\text{in}_F x). x$
1333	$\text{rec}_F e_1 e_2 = f \quad \text{where } f = e_1 \circ F f \circ e_2$
1334	

Figure 8. Semantics of Alg combinators.

derived from the hylomorphism equation, and the properties of functors. Particularly, it is an instance of Equations A.11 and A.7 in Fig. 9. From the universal properties of Δ and ∇ (Fig. 9a), a number of equivalences can be derived, e.g.: $\pi_i \circ e_1 \Delta e_2 = e_i$; $(\pi_1 \circ e) \Delta (\pi_2 \circ e) = e$; $\pi_1 \Delta \pi_2 = \text{id}$; $(e_1 \times e_2) \circ (e_3 \Delta e_4) = (e_1 \circ e_3) \Delta (e_2 \circ e_4)$; $e_1 \nabla e_2 \circ \iota_i = e_i$; $(e \circ \iota_1) \nabla (e \circ \iota_2) = e \iota_1 \nabla \iota_2 = \text{id}$; and $(e_1 \nabla e_2) \circ (e_3 + e_4) = (e_1 \circ e_3) \nabla (e_2 \circ e_4)$ where $i \in \{1, 2\}$. The properties of combinators provide a formal framework for equational reasoning that can be used as a basis for doing program transformations [31; 34; 53]. These properties have been used for parallelising functions, e.g. [17; 33; 55]. In this paper, we use $=_{\text{ext}}$ for the equations in 9, to distinguish them from the syntactic equality ($=$).

A.2 Parallel Algebraic Language

A.2.1 Typing rules

Rules **ID**, **COMP**, **PROJ_i** and **SPLIT** are standard. The main feature of this type system is the use of eta-expanded sum-types and unions of interfaces to deal with choices. Rule **CHOICE** specifies that a choice point may be introduced at any point when a participant contains a value of a sum-type. In such cases p sends the tag of the sum-type value to any other participant whose behaviour depends on it.

After the choice point, the interface is $I[\iota_1 p] \cup^{\vec{p}} I[\iota_2 p]$, with the constraint that the participants in $I[p]$ must be in \vec{p} . Rule **ALT** specifies that e must be the parallelisation of e , considering both A_1 and A_2 as input interfaces. The output interface is the union of B_1 and B_2 . Any participant in e must be notified of the choice $\text{pids}(e) \subseteq \vec{p}$, to make sure that they perform the interactions that correspond to the correct A_i . Rule **JOIN** is the same as rule **ALT**, but we do not require the participants in e to be notified of the choice, since the input interface is the same in both branches of the choice. Rule **INJ_i** is used to *tag* an interface with the i -th injection. Then, rule **CASE_i** specifies that if e_i is the parallelisation of e_i , then $e_1 \nabla e_2$ is a parallelisation of $e_1 \nabla e_2$, given the tagged

input interface $\iota_i I_a$. Note that e_j with $j \neq i$ is free in rule **CASE_i**. We solve these cases by unification (see rule **ALT**). If this is not possible, this means that no branch in the control flow will ever reach e_j , and so we can set it to any arbitrary parallelisation of e_j , or even optimise e_j away. Rule **ALG** specifies that given any e and participant p , $e@p$ is a valid parallelisation, with output interface $b@p$. Finally, rule **EXT** is crucial for exploring potential parallelisations. It states that if e is the parallelisation of e_2 , and e_2 is extensionally equal to e_1 , then e is also a parallelisation of e_1 . The undecidability of this rule requires that the programmer specifies *rewriting strategies* both for checking and inference.

Example A.1 (Rewriting and annotation strategies). We illustrate how rewriting and annotation strategies work by showing the mergesort (**ms**) example. Consider the mergesort definition $\text{ms} = \text{rec}_T \text{mrg spl}$. Solutions to the inference problem $\vdash \text{ms} \Rightarrow ?0 : \text{Ls}@p_0 \rightarrow ?1$ provide the alternative parallelisations of **ms**. The only two rules that can be applied are **ALG** or **EXT**. By rule **ALG**, we can annotate **ms** at some p_1 : $\vdash \text{ms} \Rightarrow \text{ms}@p_1 : \text{Ls}@p_0 \rightarrow \text{Ls}@p_1$. Alternatively, we can use the hylomorphism equation, and apply rule **EXT**:

$$\text{ms} = \text{rec}_T \text{mrg spl} = \text{mrg} \circ T \text{ms} \circ \text{spl} = \text{mrg} \circ (\text{id} + \text{ms} \times \text{ms}) \circ \text{spl}$$

We decide which of the rules to apply by querying a collection of rewriting hints, that we call *rewriting strategy*. This collection of hints is of the form $[e_1 : \text{rw}_1, \dots]$, and must be specified by the programmer. The rewritings rw_i are essentially proofs that $e_i =_{\text{ext}} e'_i$, by applying equations in Fig. 9. Once a hint is used, it is removed from the collection of hints. For the mergesort example, if we use the rewriting strategy $[\text{ms} : \text{unroll } 1]$, we will apply rule **EXT**, unroll the hylomorphism equation once, and continue with an empty strategy $[]$.

$$\frac{\vdash \text{mrg} \circ (\text{id} + \text{ms} \times \text{ms}) \circ \text{spl} \Rightarrow ?0 : \text{Ls}@p_0 \rightarrow ?1 \quad \text{ms} =_{\text{ext}} \text{mrg} \circ (\text{id} + \text{ms} \times \text{ms}) \circ \text{spl}}{\vdash \text{ms} \Rightarrow ?0 : \text{Ls}@p_0 \rightarrow ?1}$$

With an empty rewriting strategy, the only possibility once we find the atomic function **spl** is to use rule **ALG**. To select a participant, we query the *annotation strategy*. The annotation strategy is a collection of expressions that we require to place at distinct participants. Suppose that our annotation strategy is $\{\text{spl}\}$. Then, we would need to select a fresh participant p_1 : $\vdash \text{spl} \Rightarrow \text{spl}@p_1 : \text{Ls}@p_0 \rightarrow ((1 + a) + \text{Ls} \times \text{Ls})@p_1$. If the annotation strategy does not contain **spl**, then we would select any participant from the input interface, to minimise the amount of messages exchanged: $\vdash \text{spl} \Rightarrow \text{spl}@p_0 : \text{Ls}@p_0 \rightarrow ((1 + a) + \text{Ls} \times \text{Ls})@p_0$. Suppose that the annotation strategy is $\{\text{spl}\}$. Then, after **spl** we have a sum type at p_1 . This requires us to introduce a choice point:

$$\begin{array}{ll}
\text{id} \circ e = e \circ \text{id} = e & \text{(A.1)} \\
F(e_1 \circ e_2) = F e_1 \circ F e_2 & \text{(A.2)} \\
F \text{id} = \text{id} & \text{(A.3)} \\
e_1 \circ (e_2 \circ e_3) = (e_1 \circ e_2) \circ e_3 & \text{(A.4)} \\
f = e_1 \Delta e_2 \Leftrightarrow \pi_1 \circ e = e_1 \wedge \pi_2 \circ e = e_2 & \text{(A.5)} \\
e = e_1 \nabla e_2 \Leftrightarrow e \circ \iota_1 = e_1 \wedge e \circ \iota_2 = e_2 & \text{(A.6)}
\end{array}$$

(a) Basic Properties of Combinators

$$\begin{array}{ll}
e_3 \circ e_4 = \text{id} \Rightarrow (\text{rec}_F e_1 e_3) \circ (\text{rec}_F e_4 e_2) = \text{rec}_F e_1 e_2 & \text{(A.7)} \\
\eta : F_1 \rightarrow F_2 \Rightarrow \text{rec}_{F_1}(e_1 \circ \eta) e_2 = \text{rec}_{F_2} e_1 (\eta \circ e_2) & \text{(A.8)} \\
e_1 \text{ strict} \wedge e_1 \circ e_3 = e_5 \circ F e_1 \wedge e_4 \circ e_2 = F e_2 \circ e_6 \Rightarrow e_1 \circ (\text{rec}_F e_3 e_4) \circ e_2 = \text{rec}_F e_5 e_6 & \text{(A.11)}
\end{array}$$

(b) Hylomorphism Laws

$$\begin{array}{ll}
\text{rec}_F \text{ in out} = \text{id}_{\mu F} & \text{(A.9)} \\
e_1, e_2 \text{ strict} \Rightarrow \text{rec}_F e_1 e_2 \text{ strict} & \text{(A.10)}
\end{array}$$

Figure 9. Properties of point-free combinators

$$\frac{\vdash \text{mrg} \circ (\text{id} + \text{ms} \times \text{ms}) \Rightarrow ?0 \quad \{p_1\} \subseteq ?3 \quad : ((1 + a) + \text{Ls} \times \text{Ls})@(\iota_1 p_1 \cup^{?3} \iota_2 p_1) \rightarrow ?1}{\vdash \text{mrg} \circ (\text{id} + \text{ms} \times \text{ms}) \Rightarrow ?0 \circ [p_1 \oplus ?3] \quad : ((1 + a) + \text{Ls} \times \text{Ls})@p_1 \rightarrow ?1 \cup^{?3} ?2}$$

By collecting all constraints of the form $\{p_1\} \subseteq ?3$, we can fully determine what is the minimum list of participants $?3$ that we require. To conclude our example, we show the end result of applying the rest of the rules. The final structure follows a divide-and-conquer parallel structure, that may bypass p_2 and p_3 if the input list is empty or singleton.

$$\begin{array}{l}
\vdash \text{mrg} \circ (\text{id} + \text{ms} \times \text{ms}) \\
\Rightarrow \text{mrg}@p_1 \circ (\text{id} + (\text{ms}@p_2 \circ \pi_1 @p_1) \Delta (\text{ms}@p_3 \circ \pi_2 @p_1)) \\
\quad \circ [p_1 \oplus p_1 p_2 p_3] \circ \text{spl}@p_1 \\
: \text{Ls}@p_0 \rightarrow \text{Ls}@p_1 \cup^{p_1 p_2 p_3} \text{Ls}@p_1
\end{array}$$

Definition A.2 (Product and Injection of Choice Interfaces). We sometimes write $A \times B$ to represent the product of interfaces that contain choices. We do the product of the respective interfaces, after performing first the choices in A , and then the choices in B :

$$\begin{array}{l}
(R_1 \cup^{\bar{p}} R_2) \times R_3 = (R_1 \times R_3) \cup^{\bar{p}} (R_2 \times R_3) \\
I_1 \times (R_1 \cup^{\bar{p}} R_2) = (I_1 \times R_1) \cup^{\bar{p}} (I_1 \times R_2)
\end{array}$$

We also write injections of interfaces that contain choices:

$$\iota_i (R_1 \cup^{\bar{p}} R_2) = \iota_i R_1 \cup^{\bar{p}} \iota_i R_2$$

Definition A.3 (Well-formedness of interfaces: $\text{WF}(a@R)$). Interface $a@R$ is well formed if R matches the structure of a :

$$\frac{\text{WF}(a_i @ I)}{\text{WF}(a @ p)} \quad \frac{\text{WF}(a_i @ I)}{\text{WF}((a_1 + a_2) @ (\iota_i I))} \quad \frac{\text{WF}(a @ I_1) \quad \text{WF}(b @ I_2)}{\text{WF}((a \times b) @ (I_1 \times I_2))}$$

$$\frac{\text{WF}(a @ R_1) \quad \text{WF}(a @ R_2)}{\text{WF}(a @ (R_1 \cup^{\bar{p}} R_2))}$$

For well-formed interfaces, we sometimes propagate the annotation down the type structure, e.g. $a @ R_1 \times b @ R_2$ is notation for $(a \times b) @ (R_1 \times R_2)$. We also define sums of interfaces $a @ R_1 +^{\bar{p}} b @ R_2$ as notation for $(a + b) @ (\iota_1 R_1 \cup^{\bar{p}} \iota_2 R_2)$.

$$\begin{array}{c}
\text{JOIN} \\
\frac{\vdash e \Rightarrow e : A \rightarrow B}{\vdash e \Rightarrow e : A \cup^{\bar{p}} A \rightarrow B \cup^{\bar{p}} B} \\
\\
\text{ALT} \\
\frac{\vdash e \Rightarrow e : A_1 \rightarrow B_1 \quad \vdash e \Rightarrow e : A_2 \rightarrow B_2 \quad A_1 \neq A_2 \quad \text{pids}(e) \subseteq \bar{r}}{\vdash e \Rightarrow e : A_1 \cup^{\bar{p}} A_2 \rightarrow B_1 \cup^{\bar{p}} B_2} \\
\\
\text{ALG} \quad \text{EXT} \\
\frac{\vdash e : a \rightarrow b}{\vdash e \Rightarrow e @ p : a @ I \rightarrow b @ p} \quad \frac{\vdash e_2 \Rightarrow e : a @ I \rightarrow B \quad e_1 =_{\text{ext}} e_2}{\vdash e_1 \Rightarrow e : a @ I \rightarrow B} \\
\\
\text{INJ}_i \quad \text{ID} \\
\frac{}{\vdash \iota_i \Rightarrow \iota_i : a @ I \rightarrow a @ (\iota_i I)} \quad \frac{}{\vdash \text{id} \Rightarrow \text{id} : A \rightarrow A} \\
\\
\text{PROJ}_i \\
\frac{i \in [1, 2]}{\vdash \pi_i \Rightarrow \pi_i : (a_1 \times a_2) @ (I_1 \times I_2) \rightarrow a_i @ I_i} \\
\\
\text{COMP} \\
\frac{\vdash e_1 \Rightarrow e_1 : B \rightarrow C \quad \vdash e_2 \Rightarrow e_2 : A \rightarrow B}{\vdash e_1 \circ e_2 \Rightarrow e_1 \circ e_2 : A \rightarrow C} \\
\\
\text{CASE}_i \\
\frac{\vdash e_i \Rightarrow e_i : a_i @ I \rightarrow B}{\vdash e_1 \nabla e_2 \Rightarrow e_1 \nabla e_2 : (a_1 + a_2) @ (\iota_i I) \rightarrow B} \\
\\
\text{SPLIT} \\
\frac{\vdash e_1 \Rightarrow e_1 : a @ I \rightarrow B \quad \vdash e_2 \Rightarrow e_2 : a @ I \rightarrow C}{\vdash e_1 \Delta e_2 \Rightarrow e_1 \Delta e_2 : a @ I \rightarrow B \times C} \\
\\
\text{CHOICE} \\
\frac{\vdash e \Rightarrow e : a @ (I[\iota_1 p] \cup^{\bar{p}} I[\iota_2 p]) \rightarrow B \quad \text{pids}(I[p]) \subseteq \bar{p} \quad \text{tyAt}(I, a) = b + c}{\vdash e \Rightarrow e \circ [p \oplus \bar{p}] : a @ I[p] \rightarrow B}
\end{array}$$

Figure 10. Typing rules of PALg

Definition A.4 (Projection Rules ($G \upharpoonright p$) and Merging ($L_1 \sqcap L_2$)). Projection defines how to obtain the local type L of a participant p in a global type G : $G \upharpoonright p$.

$$\begin{array}{l}
(p_1 \rightarrow p_2 : a. G) \upharpoonright p \\
= \begin{cases} p_2!(a). (G \upharpoonright p) & \text{if } p = p_1 \neq p_2 \\ p_1?(a). (G \upharpoonright p) & \text{if } p = p_2 \neq p_1 \\ G \upharpoonright p & \text{otherwise} \end{cases}
\end{array}$$

$$\begin{array}{l}
(p_1 \rightarrow p_2 : \{l_i.G_i\}_{i \in I}) \upharpoonright p \\
= \begin{cases} p_2 \oplus \{l_i.G_i \upharpoonright p\} & \text{if } p = p_1 \\ p_1 \& \{l_i.G_i \upharpoonright p\} & \text{if } p = p_2 \\ \prod_{i \in I} (G_i \upharpoonright p) & \text{otherwise} \end{cases}
\end{array}$$

$$\left(\mu X. (G \upharpoonright p) \right) \text{ if } G \upharpoonright p \neq X'. \forall X'$$

1541 $p \& \{l_i.L_i\}_{i \in I} \sqcap p \& \{l_j.L'_j\}_{j \in J}$
 1542 $= p \& \{l_k.L_k \sqcap L'_k\}_{k \in I \cap J} \cup \{l_i.L_i\}_{i \in I \setminus J} \cup \{l_m.L_m\}_{m \in J \setminus I}$
 1543 $\mu X.L_1 \sqcap \mu X.L_2 = \mu X.(L_1 \sqcap L_2) \quad L \sqcap L = L$
 1544

1545 Projection onto a role is not necessarily defined. Particularly,
 1546 projecting $p_1 \rightarrow p_2 : \{l_i.G_i\}$ onto p , with $p \neq p_1$ and $p \neq p_2$,
 1547 is only defined if the projection of all G_i onto p can be *merged*.
 1548 Two local types can be merged only if they are the same, or
 1549 if they branch on the same role p , and their continuations
 1550 can be merged. For example, p_3 's local type of the global
 1551 type: $\mu X.p_1 \rightarrow p_2 \{l_1.p_2 \rightarrow p_3 : l_2.p_1 \rightarrow p_3 : l_3.X, l_4.p_2 \rightarrow$
 1552 $p_3 : l_5.p_1 \rightarrow p_3 : l_6.\text{end}\}$ is $\mu X.p_2 \& \{l_2.p_1 \& \{l_3.X\}, l_4.p_2 \&$
 1553 $\{l_5.p_1 \& \{l_6.\text{end}\}\}$.

1554 **Definition A.5.** LTS for Local Type Configurations

1555 $\langle C, Q \rangle \rightsquigarrow^\ell \langle C', Q' \rangle \quad C = [p_i \mapsto L_i]_{i \in I} \quad Q = [p_i p_j \mapsto w]_{i \in I, j \in I}$
 1556
 1557 $\langle C[p_0 \mapsto p_1!(a).L], Q[p_0 p_1 \mapsto w] \rangle$
 1558 $\rightarrow_{p_0 p_1!(a)} \langle C[p_0 \mapsto L], Q[p_0 p_1 \mapsto a \cdot w] \rangle$
 1559 $\langle C[p_0 \mapsto p_1?(a).L], Q[p_1 p_0 \mapsto w \cdot a] \rangle$
 1560 $\rightarrow_{p_0 p_1?(a)} \langle C[p_0 \mapsto L], Q[p_1 p_0 \mapsto w] \rangle$
 1561 $\langle C[p_0 \mapsto p_1 \oplus \{l_i.L_i\}_{i \in I}], Q[p_0 p_1 \mapsto w] \rangle$
 1562 $\rightarrow_{p_0 p_1 \oplus l_i} \langle C[p_0 \mapsto L_i], Q[p_0 p_1 \mapsto l_i \cdot w] \rangle$
 1563 $\langle C[p_0 \mapsto p_1 \& \{l_i.L_i\}_{i \in I}], Q[p_1 p_0 \mapsto w \cdot l_i] \rangle$
 1564 $\rightarrow_{p_0 p_1 \& l_i} \langle C[p_0 \mapsto L_i], Q[p_1 p_0 \mapsto w] \rangle$
 1565
 1566

1567 **Definition A.6** (Interface projection: $A \upharpoonright p$). The projection
 1568 of interface A onto role p is the part of interface A that is
 1569 located at p . We define the projection for $a \in R$ inductively on
 1570 the structure of R :

1571 $a \upharpoonright p_0 \upharpoonright p_1 = \{a, \text{ if } p_0 = p_1; 1, \text{ otherwise}\}$
 1572 $(a \times b) \upharpoonright (R_1 \times R_2) \upharpoonright p = (a \upharpoonright R_1 \upharpoonright p) \times (b \upharpoonright R_2 \upharpoonright p)$
 1573 $(a_1 + a_2) \upharpoonright (l_i R) \upharpoonright p = (a_i \upharpoonright R \upharpoonright p)$
 1574 $a \upharpoonright (R_1 \cup \bar{p} R_2) \upharpoonright p =$
 1575 $\begin{cases} (a \upharpoonright R_1) \upharpoonright p \uplus (a \upharpoonright R_2) \upharpoonright p & \text{if } p \in \bar{p} \\ a' & \text{if } a' = (a \upharpoonright R_1) \upharpoonright p = (a \upharpoonright R_2) \upharpoonright p \end{cases}$
 1576
 1577

1578 A.3 MPST

1579 **Definition A.7** (Label Broadcasting). We define a macro
 1580 that represents the broadcasting of a label to multiple partic-
 1581 ipants in a choice. We write

1582 $\bullet p \rightarrow \{p_j\}_{j \in [1, n]} : \{l_i.G_i\}_{i \in I}$ for $p \rightarrow p_1 \{l_i.p \rightarrow p_2 \{l_i$
 1583 $\dots p \rightarrow p_n \{l_i.G_i\} \dots\}_{i \in I}$; and
 1584 $\bullet \{p_j\}_{j \in [1, n]} \oplus \{l_i.L_i\}_{i \in I}$ for $p_1 \oplus \{l_i.p_2 \oplus \{l_i \dots p_n \oplus$
 1585 $\{l_i.L_i\} \dots\}_{i \in I}$.

1586 It is straightforward to show that $(p_1 \rightarrow \{p_j\}_{j \in J} : \{l_i.G_i\}_{i \in I}) \upharpoonright$
 1587 $p_2 = \{p_j\}_{j \in J} \oplus \{l_i.G_i \upharpoonright p_2\}_{i \in I}$, if $p_1 = p_2$, and $(p_1 \rightarrow \{p_j\}_{j \in J} :$
 1588 $\{l_i.G_i\}_{i \in I}) \upharpoonright p_2 = p_1 \& \{l_i.G_i \upharpoonright p_2\}_{i \in I}$, if $p_2 = p_j$ for some
 1589 $j \in J$.

1591 The relation $\models p \Leftarrow A \sim G$ associates p and A with the
 1592 global type G (Fig. 11).

1593 Rules **ID**, **INJ_i**, **PROJ_i**, and **CASE_i** are straightforward. Rule
 1594 **COMP** associates two PALg expressions with the *sequencing*

1595

1596 of their respective global types, $G_1 \wp G_2$. The sequencing
 1597 produces the global type that results of performing first G_1 ,
 1598 and then G_2 , by taking into account branching and choices:

1599 $\text{end} \wp G = G$
 1600 $(p_1 \rightarrow p_2 : a.G_1) \wp G_2 = p_1 \rightarrow p_2 : a.(G_1 \wp G_2)$
 1601 $(p_1 \rightarrow p_2 \{l_i.G_i\}) \wp G_2 = p_1 \rightarrow p_2 \{l_i.G_i \wp G_2\}$
 1602 $(p_1 \rightarrow p_2 \begin{Bmatrix} l_1.G_1 \\ l_2.G_2 \end{Bmatrix}) \wp (G_3 \cup G_4) = p_1 \rightarrow p_2 \begin{Bmatrix} l_1.G_1 \wp G_3 \\ l_2.G_2 \wp G_4 \end{Bmatrix}$
 1603 $(G_1 \cup G_2) \wp (G_3 \cup G_4) = (G_1 \wp G_3) \cup (G_2 \wp G_4)$
 1604
 1605

1606 Rule **CHOICE** turns a choice point of p with dependencies
 1607 \bar{p} into a global type choice: p notifies all participants in \bar{p}
 1608 of the branch in the protocol that they need to take. Rule
 1609 **ALT** associates p with two protocols, G_1 and G_2 , whenever
 1610 the input interface is a choice of A_1 and A_2 . Each G_i is the
 1611 continuation that corresponds to the i -th branch of the choice
 1612 that led to interface A_i . Note that contrary to the typing
 1613 rules of Fig. 2, there is no rule **JOIN**. This is because **JOIN**
 1614 was only used to avoid adding too many participants to a
 1615 choice. However, at this point, these participants are known,
 1616 and specified at the choice points. Rule **SPLIT** associates a
 1617 split of PALg expressions with the sequence of the respective
 1618 interactions. The rule uses $[G_2/\text{end}]G_1$ instead of \wp , because
 1619 if G_1 contains a global type choice, then the interactions
 1620 described by G_2 must be done after every branch in G_1 . Since
 1621 both G_1 and G_2 start from the same input interface, any
 1622 choice in either G_i must be independent of any choice in the
 1623 other G_j . Finally, rule **ALG** specifies that if the input interface
 1624 is $a \in I$, and the expression is $e \oplus p$, then the protocol comprises
 1625 the sequence of interactions from all participants in I to
 1626 participant p .

1627 **Example A.8** (Mergesort Protocol). Recall the PALg expres-
 1628 sion inferred for **ms** in Example A.1:

1629 $\vdash \text{mrg} \circ (\text{id} + \text{ms} \times \text{ms})$
 1630 $\Rightarrow \text{mrg} \oplus p_1 \circ (\text{id} + (\text{ms} \oplus p_2 \circ \pi_1 \oplus p_1) \Delta (\text{ms} \oplus p_3 \circ \pi_2 \oplus p_1))$
 1631 $\circ [p_1 \oplus p_1 p_2 p_3] \circ \text{spl} \oplus p_1$
 1632 $: \text{Ls} \oplus p_0 \rightarrow \text{Ls} \oplus p_1 \cup^{p_1 p_2 p_3} \text{Ls} \oplus p_1$
 1633
 1634

1635 We need to solve $\models \text{mrg} \oplus p_1 \circ (\text{id} + (\text{ms} \oplus p_2 \circ \pi_1 \oplus p_1) \Delta (\text{ms} \oplus p_3 \circ$
 1636 $\pi_2 \oplus p_1)) \circ [p_1 \oplus p_1 p_2 p_3] \circ \text{spl} \oplus p_1 \Leftarrow \text{Ls} \oplus p_0 \sim \text{?0}$.

1637 The first step is a straightforward application of **COMP**. The
 1638 case for $[p_1 \oplus p_1 p_2 p_3]$ is the result of applying rule **CHOICE**.
 1639 To help readability, we use different colors for the left and
 1640 the right branches:

1641 $\models [p_1 \oplus p_1 p_2 p_3]$
 1642 $\Leftarrow ((1 + a) + \text{Ls} \times \text{Ls}) \oplus p_1 \sim p_1 \rightarrow \{p_2 p_3\} \{l_1.\text{end}; l_2.\text{end}\}$
 1643

1644 At this point, the input interface is:

1645 $((1 + a) + \text{Ls} \times \text{Ls}) \oplus (l_1 p_1) \cup^{p_1 p_2 p_3} ((1 + a) + \text{Ls} \times \text{Ls}) \oplus (l_2 p_1)$
 1646

1647 This means that we need to obtain two sub-protocols, for
 1648 the left and the right branches respectively. The left branch is
 1649 solved by applying rule **INJ₁**, while the right branch is solved
 1650 by rules **CASE**, **INJ₂**, **SPLIT**, **COMP** and **ALG**:

$$\begin{array}{c}
\text{ID} \\
\hline
\vdash \text{id} \Leftarrow a@I \sim \text{end} \\
\text{INJ}_i \\
\hline
\vdash \iota_i \Leftarrow a_i@I \sim \text{end} \\
\text{PROJ}_i \\
\hline
\begin{array}{c}
i \in [1, 2] \\
\vdash \pi_i \Leftarrow (a_1 \times a_2)@(I_1 \times I_2) \sim \text{end}
\end{array} \\
\text{CHOICE} \\
\hline
\vdash [p \oplus \bar{p}] \Leftarrow a[b + c]@I[p] \sim p \rightarrow \{\bar{p} \setminus p\} \{t_1. \text{end}; t_2. \text{end}\} \\
\text{CASE}_i \\
\hline
\begin{array}{c}
\vdash e_i \Leftarrow a_i@I \sim G \\
\vdash e_1 \nabla e_2 \Leftarrow (a_1 + a_2)@(I_1 I) \sim G
\end{array} \\
\text{ALT} \\
\hline
\begin{array}{c}
\vdash e \Leftarrow A_1 \sim G_1 \quad \vdash e \Leftarrow A_2 \sim G_2 \\
\vdash e \Leftarrow A_1 \cup \bar{p} A_2 \sim G_1 \cup G_2
\end{array} \\
\text{ALG} \\
\hline
\begin{array}{c}
\vdash e : a \rightarrow b \\
\vdash e@p \Leftarrow a@I \sim [a@I \rightsquigarrow p]
\end{array} \\
\text{COMP} \\
\hline
\begin{array}{c}
\vdash e_2 \Leftarrow A \sim G_2 \quad \vdash e_1 \Leftarrow B \sim G_1 \quad \vdash e_2 : A \rightarrow B \\
\vdash e_1 \circ e_2 \Leftarrow A \sim G_2 \circ G_1
\end{array} \\
\text{SPLIT} \\
\hline
\begin{array}{c}
\vdash e_i \Leftarrow a@I \sim G_i \quad (i = 1, 2) \\
\vdash e_1 \Delta e_2 \Leftarrow a@I \sim [G_2/\text{end}]G_1
\end{array}
\end{array}$$


$$\begin{array}{l}
[a@p_1 \rightsquigarrow p_2] = p_1 \rightarrow p_2 : a. \text{end}, \text{ if } p_1 \neq p_2; [a@p \rightsquigarrow p] = \text{end}; \\
[(a \times b)@(I_a \times I_b) \rightsquigarrow p] = [a@I_a \rightsquigarrow p] \circ [b@I_b \rightsquigarrow p]; \text{ and} \\
[(a_1 + a_2)@(I_1 I) \rightsquigarrow p] = [a_i@I \rightsquigarrow p]
\end{array}$$

Figure 11. Protocol relation

$$\begin{array}{l}
\vdash \text{id} + (\text{ms}@p_2 \circ \pi_1@p_1) \Delta (\text{ms}@p_3 \circ \pi_2@p_1) \\
\Leftarrow ((1 + a) + \text{Ls} \times \text{Ls})@(I_1 p_1) \sim \text{end} \\
\vdash \text{id} + (\text{ms}@p_2 \circ \pi_1@p_1) \Delta (\text{ms}@p_3 \circ \pi_2@p_1) \\
\Leftarrow ((1 + a) + \text{Ls} \times \text{Ls})@(I_2 p_1) \sim p_1 \rightarrow p_2 : \text{Ls}.p_1 \rightarrow p_3 : \text{Ls}. \text{end}
\end{array}$$

The interface at this point is $((1 + a) + \text{Ls} \times \text{Ls})@(I_1 p_1 \cup^{p_1 p_2 p_3} I_2 (p_2 \times p_3))$. For the last expression, $\text{mrg}@p_1$, we produce the following protocol: $\text{end} \cup p_2 \rightarrow p_1 : \text{Ls}.p_3 \rightarrow p_1 : \text{Ls}. \text{end}$. This is because, on the left branch, the input is still at p_1 , so no communication is required. On the right branch, the input is a product of lists, one at p_2 , and another one at p_3 , and so they need to communicate with p_1 . The final protocol is obtained by applying sequencing:

$$\begin{array}{l}
p_1 \rightarrow \{p_2 p_3\} \left\{ \begin{array}{l} t_1. \text{end}; \\ t_2. p_1 \rightarrow p_2 : \text{Ls}.p_1 \rightarrow p_3 : \text{Ls}. \text{end} \end{array} \right\} \circ \\
\quad (\text{end} \cup (p_2 \rightarrow p_1 : \text{Ls}. p_3 \rightarrow p_1 : \text{Ls}. \text{end})) \\
= p_1 \rightarrow \{p_2 p_3\} \left\{ \begin{array}{l} t_1. \text{end}; \\ t_2. p_1 \rightarrow p_2 : \text{Ls}.p_1 \rightarrow p_3 : \text{Ls}. \\ \quad p_2 \rightarrow p_1 : \text{Ls}. p_3 \rightarrow p_1 : \text{Ls}. \text{end} \end{array} \right\}
\end{array}$$



A.4 Mp

Mp comprises four basic operations: send, receive, choice and branching, with a standard (asynchronous) semantics. Additionally, for composing actions that depend on the same choice, we introduce case expressions.

$$\begin{array}{l}
v ::= x \mid (v, v) \mid \iota_i v \mid \dots \mid e v \\
m_i ::= \text{ret } v \mid m \gg f \mid \text{send } p v \mid \text{recv } p a \mid \text{sel } \bar{p} v f_1 f_2 \\
\quad \mid \text{brn } p m_1 m_2 \mid \text{case } f_1 f_2 \quad f ::= \lambda x. m
\end{array}$$

Values v are either primitive values, tagged values $\iota_i v$, pairs of values, or the result of applying an Alg expression e to a value. We use standard notation for the monadic unit (**ret**) and bind (\gg). The term $\lambda x. m$ is a monadic continuation. We write $\lambda_. m$ when the continuation discards the result of the previous monadic action. We use the standard Kleisli composition: $f_1 \gg f_2 = \lambda x. f_1 x \gg f_2$.

The message-passing constructs are standard, except **sel**, **brn** and **case**, which are used for performing choices, and composing actions that depend on the same choice. We explain them in detail below. We include select and branching as syntactic constructs to simplify the typeability of parallel code against local types, but their semantics can be defined in terms of standard pattern matching, plus *send* and *receive* operations.

Each monadic computation f or m has a type $m : \text{Mp } L a$, where a is the return type of m , and L is the type index of Mp, and it represents the local type that corresponds to the behaviour of the term m . There is almost a one to one correspondence between the terms L and the monadic actions m , so we refer the reader to Appendix A (Fig. 16) for the full definition.

Composing Choices The types of the constructs that deal with choices use a new type, \uplus , that is isomorphic to sum types, but that can only be constructed and eliminated by using the following monadic constructs:

$$\begin{array}{l}
\text{sel } \bar{p} : a + b \rightarrow (a \rightarrow \text{Mp } L_1 c_1) \rightarrow (b \rightarrow \text{Mp } L_2 c_2) \\
\quad \rightarrow \text{Mp } (\bar{p} \oplus \{t_1.L_1; t_2.L_2\}) (c_1 \uplus c_2) \\
\text{brn } p : \text{Mp } L_1 a_1 \rightarrow \text{Mp } L_2 a_2 \\
\quad \rightarrow \text{Mp } (p \& \{t_1.L_1; t_2.L_2\}) (a_1 \uplus a_2) \\
\text{case} : (a \rightarrow \text{Mp } L_1 c) \rightarrow (b \rightarrow \text{Mp } L_2 d) \rightarrow a \uplus b \\
\quad \rightarrow \text{Mp } (L_1 \cup L_2) (c \uplus d)
\end{array}$$

These constructs ensure that the tag used to build $a \uplus b$ indeed corresponds to the correct branch of the right choice. We use **case** to compose actions that depend on a previous choice. It may seem that this treatment of \uplus leads to unnecessary code duplication, e.g. the only possibility to compose a single action f after a branch is using **case**: $\text{brn } p m_1 m_2 \gg \text{case } f f$. Our back-end easily optimises those cases to avoid code duplication.

By the definition of the monadic bind, when we compose a branch or select with a case expression, the final local type cannot contain \cup . To illustrate this, consider $m_1 : \text{Mp } L_1 (a \uplus b)$ and $f_2 : a \uplus b \rightarrow \text{Mp } (L_2 \cup L_3) (c \uplus d)$. The local type of $m_1 \gg f_2$ must be $L_1 \uplus (L_2 \cup L_3)$. But that is only defined if L_1 contains a branch or select. Therefore, $m_1 \gg f_2$ is only well-typed if f_2 is a case expression on the tag introduced by the topmost branch or select of m_1 .

Parallel programs We define the basic constructs of PALg in a bottom-up way by manipulating *parallel programs*. Parallel programs are mappings from participants to their monadic action: $\mathbf{E} ::= [p_i \mapsto m_i]_{i \in I}$. If $m_i : \text{Mp } L_i a_i$ for all $i \in I$, then we write $[p_i \mapsto m_i]_{i \in I} : \text{Mp } [p_i \mapsto L_i]_{i \in I} [p_i \mapsto a_i]_{i \in I}$. The semantics of both local types and monadic actions is defined in terms of such collections of actions or local types, and shared queues of values W , or queues of types Q , e.g. $\langle \mathbf{E}, W \rangle \xrightarrow{\ell} \langle \mathbf{E}', W' \rangle$ is a transition from \mathbf{E} to \mathbf{E}' , and shared queues W to W' with *observable action* ℓ . We prove a standard safety theorem (Theorem 5.1 below) that guarantees that if a participant does a transition with some observable action, then so does the type index.

Theorem 5.1. [Soundness] Assume $\mathbf{E} : \text{Mp } C A$, $m : \text{Mp } L a$ and $W : Q$. Suppose $\langle \mathbf{E}[r \mapsto m], W \rangle \xrightarrow{\ell} \langle \mathbf{E}'[r \mapsto m'], W' \rangle$. Then there exists $\langle C[r \mapsto L], Q \rangle \xrightarrow{\ell} \langle C'[r \mapsto L'], Q' \rangle$ such that $W' : Q'$ and $m' : \text{Mp } L' a$.

Notations and Operations for Parallel Programs We simplify the notation for \mathbf{E} , when all L_i are projections of the same global type, and the a_i are *projections of the same interface*. We define the projection of an interface at a participant, $A \upharpoonright p$, to be the part of A that is at p (Appendix A.6). Whenever we have $m_p : \text{Mp } (G \upharpoonright p) (A \upharpoonright p)$ for all participants in $p \in G$, we use the notation $[p \mapsto m_p]_{p \in \text{pids}(G)} : \text{Mp } G A$. This means that the collection of all actions m_p *behave* as prescribed by G , and produce their result in interface A . Finally, if we have $\mathbf{E} = [p \mapsto f_p : A \upharpoonright p \rightarrow \text{Mp } (G \upharpoonright p) (B \upharpoonright p)]_{p \in \text{pids}(G)}$, we write $\mathbf{E} : A \rightarrow \text{Mp } G B$.

Parallel programs have a default value for participants that are not in their domain. Unless otherwise specified, this default value is the identity. For example, $\mathbf{E}(p) = f$ if $\mathbf{E} = \mathbf{E}'[p \mapsto f]$, and $\mathbf{E}(p) = \lambda x. \text{ret } x$ if $p \notin \mathbf{E}$. We specify the default value using the underscore character as a key in the mapping from participants to monadic actions: $[_ \mapsto f]$.

Distributed Values and Execution We define the execution of a parallel program on a distributed value below. A distributed value $V : a @ R$ is a mapping from participants to the value that they hold in the respective interface: $[p_i \mapsto (v_i : (a @ R) \upharpoonright p_i)]_{i \in I} : (a @ R)$. Additionally, we require unit to be the default value, so if $p \notin \text{pids}(R)$, then $V(p) = ()$.

Definition A.9 (Execution). Given $\mathbf{E} = [p_i \mapsto f_i]_{i \in I}$ and $X = [p_j \mapsto x_j]_{j \in J}$, we define $\mathbf{E}(X) = [p_i \mapsto f_i X(p_i)]_{i \in I}$, with $X(p_i) = x_i$ if $i \in J$, or $X(p_i) = ()$ otherwise. Given $Y =$

$[p_k \mapsto y_k]_{k \in K}$, we say that $P(X)$ executes to Y , $P(X) \rightsquigarrow^* Y$, if there is a trace $\langle P(X), \emptyset \rangle \rightsquigarrow^* \langle [p_i \mapsto \text{ret } Y(p_i)]_{i \in I}, \emptyset \rangle$. We write $P(X) = Y$, whenever there is a unique Y s.t. for all Z , $P(X) \rightsquigarrow^* Z$ implies that $Z = Y$.

Composition and Identity Composition is defined as the standard Kleisli composition, extended to parallel programs as follows: $\mathbf{E}_1 \gg \mathbf{E}_2 = [p \mapsto \mathbf{E}_1(p) \gg \mathbf{E}_2(p)]_{p \in \mathbf{E}_1 \cup \mathbf{E}_2}$. Then, $\mathbf{E}_2 \circ \mathbf{E}_1 = \mathbf{E}_1 \gg \mathbf{E}_2$. Identity is simply the empty program with just the default value, $\text{id} = []$.

Split and Projection The split operation is the participant-wise split, and the i -th projection is the environment with the projection i as the default value:

$$\mathbf{E}_1 \Delta \mathbf{E}_2 = [p \mapsto \lambda x. \mathbf{E}_1(p) x \gg \lambda y. \mathbf{E}_2(p) x \gg \lambda z. \text{ret}(y, z)]_{p \in \mathbf{E}_1 \cup \mathbf{E}_2},$$

$$\pi_i = [_ \mapsto \lambda x. \text{ret}(\pi_i x)]$$

Case and Injection Case expressions will never occur during code generation, since they will be resolved by choices. Injections only *tag* a branch in the protocol, and so we define them as the identity: $\iota_i = []$.

Choices Choices are performed by the participant holding a value of a sum-type, and the tag is notified to the list of participants that depend on them. The definition uses functions $\text{get}_I(x)$ and $\text{put}_I(y, x)$ to extract the value of a sum-type from the hole of a one-hole context I (§3.1), and to replace the value at the hole respectively.

$$[p_0 \oplus p_0 p_1 \cdots p_n] = \left[\begin{array}{l} p_0 \mapsto \lambda x. \text{sel}(p_1 \cdots p_n) (\text{get}_I(x)) \\ \quad (\lambda y. \text{ret}(\text{put}_I(y, x))) (\lambda y. \text{ret}(\text{put}_I(y, x))) \\ p_1 \mapsto \lambda x. \text{brn } p (\text{ret } x) (\text{ret } x); \\ \dots \\ p_n \mapsto \lambda x. \text{brn } p (\text{ret } x) (\text{ret } x) \end{array} \right]$$

The presence of type \uplus means that we might require to perform a case expression to inspect the result of a previous choice: we define $\mathbf{E}_1 \uplus^{\vec{p}} \mathbf{E}_2$ for this.

$$\mathbf{E}_1 \uplus^{\vec{p}} \mathbf{E}_2 = [p \mapsto \lambda x. \text{case } \mathbf{E}_1(p) \mathbf{E}_2(p)]_{p \in \vec{p} \cup (\mathbf{E}_1 \cup \mathbf{E}_2) \setminus \{\vec{p}\}}$$

The definition of $\uplus^{\vec{p}}$ means that the participants involved in a choice will perform a case expression to inspect which branch they need to take, while the rest of the participants will continue as specified by either \mathbf{E}_1 or \mathbf{E}_2 . Note $(\mathbf{E}_1 \cup \mathbf{E}_2)(p)$ will produce $\mathbf{E}_2(p)$, if $p \in \mathbf{E}_1 \cap \mathbf{E}_2$. This will not be an issue during code generation: any participant that is not involved in a choice will have the same continuation in both branches.

A.5 Mp code generation

The translation scheme for Mp code generation (Fig. 12) is done recursively on the structure of PALg expressions. It takes a PALg expression e , an interface A , and produces a mapping from all participants in e and A to their respective monadic continuations. We write $\llbracket e \rrbracket(A)$, and guarantee that $\llbracket e \rrbracket(A) : A \rightarrow \text{Mp } G B$, if $\models e \Leftarrow A \sim (G, B)$. This means that if e induces protocol G with interfaces $A \rightarrow B$, then the generated code behaves as G , with interfaces A and B .

Code generation follows a similar structure to global type inference. For code generation, we require a partial function $\text{cod}(e, A)$ that infers the codomain interface of e using Fig. 2. The translation to Mp requires to define the interactions from an interface I that gathers a type a at p : $(a@I \rightsquigarrow p) : a@I \rightarrow \text{Mp } [a@I \rightsquigarrow p] (a@p)$. The definition is analogous to that of $[a@I \rightsquigarrow p]$. The remaining of the translation is straightforward, built on top of the previous definitions.

$$\begin{aligned} (a@p_1 \rightsquigarrow p_0) &= [p_1 \mapsto \lambda x. \text{send } x \text{ p}_0 \text{ p}_0 \mapsto \lambda _ . \text{recv } p_1 \ a] \\ ((a_1 + a_2)@(\iota_i \ I) \rightsquigarrow p) &= (a_i@I \rightsquigarrow p) \ggg [p \mapsto \lambda x. \text{ret } (\iota_i \ x)] \\ ((a \times b)@(\iota_1 \times \iota_2) \rightsquigarrow p) &= (a@I_1 \rightsquigarrow p) \times (b@I_2 \rightsquigarrow p) \\ &\ggg [p_i \mapsto \lambda _ . \text{ret}()]_{p_i \in \text{pids}(I_1 \times I_2) \setminus \{p\}} \end{aligned}$$

$$\begin{aligned} \llbracket \text{id} \rrbracket (a@I) &= [] & \llbracket \iota_i \rrbracket (a@I) &= [] \\ \llbracket e@p \rrbracket (a@I) &= (a@I \rightsquigarrow p) \ggg [p \mapsto \lambda x. \text{ret } (e \ x)] \\ \llbracket e_1 \Delta e_2 \rrbracket (a@I) &= \llbracket e_1 \rrbracket (a@I) \Delta \llbracket e_2 \rrbracket (a@I) \\ \llbracket e_1 \circ e_2 \rrbracket (A) &= \llbracket e_2 \rrbracket (A) \ggg \llbracket e_1 \rrbracket (\text{cod}(e_2, A)) \\ \llbracket e_1 \nabla e_2 \rrbracket ((a_1 + a_2)@(\iota_i \ I)) &= \llbracket e_i \rrbracket (a_i@I) \\ \llbracket \pi_i \rrbracket ((a \times b)@(\iota_1 \times \iota_2)) &= \pi_i \\ \llbracket [p \oplus \vec{p}] \rrbracket (a@I [p]) &= [p \oplus \vec{p}] \\ \llbracket e \rrbracket (a@(R_1 \cup^{\vec{p}} R_2)) &= \llbracket e \rrbracket (a@R_1) \uplus^{\vec{p}} \llbracket e \rrbracket (a@R_2) \end{aligned}$$

Figure 12. Translating PAIg to Mp code.

Protocol Compliance Theorem 5.2 guarantees that the generated code follows the protocol inferred using the relation in Fig. 3. This fact is enough to guarantee that the generated code is deadlock-free. Moreover, we can use it to prove that the generated code is extensionally equal to the input Alg expression. We state this in Theorem 5.3.

Theorem 5.2. [Protocol Conformance of the Generated Code] *If $\models e \Leftarrow A \sim G$, then $\llbracket e \rrbracket (A)$ complies with protocol G .*

This theorem is proved by induction on the structure of the derivation \models , and by the definition of \uparrow . This result guarantees that the generated code corresponds to the protocol inferred from e . Since the protocol inferred from e is deadlock-free, then so is the generated code. See Appendix B.3.

Extensional Equivalence Additionally to deadlock-freedom and protocol compliance, we prove that if e is the annotation of e , then running the code generated from e on x produces the same result as evaluating e on x . This guarantees that, regardless of the annotations and interfaces chosen for e , the parallel code always produces the same result as the sequential implementation. We show the statement below, in Theorem 5.3, and refer to **Appendix C** for the full proof.

We specify the extensionality theorem on *runnable* parallel programs, which are those with a single entry/exit point, i.e. a *master* worker p_m that starts the computation, and gathers the results. Suppose we call this master worker p_m . Given any e , we can guarantee that p_m is the entry point by setting it to be the domain interface: $\vdash e : a@p_m \rightarrow b@R$. To make it the exit point, we need to make sure that it is the codomain interface. We can do this by forcing the participants in R to

send their values to p_m as follows: $\vdash \text{id}@p_m \circ e : a@p_m \rightarrow b@R$. However, and due to the presence of choices, the codomain interface may contain \cup . For example, if $e : a@p_m \rightarrow b@(I_1 \cup^{\vec{p}} I_2)$, with $I_1 \neq I_2$, then $\vdash \text{id}@p_m \circ e : a@p_m \rightarrow b@(p_m \cup^{\vec{p}} p_m)$, with $p_m \in \vec{p}$. This means that $b@(p_m \cup^{\vec{p}} p_m) \uparrow p_m = b \uplus b$. To obtain a single value of type b , we use function $\text{join} : a \uplus a \rightarrow a$, which is equivalent to $\text{id} \nabla \text{id}$ for regular sum-types. We lift it to a monadic action, $\text{join}(R)$, to join all branches in R :

$$\text{join}(I) = [] \quad \text{join}(R_1 \cup^{\vec{p}} R_2) = \text{join}(R_1) \uplus^{\vec{p}} \text{join}(R_2) \ggg [p \mapsto \lambda x. \text{ret } (\text{join } x)]_{p \in \vec{p}}$$

Note that $\text{join}(R_1 \cup^{\vec{p}} R_2)$ is only defined if $(a@R_1 \uparrow p) = (a@R_2 \uparrow p)$ for all roles. The *runnable parallel program* for $e : a@p_m \rightarrow b@R$, $J\llbracket e \rrbracket (p_m)$, is defined as follows:

$$J\llbracket e \rrbracket (p_m) = \llbracket \text{id}@p_m \circ e \rrbracket (a@p_m) \ggg [p_m \mapsto \text{join}(R)].$$

Our extensionality statement specifies that executing the runnable parallel program for e , with master p and value x , produces value $e \ x$ at p .

Theorem 5.3. [Extensionality] *Assume $e \Rightarrow e : a@p \rightarrow b@R$ and $x : a$ initially at p . If $e \ x = y$, then the execution of $\llbracket e \rrbracket (p)$ also produces y , distributed across R .*

Example A.10 (MergeSort Code Generation). We start with the annotated ms from Example A.1, and we use p_1 as master role to avoid the initial communication from p_0 to p_1 :

$$\begin{aligned} \text{pms} &= \text{mrg}@p_1 \circ (\text{id} + (\text{ms}@p_2 \circ \pi_1@p_1) \Delta (\text{ms}@p_3 \circ \pi_2@p_1)) \\ &\circ [p_1 \oplus p_1 p_2 p_3] \circ \text{spl}@p_1 : \text{Ls}@p_1 \rightarrow \text{Ls}@p_1 \cup^{p_1 p_2 p_3} \text{Ls}@p_1 \end{aligned}$$

Note that $\llbracket \text{id}@p_1 \circ \text{pms} \rrbracket$ would be equivalent to $\llbracket \text{pms} \rrbracket$ because the output interface is already of the form $p_1 \cup p_1$. Therefore, for simplicity, we show $\llbracket \text{pms} \rrbracket$, and use it to produce the runnable parallel program. Fig. 14 show the code generation process using a table, where the i -th column is the current code for participant p_i , and the last column shows the expression and input interface that we are translating next.

From this point, we need to produce the code for the two branches. The left branch is straightforward, and is simply $\lambda x. \text{ret } x$ for all participants. The result for the right branch is show in Fig. 15:

Next, we combine both branches using **case**, and compose it with the previous result. To avoid unnecessarily pattern-matching expressions such as **case**, we optimise the code using rules of the form:

$$\text{sel } \vec{p} \ f_1 \ f_2 \ggg \text{case } f_3 \ f_4 = \text{sel } \vec{p} \ (f_1 \ggg f_3) \ (f_2 \ggg f_4).$$

Additionally, we optimise all instances of e.g. $(x, ())$ using the fact that $1 \times a \cong a \times 1 \cong a$. We show below the code for all p_i , after composing it with $\text{join}(p_1 \cup^{p_1 p_2 p_3} p_1)$, and applying these optimisations. We use different colours to highlight the different branches of the protocol:

1981 $p_1 \mapsto \mathbf{sel} \{p_2, p_3\} (\mathbf{spl} \ v) (\lambda x. \mathbf{ret} (\mathbf{mrg} (t_1 \ x))) (\lambda x. \mathbf{send} \ p_2 (\pi_1 \ x) \gg \lambda y. \mathbf{send} \ p_3 (\pi_2 \ x) \gg \lambda _.$ 2036
1982 $\mathbf{recv} \ p_2 \ Ls \gg \lambda x. \mathbf{recv} \ p_3 \ Ls \gg \lambda y. \mathbf{ret} (\mathbf{mrg} (t_2 \ (x, y)))) \gg \lambda x. \mathbf{ret} (\mathbf{join} \ x)$ 2037
1983 $p_2 \mapsto \mathbf{brn} \ p_1 (\mathbf{ret} \ x) (\mathbf{recv} \ p_1 \ Ls \gg \lambda x. \mathbf{ret} (\mathbf{ms} \ x) \gg \lambda x. \mathbf{send} \ p_1 \ x) \gg \lambda x. \mathbf{ret} (\mathbf{join} \ x)$ 2038
1984 $p_3 \mapsto \mathbf{brn} \ p_1 (\mathbf{ret} \ x) (\mathbf{recv} \ p_1 \ Ls \gg \lambda x. \mathbf{ret} (\mathbf{ms} \ x) \gg \lambda x. \mathbf{send} \ p_1 \ x) \gg \lambda x. \mathbf{ret} (\mathbf{join} \ x)$ 2039
1985 $p_1 \mapsto \mathbf{send} \ p_2 (\pi_1 (v_1, v_2)) \gg \lambda y. \mathbf{send} \ p_3 (\pi_2 (v_1, v_2)) \gg \lambda _.$ 2040
1986 $\mathbf{recv} \ p_2 \ Ls \gg \lambda x. \mathbf{recv} \ p_3 \ Ls \gg \lambda y. \mathbf{ret} (\mathbf{br}_2 (\mathbf{mrg} (t_2 \ (x, y)))) \gg \lambda x. \mathbf{ret} (\mathbf{join} \ x)$ 2041
1987 $p_2 \mapsto \mathbf{recv} \ p_1 \ Ls \gg \lambda x. \mathbf{ret} (\mathbf{ms} \ x) \gg \lambda x. \mathbf{send} \ p_1 \ x \gg \lambda x. \mathbf{ret} (\mathbf{br}_2 \ x) \gg \lambda x. \mathbf{ret} (\mathbf{join} \ x)$ 2042
1988 $p_3 \mapsto \mathbf{recv} \ p_1 \ Ls \gg \lambda x. \mathbf{ret} (\mathbf{ms} \ x) \gg \lambda x. \mathbf{send} \ p_1 \ x \gg \lambda x. \mathbf{ret} (\mathbf{br}_2 \ x) \gg \lambda x. \mathbf{ret} (\mathbf{join} \ x)$ 2043
1989 $p_1 \mapsto \mathbf{recv} \ p_2 \ Ls \gg \lambda x. \mathbf{recv} \ p_3 \ Ls \gg \lambda y. \mathbf{ret} (\mathbf{br}_2 (\mathbf{mrg} (t_2 \ (x, y)))) \gg \lambda x. \mathbf{ret} (\mathbf{join} \ x)$ 2044
1990 $p_2 \mapsto \mathbf{ret} (\mathbf{ms} \ v_1) \gg \lambda x. \mathbf{send} \ p_1 \ x \gg \lambda x. \mathbf{ret} (\mathbf{br}_2 \ x) \gg \lambda x. \mathbf{ret} (\mathbf{join} \ x)$ 2045
1991 $p_3 \mapsto \mathbf{ret} (\mathbf{ms} \ v_2) \gg \lambda x. \mathbf{send} \ p_1 \ x \gg \lambda x. \mathbf{ret} (\mathbf{br}_2 \ x) \gg \lambda x. \mathbf{ret} (\mathbf{join} \ x)$ 2046
1992 $p_1 \mapsto \mathbf{ret} (\mathbf{mrg} (t_2 (\mathbf{ms} \ v_1, \mathbf{ms} \ v_2))) = \mathbf{ret} (\mathbf{mrg} ((\mathbf{id} + \mathbf{ms} \times \mathbf{ms})(\mathbf{spl} \ v))) = \mathbf{ret} (\mathbf{ms} \ v)$ 2047
1993 $p_2 \mapsto \mathbf{ret} () \mid p_3 \mapsto \mathbf{ret} ()$ 2048
1994 2049
1995 2050

Figure 13. Step-by-step execution of the parallel code for \mathbf{ms} . Input is $[p_1 \mapsto v]$, with $\mathbf{spl} \ v = t_2 (v_1, v_2)$.

1996	p_1	p_2	p_3	
1997	$\lambda x. \mathbf{ret} \ x$	$\lambda x. \mathbf{ret} \ x$	$\lambda x. \mathbf{ret} \ x$	$\llbracket \mathbf{spl} @ p_1 \rrbracket (Ls @ p_1)$
1998	$\lambda x. \mathbf{ret} (\mathbf{spl} \ x)$	$\lambda x. \mathbf{ret} \ x$	$\lambda x. \mathbf{ret} \ x$	$\llbracket [p_1 \oplus p_1 p_2 p_3] \rrbracket (((1 + a) + Ls \times Ls) @ p_1)$
1999	$\lambda x. \mathbf{ret} (\mathbf{spl} \ x) \gg \lambda x. \mathbf{sel} \{p_2, p_3\}$	$\lambda x. \mathbf{brn} \ \{p_1\}$	$\lambda x. \mathbf{brn} \ \{p_1\}$	$\llbracket \mathbf{id} + \dots \rrbracket (((1 + a) + Ls \times Ls)$
2000	$(\lambda x. \mathbf{ret} \ x) (\lambda x. \mathbf{ret} \ x)$	$(\lambda x. \mathbf{ret} \ x) (\lambda x. \mathbf{ret} \ x)$	$(\lambda x. \mathbf{ret} \ x) (\lambda x. \mathbf{ret} \ x)$	$@ (t_1 \ p_1 \cup^{p_1 p_2 p_3} t_2 \ p_2))$

Figure 14. Example Translation

2004	p_1	p_2	p_3	
2005	$\lambda x. \mathbf{ret} (\pi_1 \ x) \gg \lambda y. \mathbf{send} \ p_2 \ y \gg$	$\lambda _ . \mathbf{recv} \ p_1 \ Ls \gg$	$\lambda _ . \mathbf{recv} \ p_1 \ Ls \gg$	$\llbracket (\mathbf{ms} @ p_2 \circ \pi_1 @ p_1) \Delta$
2006	$\lambda z. \mathbf{ret} (\pi_2 \ x) \gg \lambda t. \mathbf{send} \ p_3 \ z \gg$	$\lambda x. \mathbf{ret} (\mathbf{ms} \ x) \gg \lambda x. \mathbf{ret} (x, ())$	$\lambda x. \mathbf{ret} (\mathbf{ms} \ x) \gg \lambda x. \mathbf{ret} ((), x)$	$(\mathbf{ms} @ p_3 \circ \pi_2 @ p_1) \rrbracket ((Ls \times Ls) @ p_1)$
2007	$\lambda r. \mathbf{ret}(z, r)$			

Figure 15. Right branch for mergesort.

2011 $p_1 \mapsto \lambda x. \mathbf{sel} \{p_2, p_3\} (\mathbf{spl} \ x) (\lambda x. \mathbf{ret} (\mathbf{mrg} (t_1 \ x)))$
2012 $(\lambda x. \mathbf{send} \ p_2 (\pi_1 \ x) \gg \lambda y. \mathbf{send} \ p_3 (\pi_2 \ x) \gg \lambda _.$
2013 $\mathbf{recv} \ p_2 \ Ls \gg \lambda x. \mathbf{recv} \ p_3 \ Ls \gg \lambda y. \mathbf{ret} (\mathbf{mrg} (t_2 \ (x, y))))$
2014 $\gg \lambda x. \mathbf{ret} (\mathbf{join} \ x)$
2015 $p_2 \mapsto \lambda x. \mathbf{brn} \ p_1 (\mathbf{ret} \ x) (\mathbf{recv} \ p_1 \ Ls \gg \lambda x. \mathbf{ret} (\mathbf{ms} \ x) \gg \lambda x.$
2016 $\mathbf{send} \ p_1 \ x) \gg \lambda x. \mathbf{ret} (\mathbf{join} \ x)$
2017 $p_3 \mapsto \lambda x. \mathbf{brn} \ p_1 (\mathbf{ret} \ x) (\mathbf{recv} \ p_1 \ Ls \gg \lambda x. \mathbf{ret} (\mathbf{ms} \ x) \gg \lambda x.$
2018 $\mathbf{send} \ p_1 \ x) \gg \lambda x. \mathbf{ret} (\mathbf{join} \ x)$

2019 We show in Figure 13 the step-by-step execution of this code
2020 on distributed value $[p_1 \mapsto v]$, with $\mathbf{spl} \ v = t_2 (v_1, v_2)$. The
2021 final result is equal to p_1 applying \mathbf{ms} directly on the input.
2022

2023 **Typing Mp against Local Types** We define a relation between Mp code and the local type that captures their communication behaviour (Fig. 16). We define a judgement of the form $\Gamma \vdash m : \mathbf{Mp} \ L \ a$, where $\mathbf{Mp} \ L \ a$ is the type of an Mp expression that conforms to protocol L and returns a value of type a . The types are parameterised by a variable l that represents a local type continuation. The rules in Fig. 16 are straightforward, since they relate in a one-to-one way to the constructs of local types.

2033 **Semantics** The operational semantics of Mp terms is standard, and mirrors that of the local type configurations in [27].
2034
2035

The operational semantics is defined as an LTS with transitions of the form $\langle [p_i \mapsto m_i]_{i \in I}, W \rangle \xrightarrow{\ell} \langle [p_i \mapsto m'_i]_{i \in I}, W' \rangle$. Here $\ell ::= p_0 p_1 ! \langle a \rangle \mid p_0 p_1 ? \langle a \rangle \mid p_0 p_1 \oplus t_i \mid p_0 p_1 \& t_i$ is the observable action that takes place, and represents, respectively, p_0 sends to p_1 a value of type a , p_1 receives from p_0 , p_0 sends label i to p_1 , and p_1 receives label i from p_0 . We use the special symbol ϵ to represent that no communication took place. Finally, W is a mapping from ordered pairs of roles to unbounded buffers that contain the data sent between participants.

2077 **Definition A.11.** LTS for Mp Terms $\langle [P, W] \xrightarrow{\ell} [P', W'] \rangle$,
2078 P, W transitions to P', W' with action ℓ . The transition rules are defined in Fig. 17.
2079
2080

2081 Similarly, we define $\langle C, Q \rangle \rightarrow^\ell \langle C', Q' \rangle$ for the LTS of local type configurations (App. A). Here, C is a collection of local types, $C = [p_i \mapsto L_i]_{i \in I}$, and Q is a mapping from ordered pairs of roles to unbounded buffers that contain types of the data exchanged. We also say that W is compatible with $Q, W : Q$, if for all pair $p_1 p_2$, if $w_1 \cdots w_n = W(p_1 p_2)$ then $a_1 \cdots a_n = Q(p_1 p_2)$, and for all $i, w_i : a_i$.
2082
2083
2084
2085
2086
2087

2088 **Definition A.12** (Get/Set for Types from One Hole Contexts). Whenever a role performs a choice, we have a type
2089
2090

2091	RET	$\frac{\Gamma \vdash v : a}{\Gamma \vdash \text{ret } v : \text{Mp end } a}$	SEND	$\frac{\Gamma \vdash v : a}{\Gamma \vdash \text{send } r v : \text{Mp } (p!(a). \text{end}) 1}$
2095	RECV	$\frac{\text{ABS}}{\Gamma, x : a \vdash m : \text{Mp } L b}$	$\frac{\Gamma \vdash \text{recv } r a : \text{Mp } (p?(a). \text{end}) a}{\Gamma \vdash \lambda x. m : a \rightarrow \text{Mp } L b}$	
2099	BIND	$\frac{\Gamma \vdash m : \text{Mp } L_1 a \quad \Gamma \vdash f : a \rightarrow \text{Mp } L_2 b}{\Gamma \vdash m \gg f : \forall l_2. \text{Mp } (L_1 \wp L_2) a}$		
2103	BRANCH	$\frac{\Gamma \vdash m_1 : \text{Mp } L_1 a_1 \quad \Gamma \vdash m_2 : \text{Mp } L_2 a_2}{\Gamma \vdash \text{brn } r m_1 m_2 : \text{Mp } (r \& \{l_1.L_1; l_2.L_2\}) (a_1 \uplus a_2)}$		
2107	SELECT	$\frac{\Gamma \vdash v : a + b \quad \Gamma \vdash f_1 : a \rightarrow \text{Mp } L_1 c_1 \quad \Gamma \vdash f_2 : b \rightarrow \text{Mp } L_2 c_2}{\Gamma \vdash \text{sel } v \{r_j\}_{j \in J} f_1 f_2 : \text{Mp } (\{p_j\}_{j \in J} \oplus \{l_1.L_1; l_2.L_2\}) (c_1 \uplus c_2)}$		
2111	CASE	$\frac{\Gamma \vdash f_1 : a_1 \rightarrow \text{Mp } L_1 b_1 \quad \Gamma \vdash f_2 : a_2 \rightarrow \text{Mp } L_2 b_2}{\Gamma \vdash \text{case } f_1 f_2 : a_1 \uplus a_2 \rightarrow \text{Mp } (L_1 \cup L_2) (b_1 \uplus b_2)}$		

Figure 16. Typing rules for Mp code.

with one hole, $a@I[p]$, with a sum-type at the hole $(b + c)@p$. The code for p requires to extract the sum type from the type a , and to set the value at the hole pointed by I . This is because p may occur deep in $I[p]$ and, therefore $a@I[p] \upharpoonright p$ may be different to $b + c$. We achieve this with the following families of functions, $\text{get}_I : a \rightarrow \text{typeAt}(I, a)$, and $\text{put}_I : c \times a \rightarrow \text{substTy}(I, a, c)$, where typeAt and substTy get/set the type at the hole in I .

B Deadlock-Freedom

B.1 Proof of Lemma 4.2

Lemma 4.2. [Existence of Associated Global Type] For all $\text{WF}(A)$, if $\vdash e : A \rightarrow B$, then there exists G s.t. $\vdash e \Leftarrow A \sim G$.

Proof. By induction on the structure of the derivation $\vdash e : A \rightarrow B$.

Case JOIN. $\vdash e \Rightarrow e : A \cup^{\vec{p}} A \rightarrow B \cup^{\vec{p}} B$. By the IH, $\vdash e \Leftarrow A \sim G$. By rule ALT, $\vdash e \Leftarrow A \cup^{\vec{p}} A \sim G \cup G$.

Case ALG. $\vdash e \Rightarrow e@p : a@I \rightarrow b@p$. By ALG, $\vdash e@r \Leftarrow a@I \sim [a@I \rightsquigarrow r]$.

Case INJ_i. By INJ_i, $\vdash l_i \Leftarrow A \sim \text{end}$.

Case ALT. $\vdash e \Rightarrow e : A_1 \cup^{\vec{p}} A_2 \rightarrow B_1 \cup^{\vec{p}} B_2$, with $\text{pids}(e) \subseteq \vec{p}$. By the IH, $\vdash e \Leftarrow A_1 \sim G_1$ and $\vdash e \Leftarrow A_2 \sim G_2$. By ALT, $\vdash e \Leftarrow A_1 \cup A_2 \sim G_1 \cup^{\vec{p}} G_2$.

Case ID. By $\text{id} \Leftarrow a@I \sim \text{end}$.

Case CHOICE. $\vdash e \Rightarrow [p \oplus \vec{p}] : a@I[p] \rightarrow a@(I[l_1 p] \cup^{\vec{p}} I[l_2 p])$. By rule CHOICE, $\vdash [p \oplus \vec{p}] \Leftarrow a@I[p] \sim p \rightarrow \{ \vec{p} \} \{l_1. \text{end}; l_2. \text{end}\}$.

Case PROJ_i. By rule PROJ_i, $\vdash \pi_i \Leftarrow A_1 \times A_2 \sim \text{end}$.

Case COMP. $\vdash e_1 \circ e_2 \Rightarrow e_1 \circ e_2 : A \rightarrow C$. This implies that $\vdash e_1 \Rightarrow e_1 : B \rightarrow C$ and $\vdash e_2 \Rightarrow e_2 : A \rightarrow B$. By the IH, $\vdash e_2 \Leftarrow A \sim G_2$, and $\vdash e_1 \Leftarrow B \sim G_1$. We proceed by induction on the size of B . The base case is $b@I_2$. In this case, A must be $a@I_1$, so G_1 must not be \cup or contain any choices. Therefore, $G_1 \wp G_2$ is defined, and equal to $[G_2/\text{end}]G_1$. If $B = b@(R_{21} \cup R_{22})$, then $G_2 = G_{21} \cup G_{22}$. There are two cases: (1) $G_1 = G_{11} \cup G_{12}$, or (2) there is a choice in G_1 , i.e. $G_1 = [p \rightarrow \vec{p}\{l_1. G_{11}; l_2. G_{12}\}/\text{end}]G'_1$. In the first case, we have that $\vdash e_1 \Leftarrow a@R_{1i} \sim G_{1i}$, and $\vdash e_2 \Leftarrow a@R_{2i} \sim G_{2i}$, which by the IH implies that $G_{1i} \wp G_{2i}$ is defined. Therefore, $(G_{11} \cup G_{12}) \wp (G_{21} \cup G_{22}) = (G_{11} \wp G_{21}) \cup (G_{12} \wp G_{22})$. In the second case, there must be two sub-expressions of e_1 , e_{11} and e_{12} s.t. $e_{11} \Leftarrow a@I \sim G'_1$, and $e_{12} \circ [p \oplus \vec{p}] \Leftarrow d@I[p] \sim p \rightarrow \vec{p}\{l_1. G_{11}; l_2. G_{12}\}$, with $e_{12} : d@I[l_i p] \rightarrow b@R_{2i}$ and $e_{12} \Leftarrow d@I[l_i p] \sim G_{1i}$. By the IH, $G_{1i} \wp G_{2i}$ must be defined, which implies that $p \rightarrow \vec{p}\{l_1. G_{11}; l_2. G_{12}\} \wp (G_{21} \cup G_{22})$ is defined, and therefore $G_1 \wp (G_{21} \cup G_{22})$ is also defined.

Case CASE_i. $\vdash e_1 \nabla e_2 \Rightarrow e_1 \nabla e_2 : l_i A \rightarrow B$. By inversion, $\vdash e_i \Rightarrow e_i : A \rightarrow B$. By the IH, $\vdash e_i \Leftarrow A \sim G$. By CASE_i, $\vdash e_1 \nabla e_2 \Leftarrow l_i A \sim G$.

Case SPLIT. $\vdash_{\text{det}} e_1 \Delta e_2 \Rightarrow e_1 \Delta e_2 : a@I \rightarrow B \times C$. By inversion, $\vdash e_1 \Rightarrow e_1 : a@I \rightarrow B$ and $\vdash e_2 \Rightarrow e_2 : a@I \rightarrow C$. By the IH, $\vdash e_1 \Leftarrow a@I \sim G_1$ and $\vdash e_2 \Leftarrow a@I \sim G_2$. Therefore, $\vdash e_1 \Delta e_2 \Leftarrow a@I \sim [G_2/\text{end}]G_1$. \square

B.2 Proof of Lemma 4.3

Lemma 4.3. [Protocol Deadlock-Freedom] For all $\text{WF}(A)$, if $\vdash e : A \rightarrow B$ and $\vdash e \Leftarrow A \sim G$, then $\text{WF}(G)$.

Proof. By induction on the structure of the derivation $\vdash e : a@I \rightarrow B$.

Case JOIN. $\vdash e \Rightarrow e : A \cup^{\vec{p}} A \rightarrow B \cup^{\vec{p}} B$. By case analysis, the only possibility for G is that it is obtained via the JOIN protocol rule: $\vdash e \Leftarrow A \cup A \sim G$. By the JOIN typing rule, $\vdash e \Rightarrow e : A \rightarrow B$. By the JOIN protocol rule, $\vdash e \Leftarrow A \cup^{\vec{p}} A \sim G \cup G$. By the IH, $\text{WF}(G)$, therefore $\text{WF}(G \cup G)$.

Case ALG. $\vdash e \Rightarrow e@p : a@I \rightarrow b@p$. By case analysis, $\vdash e@p \Leftarrow a@I \sim [a@I \rightsquigarrow p]$. By definition, $\text{WF}([a@I \rightsquigarrow p])$.

Case INJ_i. $\vdash l_i \Leftarrow a@I \sim \text{end}$. Trivial $\text{WF}(\text{end})$.

Case ALT. $\vdash e \Rightarrow e : A_1 \cup^{\vec{p}} A_2 \rightarrow B_1 \cup^{\vec{p}} B_2$, with $\text{getRoles}(e) \subseteq \vec{p}$, $\vdash e \Rightarrow e : A_1 \rightarrow B_1$, $\vdash e \Rightarrow e : A_2 \rightarrow B_2$, and $A_1 \neq A_2$. We know, by a straightforward induction on the typing rules for PALG, that if $\vdash e : A_1 \cup^{\vec{p}} A_2 \rightarrow B$, and $A_1 \neq A_2$, then $\text{pids}(A_i) \subseteq \vec{p}$. By the ALT protocol rule, $\vdash e \Leftarrow A_1 \cup^{\vec{p}} A_2 \sim G_1 \cup^{\vec{p}} G_2$, with $\vdash e \Leftarrow A_1 \sim G_1$ and $\vdash e \Leftarrow A_2 \sim G_2$. By the IH, $\text{WF}(G_1)$ and

$$\begin{array}{c}
 \boxed{\langle P, W \rangle \rightsquigarrow^\ell \langle P', W' \rangle} \quad P = [p_i \mapsto m_i]_{i \in I} \quad W = [p_i p_j \mapsto w]_{i \in I, j \in I} \\
 \hline
 \langle P[p \mapsto m], W \rangle \rightsquigarrow^\ell \langle P[p \mapsto m'], W' \rangle \quad \langle P[p \mapsto \text{ret } v \gg f], W \rangle \rightsquigarrow^\epsilon \langle P[p \mapsto f v], W \rangle \\
 \langle P[p \mapsto m \gg f], W \rangle \rightsquigarrow^\ell \langle P[p \mapsto m' \gg f], W' \rangle \\
 \langle P[p_1 \mapsto \text{send } p_2 (v : a)], W[p_1 p_2 \mapsto w] \rangle \rightsquigarrow^{p_1 p_2!(a)} \langle P[p_1 \mapsto \text{ret } ()], W[p_1 p_2 \mapsto v \cdot w] \rangle \\
 \langle P[p_1 \mapsto \text{recv } p_2 a], W[p_2 p_1 \mapsto w \cdot v] \rangle \rightsquigarrow^{p_2 p_1?(a)} \langle P[p_1 \mapsto \text{ret } v], W[p_2 p_1 \mapsto w] \rangle \\
 \langle P[p_0 \mapsto \text{sel } (t_i v) \{ \} f_1 f_2], W \rangle \rightsquigarrow^\epsilon \langle P[p_0 \mapsto f_i v \gg \lambda x. \text{ret } (\text{br}_i x)], W \rangle \\
 \langle P[p_0 \mapsto \text{sel}(t_i v)\{p_1 \dots p_n\} f_1 f_2], W[p_0 p_1 \mapsto w] \rangle \rightsquigarrow^{p_0 p_1 \oplus t_i} \langle P[p_0 \mapsto \text{sel}(t_i v)\{p_2 \dots p_n\} f_1 f_2], W[p_0 p_1 \mapsto l_i \cdot w] \rangle \\
 \langle P[p_1 \mapsto \text{brn } p_2 m_1 m_2], W[p_2 p_1 \mapsto w \cdot l_i] \rangle \rightsquigarrow^{p_2 p_1 \& t_i} \langle P[p_1 \mapsto m_i \gg \lambda x. \text{ret } (\text{br}_i x)], W[p_2 p_1 \mapsto w] \rangle \\
 \langle P[p_1 \mapsto \text{case } f_1 f_2 (\text{br}_i v)], W \rangle \rightsquigarrow^\epsilon \langle P[p_1 \mapsto f_i v], W \rangle
 \end{array}$$

Figure 17. Rules for the LTS of Mp terms

$$\begin{array}{c}
 \text{get}_{[\]}(x) = x \quad \text{get}_{I \times I}(x) = \text{get}_I(\pi_1 x) \quad \text{get}_{I \times I}(x) = \text{get}_I(\pi_2 x) \quad \text{get}_{t_i I}(x) = \text{get}_I(x) \\
 \text{put}_{[\]}(x, y) = x \quad \text{put}_{I \times I}(x, y) = (\text{put}_I(x, \pi_1 y), \pi_2 y) \quad \text{put}_{I \times I}(x, y) = (\pi_1 y, \text{put}_I(x, \pi_2 y)) \quad \text{put}_{t_i I}(x, y) = \text{put}_I(x, y)
 \end{array}$$

Figure 18. Get/Set for Types from One Hole Contexts

WF(G_2). Since $\text{pids}(G_i) \subseteq \text{pids}(p) \cup \text{pids}(A_1) \cup \text{pids}(A_2) \subseteq \text{pids}(r)$, $\text{WF}(G_1 \cup^{\bar{p}} G_2)$.

Case ID. Trivial by $\text{WF}(\text{end})$.

Case CHOICE. $\vdash e \Rightarrow e \circ [p \oplus \bar{p}] : a@I[p] \rightarrow B_1 \cup \bar{p}B_2$, where $\text{pids}(e) \subseteq \bar{p}$, and $I[p] \subseteq \bar{p}$. By the CHOICE and ALT typing rules, $\vdash e \Rightarrow e : a@I[t_i p] \rightarrow B_i$. By inversion, the protocol rule must be also CHOICE: $\models [p \oplus \bar{p}] \Leftarrow a@I[p] \sim p \rightarrow \bar{p}\{t_i. G_i\}_{i \in [1,2]}$. By the CHOICE protocol rule, $\models e \Leftarrow a@I[t_i p] \sim G_i$. By the IH, $\text{WF}(G_i)$. Since $\text{pids}(G_i) \subseteq \text{pids}(e) \cup \text{pids}(I[p]) \subseteq \bar{p}$, then for all $p' \in G_i$, $(p \rightarrow \bar{p}\{t_i. G_i\}_{i \in [1,2]}) \upharpoonright p'$ must be defined. Therefore, $\text{WF}(p \rightarrow \bar{p}\{t_i. G_i\}_{i \in [1,2]})$.

Case PROJ_i. Trivial by $\text{WF}(\text{end})$.

Case COMP. $\vdash e_1 \circ e_2 \Rightarrow e_1 \circ e_2 : A \rightarrow C$, with $\vdash e_1 \Rightarrow e_1 : B \rightarrow C$ and $\vdash e_2 \Rightarrow e_2 : A \rightarrow B$. By inversion, the only possible protocol rule is also COMP. Therefore, $\models e_1 \circ e_2 \Leftarrow A \sim G_2 \wp G_1$, with $\models e_2 \Leftarrow A \sim G_2$ and $\models e_1 \Leftarrow B \sim G_1$. By the IH, $\text{WF}(G_1)$ and $\text{WF}(G_2)$. Also, by the induction on the derivation of \vdash , we know that $A_1 \cup^{\bar{p}} A_2$, if $A_1 \neq A_2$, then $\text{pids}(A_i) \subseteq \bar{p}$. This implies that if G_1 is $G_{11} \cup^{\bar{p}} G_{12}$, then either the projection onto p of G_{1i} is the same, or $p \in \bar{p}$. By the CHOICE rule, G must be of the form $G'[p \rightarrow \bar{p}\{t_i. G_i\}_{i \in [1,2]}]$, therefore, for for all $p' \in \text{pids}(G_{1i})$, the projection of $(G_2 \wp (G_{11} \cup G_{12})) \upharpoonright p'$ must be defined, which implies that $G_2 \wp G_1$ is defined.

Case CASE_i. $\vdash e_1 \nabla e_2 \Rightarrow e_1 \nabla e_2 : a@(t_i I) \rightarrow B$ and $\models e_1 \nabla e_2 \Leftarrow a@(t_i I) \sim G$. By the CASE_i protocol and typing rules, $\models e_i \Leftarrow A \sim G$ and $\vdash e_i \Rightarrow e_i : A \rightarrow B$. We conclude by the IH that $\text{WF}(G)$.

Case SPLIT. $\vdash e_1 \Delta e_2 \Rightarrow e_1 \Delta e_2 : A \rightarrow (b \times c)@(R_1 \times R_2)$ and $\models e_1 \Delta e_2 \Leftarrow A \sim [G_2/\text{end}]G_1$. By the IH, we know that

$\text{WF}(G_1)$ and $\text{WF}(G_2)$. By straightforward induction on the structure of G_1 , if $\text{WF}(G_i)$, then $\text{WF}([G_2/\text{end}]G_1)$. \square

B.3 Proof of Theorem 5.2

Theorem 5.2. *[Protocol Conformance of the Generated Code] If $\models e \Leftarrow A \sim G$, then $\llbracket e \rrbracket(A)$ complies with protocol G .*

Proof. By induction on the structure of the derivation $\models e \Leftarrow A \sim G$.

Case ALT. $\models e \Leftarrow A_1 \cup^{\bar{p}} A_2 \sim G_1 \cup^{\bar{p}} G_2$, with $\models e \Leftarrow A_1 \sim G_1$, $\models e \Leftarrow A_2 \sim G_2$. By the IH, $\llbracket e \rrbracket(A_i) : (A_i \upharpoonright p) \rightarrow \text{Mp}(G_i \upharpoonright p) (B_i \upharpoonright p)$. Moreover, we know that if $p \notin \bar{p}$, then $\llbracket e \rrbracket(A_1)(p) = \llbracket e \rrbracket(A_2)(p)$, and $G_1 \upharpoonright p = G_2 \upharpoonright p$. Therefore, by the definition of $\mathbf{E}_1 \uplus \mathbf{E}_2$, $\llbracket e \rrbracket(A_1) \uplus^{\bar{p}} \llbracket e \rrbracket(A_2) : \text{Mp}(G_1 \cup^{\bar{p}} G_2) (B_1 \cup^{\bar{p}} B_2)$.

Case ID. $\models \text{id} \Leftarrow a@I \sim \text{end}$. By the definition of $\llbracket \]$, $\llbracket \text{id} \rrbracket(a@I) = [] : a@I \rightarrow \text{Mp end } a@I$.

Case INJ_i. $\models t_i \Leftarrow a@I \sim \text{end}$. By definition, $\llbracket t_i \rrbracket(A) = [] : a@I \rightarrow \text{Mp end } (a@(t_i I))$.

Case ALG. $\models e@p_e \Leftarrow a@I \sim [a@I \rightsquigarrow p_e]$, with $e : a \rightarrow b$. We prove by straightforward induction on the structure of I that $f = (a@I \rightsquigarrow p_e) : a@I \rightarrow \text{Mp}[a@I \rightsquigarrow p_e] (a@p_e)$: if $I = p$, then $[p \mapsto \lambda x. \text{send } p_e x, p_e \mapsto \lambda _ . \text{recv } p a]$, which clearly follows $[a@p \rightsquigarrow p_e]$; if $I = I_1 \times I_2$, then a must be $a_1 \times a_2$, and we have by the IH that $(a_i@I_i \rightsquigarrow p_e) : \text{Mp}[a_i@I_i \rightsquigarrow p_e] (a_i@p_e)$; and, finally, if $I = t_i I'$, then $a = a_1 + a_2$, and $(a_i@I \rightsquigarrow p_e) : \text{Mp}[a_i@I \rightsquigarrow p_e] (a_i@p_e)$, which composed with $[p_e \mapsto \lambda x. \text{ret } (t_i x)]$ has type $\text{Mp}[(a_1 + a_2)@(t_i I) \rightsquigarrow p_e] ((a_1 + a_2)@t_i p_e)$.

Case COMP. $\models e_1 \circ e_2 \Leftarrow A \sim G_2 \wp G_1$. By the COMP rule, $\models e_2 \Leftarrow A \sim G_2$ and $\models e_1 \Leftarrow B \sim G_1$. By the IH, $\llbracket e_2 \rrbracket(A) : A \rightarrow \text{Mp } G_2 B$ and $\llbracket e_1 \rrbracket(B) : B \rightarrow \text{Mp } G_1 C$. Since $G_2 \wp G_1$ is well-formed, then $\llbracket e_2 \rrbracket(A) \gg \llbracket e_1 \rrbracket(B) : A \rightarrow \text{Mp}(G_2 \wp G_1) C$, since for all p ,

2311 $\llbracket e_2 \rrbracket (A) : A \uparrow p \rightarrow \text{Mp } G_2 (B \uparrow p)$ and $\llbracket e_1 \rrbracket (B) : B \uparrow$
 2312 $p \rightarrow \text{Mp } G_1 (C \uparrow p)$, so $\llbracket e_2 \rrbracket (A) \rightsquigarrow \llbracket e_1 \rrbracket (B) : A \uparrow p \rightarrow$
 2313 $\text{Mp } (G_1 \circledast G_2) (C \uparrow p)$.

2314 **Case CHOICE.**

2315 $\models [p \oplus \bar{p}] \Leftarrow a@I[p_c] \sim p_c \rightarrow \{\bar{p}\}\{t_i. \text{end}\}_{i \in [1,2]}$.

2316 By the definition of $[p \oplus \bar{p}]$,

2317 $p \mapsto \lambda x. \text{sel } \{\bar{p}\} (\text{get}_I(x)) (\lambda y. \text{ret } (\text{put}_I(y, x)))$

2318 $(\lambda y. \text{ret } (\text{put}_I(y, x)))$,

2319 which has type $a@I[p] \uparrow p \rightarrow \text{Mp } (\{\bar{p}\} \oplus \{t_i. \text{end}\}_{i \in [1,2]})$,

2320 and $p' \in \bar{p}$, $p' \mapsto \lambda x. \text{brn } p (\text{ret } x) (\text{ret } x)$, which has type
 2321 $a@I[p] \uparrow p' \rightarrow \text{Mp } (p \& \{t_i. \text{end}\}_{i \in [1,2]})$. Therefore $[p \oplus \bar{p}] :$
 2322 $a@I[p_c] \rightarrow \text{Mp } (p_c \rightarrow \{\bar{p}\}\{t_i. \text{end}\}_{i \in [1,2]}) a@(I[t_i p_c] \cup^{\text{pp}} I[t_i p_c])$.

2325 **Case CASE_i.** $\models e_1 \nabla e_2 \Leftarrow (a_1 + a_2)@(t_i I) \sim G$. Then,

2326 $\models e_i \Leftarrow a_i@I \sim G$. By the IH, $\llbracket e_i \rrbracket (a_i@I) : (a_i@I) \rightarrow \text{Mp } G B$.

2327 But by the definition of $\llbracket \cdot \rrbracket$, $\llbracket e_1 \nabla e_2 \rrbracket ((a_1 + a_2)@(t_i I)) : ((a_1 +$
 2328 $a_2)@(t_i I)) \rightarrow \text{Mp } G B$.

2329 **Case PROJ_i.** $\models \pi_i \Leftarrow A_1 \times A_2 \sim \text{end}$. By definition,

2330 $\llbracket \pi_i \rrbracket (A_1 \times A_2) = [_ \mapsto \lambda x. \text{ret } (\pi_i x)] : A_1 \times A_2 \rightarrow$
 2331 $\text{Mp end } A_i$.

2333 **Case SPLIT.** $\models e_1 \Delta e_2 \Leftarrow A \sim [G_2/\text{end}]G_1$. Then, $\models e_1 \Leftarrow$
 2334 $A \sim G_1$, and $\models e_2 \Leftarrow A \sim G_2$. By the IH, $\llbracket e_1 \rrbracket (A) : A \rightarrow$
 2335 $\text{Mp } G_1 B$ and $\llbracket e_2 \rrbracket (A) : A \rightarrow \text{Mp } G_2 C$. By definition,
 2336 $\llbracket e_1 \Delta e_2 \rrbracket (A) = \llbracket e_1 \rrbracket (A) \Delta \llbracket e_2 \rrbracket (A) :$
 2337 $A \rightarrow \text{Mp } ([G_2/\text{end}]G_1) (B \times C)$. \square

2339 C Extensionality

2340 Each monadic m represents the code for an individual pro-
 2341 cess. The parallel composition of the set of monadic actions
 2342 generated from a PAIg expression $e : A \rightarrow B$, each applied
 2343 to the corresponding value of type $v : A \uparrow p$ represents an
 2344 execution of the parallel algorithm on an input of type a , if
 2345 $a@R = A$. Recall from Sec. 5 that the transitions are of the
 2346 form $\langle P, W \rangle \rightsquigarrow^\ell \langle P', W' \rangle$, where P is an environment that
 2347 contains the code executed by all roles that collaborate to
 2348 compute the parallel algorithm, and W represents the shared
 2349 unbounded buffers used by each pair of participants to com-
 2350 municate. We write w for such buffers, where \emptyset is the empty
 2351 buffer, $v \cdot w$ is the buffer w extended with value v at the
 2352 leftmost position, and $w \cdot v$ is the buffer w extended with
 2353 value v at the rightmost position.

2355 **Definition C.1** (Type buffers). We write $Q = [p_i p_j \rightarrow$
 2356 $q]_{i \in I, j \in I}$, where q is a buffer of types, that can be either \emptyset ,
 2357 $a \cdot q$ or $q \cdot a$. Note that values include singleton types that
 2358 represent labels: $l_i : l_i$. We say that a buffer $w = v_1 \cdots v_n$
 2359 contains types $q = a_1 \cdots a_m$, $w : q$ if: $n = m$ and $v_i : a_i$ for
 2360 all $i \in [1, n]$. We say that $W : Q$ if for all pairs of roles, $p_i p_j$,
 2361 $W(p_i p_j) : Q(p_i p_j)$.

2363 **Theorem 5.1.** [Soundness] Assume $E : \text{Mp } C A$, $m : \text{Mp } L a$
 2364 and $W : Q$. Suppose $\langle E[r \mapsto m], W \rangle \rightsquigarrow^\ell \langle E[r \mapsto m'], W' \rangle$.

2366 Then there exists $\langle C[r \mapsto L], Q \rangle \rightarrow^\ell \langle C[r \mapsto L'], Q' \rangle$ such
 2367 that $W' : Q'$ and $m' : \text{Mp } L' a$.

2368 *Proof.* Straightforward induction on L_i , and case analysis
 2369 on m_i and the rules \rightsquigarrow and \rightarrow , since there is a one-to-one
 2370 correspondence between the rules syntactic constructs in
 2371 Mp and the local types. For \rightsquigarrow we need to take several ϵ
 2372 transitions until communication ℓ happens. \square

2374 **Lemma C.2.** Assume $G, A, B, X : A$, and $f_i : A \uparrow p_i \rightarrow$
 2375 $\text{Mp } (G \uparrow p_i) (B \uparrow p_i)$ for all $i \in I$. Let $P = [p_i \mapsto f_i]_{i \in I}$ then
 2376 there is a unique Y s.t. $P(X) = Y$.

2378 *Proof.* Straightforward consequence of Lemma 5.1, and The-
 2379 orem 3.1 in [27]. We know that the traces for G can only
 2380 differ in the order of the actions, and that this order must
 2381 preserve the dependencies laid out by G . Therefore, there
 2382 the result of any possible execution must respect the data
 2383 dependencies specified by G . \square

2384 **Lemma C.3.** If $(P, W) \Downarrow X$ and $\langle P, W \rangle \rightsquigarrow \langle P', W' \rangle$, then
 2385 $(P', W') \Downarrow X$.

2387 **Theorem 5.3.** [Extensionality] Assume $e \Rightarrow e : a@p \rightarrow b@R$
 2388 and $x : a$ initially at p . If $e \Rightarrow e : a@p \rightarrow b@R$
 2389 also produces y , distributed across R .

2391 *Proof.* We prove the following generalised statement. Let
 2392 $e : a \rightarrow b$ s.t. $e \Rightarrow e : A \rightarrow B$, $x : a$, and \vec{i} s.t. $\delta_A^{\vec{i}}(x)$ is defined.

2393 Then, there is \vec{j} , s.t. $\llbracket e \rrbracket (A)(\delta_A^{\vec{i}}(x)) = \delta_B^{\vec{j}}(\llbracket e \rrbracket x)$. We define
 2394 $\delta_A^{\vec{i}}(x) : A$ as follows:

$$\begin{aligned} \delta_I^\epsilon(x) &= \delta_I(x), \\ \delta_{R_1 \cup \bar{p} R_2}^{i_1 \cdots i_n}(x) &= \text{br}_i^{\bar{p}} (\delta_{R_{i_1}}^{i_2 \cdots i_n}(x)) \\ \delta_p(x) &= [p \mapsto x] \\ \delta_{I_1 \times I_2}(x, y) &= [p \mapsto \delta_{I_1}(x)(p) \times \delta_{I_2}(y)(p)]_{p \in \text{pids}(I_1 \times I_2)} \\ \delta_{t_i I}(t_i x) &= \delta_I(x) \end{aligned}$$

2403 We proceed by induction on the structure of the derivation
 2404 $\vdash e \Rightarrow e : A \rightarrow B$:

- 2405 • Case JOIN. We have $\vdash e \Rightarrow e : A \cup^{\bar{p}} A \rightarrow B \cup^{\bar{p}} B$ with \vdash
 2406 $e \Rightarrow e : A \rightarrow B$. By definition, $\llbracket e \rrbracket (A \cup^{\bar{p}} A) = \llbracket e \rrbracket (A) \uplus^{\bar{p}}$
 2407 $\llbracket e \rrbracket (A)$. We have that $\delta_{A \cup^{\bar{p}} A}^{i \cdot \vec{i}}(x) = \text{bf}_{i_1}^{\bar{p}} (\delta_A^{\vec{i}}(x))$. Then,
 2408 by the induction hypothesis, there exists \vec{j} s.t.
 2409 $\llbracket e \rrbracket (A \cup^{\bar{p}} A)(\delta_{A \cup^{\bar{p}} A}^{i \cdot \vec{i}}(x))$
 2410 $= (\llbracket e \rrbracket (A) \uplus^{\bar{p}} \llbracket e \rrbracket (A)) (\text{br}_i^{\bar{p}} \delta_A^{\vec{i}}(x))$
 2411 $= \text{br}_i^{\bar{p}} \llbracket e \rrbracket (A) (\delta_A^{\vec{i}}(x))$
 2412 $= \text{br}_i^{\bar{p}} (\delta_B^{\vec{j}}(\llbracket e \rrbracket x))$
 2413 $= \delta_{B \cup^{\bar{p}} B}^{i \cdot \vec{j}}(\llbracket e \rrbracket x)$
- 2414 • Case ALT. We have $\vdash e \Rightarrow e : A_1 \cup^{\bar{p}} A_2 \rightarrow B_1 \cup^{\bar{p}} B_2$
 2415 with $\vdash e \Rightarrow e : A_1 \rightarrow B_1$, $\vdash e \Rightarrow e : A_2 \rightarrow B_2$ and $A_1 \neq$
 2416 A_2 . Then, $\llbracket e \rrbracket (A_1 \cup^{\bar{p}} A_2)(\delta_{A_1 \cup^{\bar{p}} A_2}^{i \cdot \vec{i}}(x)) = (\llbracket e \rrbracket (A_1) \uplus^{\bar{p}}$
 2417 $\llbracket e \rrbracket (A_2)) (\delta_{A_1 \cup^{\bar{p}} A_2}^{i \cdot \vec{i}}(x))$

$\llbracket e \rrbracket (A_2)) (\text{br}_i^{\vec{p}} (\delta_{A_i}^{\vec{j}}(x))) = \text{br}_i^{\vec{p}} \llbracket e \rrbracket (A_i) (\delta_{A_i}^{\vec{j}}(x))$. Finally, by the induction hypothesis, there exists \vec{j} s.t. $\text{br}_i^{\vec{p}} \llbracket e \rrbracket (A_i) (\delta_{A_i}^{\vec{j}}(x)) = \text{br}_i^{\vec{p}} \delta^{\vec{j}}(B_i)(\llbracket e \rrbracket x) = \delta^{i \cdot \vec{j}}(B_1 \cup^{\vec{p}} B_2)(\llbracket e \rrbracket x)$.

- Case ALG.** We have $\vdash e \Rightarrow e@p : a@I \rightarrow b@p$, with $\vdash e : a \rightarrow b$. Then, $\llbracket e@p \rrbracket_{a@I} (\delta_I(x)) = [p_i \mapsto (a@I \rightsquigarrow p)] \delta_I(x)$, by straightforward induction on I , there exists a trace $\langle ([p_i \mapsto (a@I \rightsquigarrow p)](p_i)) \gg [p \mapsto \lambda x. \text{ret}(e \ x)] \delta_I(x), W \rangle \rightsquigarrow^{\ell_1 \dots \ell_m} \langle [p_i \mapsto \text{ret } v_i], W' \rangle$, with $v_i = ()$ for all i s.t. $p_i \neq p$, and $v_j = \llbracket e \rrbracket x$ for $p_j = p$. By Theorem 5.2, the only possible interleavings of actions of $\llbracket e@p \rrbracket$ must follow the protocol $[a@I \rightsquigarrow p]$. Since this implies that send/receive operations must happen respecting the data dependencies, any possible trace must yield the same result.
- Case INJ.** $\vdash \iota_i \Rightarrow \iota_i : A \rightarrow \iota_i \ A$ straightforward since $\llbracket \iota_i \rrbracket (A) (\delta_A(x)) = \delta_{\iota_i \ A}(\iota_i \ x)$.
- Case ID.** $\vdash \text{id} \Rightarrow \text{id} : A \rightarrow A$ straightforward, since $\llbracket \text{id} \rrbracket (A) (\delta_A(x)) = \delta_A(\text{id } x)$.
- Case PROJ.** $\vdash \pi_i \Rightarrow \pi_i : A_1 \times A_2 \rightarrow A_i$ straightforward, since $\llbracket \pi_i \rrbracket (A_1 \times A_2) (\delta_{A_1 \times A_2}(x)) = \delta_{A_i}(\pi_i \ x)$.
- Case COMP.** $\vdash e_1 \circ e_2 \Rightarrow e_1 \circ e_2 : A \rightarrow C$ with $\vdash e_1 \Rightarrow e_1 : B \rightarrow C$ and $\vdash e_2 \Rightarrow e_2 : A \rightarrow B$. A straightforward consequence of Theorem 5.2 is that if E_1 behaves as G_1 and E_2 as G_2 , then $(E_1 \ ; \ E_2)(X) = E_2(E_1(X))$, since the permutations of actions of $E_1 \ ; \ E_2$ must respect $G_1 \ ; \ G_2$. Then, by the definition of $\llbracket \cdot \rrbracket$, $\llbracket e_1 \circ e_2 \rrbracket (A) (\delta_A^{\vec{j}}(x)) = \llbracket e_1 \rrbracket (B) (\llbracket e_2 \rrbracket (A) (\delta_A^{\vec{j}}(x)))$ By the induction hypothesis: $\llbracket e_1 \rrbracket (B) (\llbracket e_2 \rrbracket (A) (\delta_A^{\vec{j}}(x))) = \llbracket e_1 \rrbracket_B (\delta_B^{\vec{j}}(\llbracket e_2 \rrbracket x)) = \delta_C^{\vec{k}}(\llbracket e_1 \rrbracket (\llbracket e_2 \rrbracket x)) = \delta_C^{\vec{k}}(\llbracket e_1 \circ e_2 \rrbracket x)$.
- Case CASE_i.** We have $\vdash e_1 \nabla e_2 \Rightarrow e_1 \nabla e_2 : \iota_i \ A \rightarrow B$, with $\vdash e_i \Rightarrow e_i : A \rightarrow B$. Note that $(\delta_{\iota_i \ A}(x))$ is only defined if $x = \iota_i \ x'$. Then, by definition, $\llbracket e_1 \nabla e_2 \rrbracket (\iota_i \ A) (\delta_{\iota_i \ A}(\iota_i \ x')) = \llbracket e_i \rrbracket (A) (\delta_A(x'))$. By the IH, $\llbracket e_i \rrbracket (A) (\delta_A(x')) = \delta(B)(\llbracket e_i \rrbracket x') = \delta(B)(\llbracket e_1 \nabla e_2 \rrbracket (\iota_i \ x')) = \delta(B)(\llbracket e_1 \nabla e_2 \rrbracket x)$.
- Case SPLIT.** We have $\vdash e_1 \Delta e_2 \Rightarrow e_1 \Delta e_2 : A \rightarrow B \times C$, with $\vdash e_2 \Rightarrow e_2 : A \rightarrow C$ and $\vdash e_1 \Rightarrow e_1 : A \rightarrow B$. By definition, $\llbracket e_1 \Delta e_2 \rrbracket (A) = \llbracket e_1 \rrbracket (A) \Delta \llbracket e_2 \rrbracket (A)$. By Theorem 5.2, we know that this behaves as $G_1 \ ; \ G_2$, if $p_1 \sim G_1$ and $p_2 \sim G_2$. Therefore, we assume again that the interleavings of the subtraces must not affect the data dependencies. Then, $(\llbracket e_1 \rrbracket (A) (\delta_A(x))) \Delta (\llbracket e_2 \rrbracket (A) (\delta_A(x))) = \delta_B^{\vec{j}}(\llbracket e_1 \rrbracket x) \Delta \delta_C^{\vec{k}}(\llbracket e_2 \rrbracket x) = \delta_{B \times C}^{\vec{j} \cdot \vec{k}}(\llbracket e_1 \Delta e_2 \rrbracket x)$.
- Case CHOICE.** We have $\vdash e \Rightarrow [p \oplus \vec{p}] : a@I[p] \rightarrow a@(\mathcal{I}[\iota_1 \ p] \cup^{\vec{p}} \mathcal{I}[\iota_2 \ p])$. We have two cases:
 - $p \mapsto \lambda x. \text{sel}(\text{get}_I(x)) \{\vec{p}\}(\lambda y. \text{put}_I(y, x))(\lambda y. \text{put}_I(y, x))$

$\forall p', p' \neq p \wedge p' \in \vec{p}, p' \mapsto \lambda x. \text{brn } p (\text{ret } x) (\text{ret } x)$ 2476
 By case analysis, if $\text{get}_I(x) = \iota_i \ v$, then we have: 2477

- $p \mapsto \text{br}_i(\text{put}_I(v, x))$ 2478
- $\forall p', p' \neq p \wedge p' \in \vec{p}, p' \mapsto \text{ret}(\text{br}_i \ x)$ 2479

 This is clearly $[p \oplus \vec{p}] \delta_{\mathcal{I}[p]}(x) = \text{br}_i^{\vec{p}} \delta_{\mathcal{I}[\iota_1 \ p]}(x) = \delta_{\mathcal{I}[\iota_1 \ p] \cup^{\vec{p}} \mathcal{I}[\iota_2 \ p]}^i(x)$. 2480–2482

□ 2483

D Generated Code

We show now the generated code for mergesort, unrolling the recursive function once.

D.1 Input Alg expression

```

fftTree :: forall n. SINat n -> Int
    -> Tree n (D [Complex Double])
    :=> Tree n (D [Complex Double])
fftTree SZ w
    = lift (intlit SZ &&& (lit w &&& id)
    >>> prim "baseFFT")
fftTree (SS x) w
    = withCDict (cdictTree @(D [Complex Double]) x)
    $
    {- Recursive FFT to EVENS and ODDS-}
    (fftTree x w
    *** fftTree x (w + 2^ toInteger x))
    {- Multiply right side by exponential -}
    >>> id
    *** mapTree x (lit ps2x &&& id >>> mapExp)
    0
    >>>
    {- zipWith add (swap arguments to force butterfly
    pattern
    - &&& zipWith sub
    -}
    zipTree x True lvl w addc
    &&& zipTree x False lvl (w + 2^ toInteger x)
    subc
    where
    lvl :: Int
    lvl = fromInteger (toInteger (SS x) + 1)
    ps2x :: Int
    ps2x = 2 ^ toInteger (SS x)
fft :: SINat n -> (D [Complex Double]) :=> D [
    Complex Double]
fft n =
    withCDict (cdictTree @(D [Complex Double]) n)
    $
    tsplit n deinterleave >>> fftTree n 0 >>> tfold
    n (append @@ 0)
fft5 :: D [Complex Double] :=> D [Complex
    Double]
    
```

```
2531     fft5 = withSize 5 fft
```

Listing 1. Fragment of FFT.hs

D.2 Main C Code and Atomic Functions

These need to be implemented by the programmer.

```
2538     #include "FFT.h"
2539     #include <inttypes.h>
2540     #include <errno.h>
2541     #include <string.h>
2542     #include <sys/time.h>
2543     #include <stdlib.h>
2544     #include <math.h>
2545
2546     #define REPETITIONS 50
2547
2548     #define BENCHMARKSEQ(s, f) { \
2549         time = 0; \
2550         time_diff = 0; \
2551         time_old = 0; \
2552         var = 0; \
2553         for(int i=0; i<REPETITIONS; i++){ \
2554             in = randvec(s, size); \
2555             start = get_time(); \
2556             out = f(in); \
2557             end = get_time(); \
2558             free_fftvec(in); \
2559             time_diff = end - start; \
2560             time_old = time; \
2561             time += (time_diff - time)/(i+1); \
2562             var += (time_diff - time) * (time_diff - \
2563                 time_old); \
2564         } \
2565         printf("\tK: %d\n", s); \
2566         printf("\t\tmean: %f\n", time); \
2567         printf("\t\tstddev: %f\n", REPETITIONS<=1? 0: sqrt(\
2568             var / (REPETITIONS - 1))); \
2569     }
2570
2571     #define WARMUP(f) { \
2572         for(int i=0; i<REPETITIONS; i++){ \
2573             in = randvec(0, size); \
2574             out = f(in); \
2575             free_fftvec(in); \
2576         } \
2577     }
2578
2579     double PI = atan2(1, 1) * 4;
2580
2581     int num_stages;
2582     int num_workers;
2583     vec_cplx_t **stages;
2584
2585
```

```
vec_cplx_t zip_add(                2586
    pair_pair_int_int_pair_vec_cplx_vec_cplx_t in)  2587
{
    2588     int lvl = in.fst.fst;          2589
    2590     int wid = in.fst.snd;          2590
    2591     vec_cplx_t l = in.snd.fst;     2591
    2592     vec_cplx_t r = in.snd.snd;     2592
    2593     vec_cplx_t lout = stages[lvl][wid]; 2593
    2594     for(int i = 0; i < l.size; i++){ 2594
    2595         lout.elems[i] = l.elems[i] + r.elems[i]; 2595
    2596     }                               2596
    2597     return lout;                   2597
}                                    2598
}                                    2599
vec_cplx_t zip_sub(                2600
    pair_pair_int_int_pair_vec_cplx_vec_cplx_t in)  2601
{
    2602     int lvl = in.fst.fst;          2603
    2603     int wid = in.fst.snd;          2604
    2604     vec_cplx_t l = in.snd.fst;     2605
    2605     vec_cplx_t r = in.snd.snd;     2606
    2606     vec_cplx_t lout = stages[lvl][wid]; 2607
    2607     for(int i = 0; i < l.size; i++){ 2608
    2608         lout.elems[i] = l.elems[i] - r.elems[i]; 2609
    2609     }                               2610
    2610     return lout;                   2611
}                                    2612
}                                    2613
vec_cplx_t cat(pair_vec_cplx_vec_cplx_t in){ 2614
    2615     in.fst.size *= 2;              2615
    2616     return in.fst;                 2616
}                                    2617
}                                    2618
void _fft(cplx_t buf[], cplx_t out[], int n, int  2619
    step)                             2620
{
    2621     if (step < n) {                2622
    2622         _fft(out, buf, n, step * 2); 2623
    2623         _fft(out + step, buf + step, n, step * 2); 2624
    2624     }                               2625
    2625     for (int i = 0; i < n; i += 2 * step) { 2626
    2626         cplx_t t = cexp(-I * PI * i / n) * out[i + 2627
            step];                    2628
    2627         buf[i / 2] = out[i] + t;     2629
    2628         buf[(i + n)/2] = out[i] - t; 2630
    2629     }                               2631
    2630     }                               2632
    2631     }                               2633
}                                    2634
}                                    2635
void show(const char * s, vec_cplx_t in) { 2636
    2637     printf("%s", s);                2636
    2638     for (int i = 0; i < in.size; i++) 2637
    2639         if (!cimag(in.elems[i]))      2638
    2640             printf("%g", creal(in.elems[i])); 2639
    2640
    2640
```



```

2641     else } 2696
2642     printf("(%g, %g) ", creal(in.elems[i]), cimag( free(stages); 2697
2643     in.elems[i])); } 2698
2644     printf("\n"); 2699
2645 } 2700
2646 void showstep(int stp, const char * s, vec_cplx_t 2701
2647 in) { pair_int_int_t iin){ 2702
2648     printf("%s", s); int wl = iin.fst; 2703
2649     for (int i = 0; i < in.size; i+=stp) int wr = iin.snd; 2704
2650     if (!cimag(in.elems[i])) int mid = stages[0][wl].size/2; 2705
2651     printf("%g", creal(in.elems[i])); stages[1][wl].size = mid; 2706
2652     else stages[1][wr].size = mid; 2707
2653     printf("(%g, %g) ", creal(in.elems[i]), cimag( stages[1][wr].elems = stages[1][wl].elems + mid; 2708
2654     in.elems[i])); 2709
2655     printf("\n"); 2710
2656 } for(int i = 0; i < stages[0][wl].size; i+= 2){ 2711
2657     stages[1][wl].elems[i/2] = stages[0][wl].elems 2712
2658     [i]; stages[1][wr].elems[i/2] = stages[0][wl].elems 2713
2659     [i+1]; 2714
2660 } memcopy(stages[0][wl].elems, stages[1][wl].elems, 2715
2661 { stages[0][wl].size * sizeof(cplx_t)); 2716
2662     int lvl = in.fst; stages[0][wr].elems = stages[0][wl].elems + mid; 2717
2663     int wid = in.snd.fst; stages[0][wr].size = mid; 2718
2664     cplx_t *buf = stages[lvl][wid].elems; 2719
2665     int n = in.snd.snd.size; 2720
2666     _fft(buf, in.snd.snd.elems, n, 1); 2721
2667     return stages[lvl][wid]; 2722
2668 } memcopy(stages[i][wl].elems, stages[1][wl]. 2723
2669     elems, stages[0][wl].size * sizeof(cplx_t)) 2724
2670 ; 2725
2671     stages[i][wl].size = mid; 2726
2672     stages[i][wr].elems = stages[i][wl].elems + 2727
2673     mid; 2728
2674     stages[i][wr].size = mid; 2729
2675     } 2730
2676     stages[0][wl].size = mid; 2731
2677     return (pair_vec_cplx_vec_cplx_t) { stages[0][wl] 2732
2678     ], stages[0][wr] }; 2733
2679 } 2734
2680 vec_cplx_t randvec(int depth, size_t s){ 2735
2681     num_workers = depth <= 1? 1 : 1 << depth - 1; 2736
2682     num_stages = depth <= 1? 2 : 1 + depth ; 2737
2683     stages = (vec_cplx_t **)malloc(num_stages * 2738
2684     sizeof(vec_cplx_t *)); 2739
2685     for (int i = 0; i < num_stages; i++){ 2740
2686     stages[i] = (vec_cplx_t *)malloc(num_workers * 2741
2687     sizeof(vec_cplx_t)); 2742
2688     stages[i][0].elems = (cplx_t *)calloc(s, sizeof 2743
2689     (cplx_t)); 2744
2690     } 2745
2691     stages[0][0].size = s; 2746
2692     } 2747
2693     } 2748
2694     } 2749
2695     } 2750

```

```

2751     srand(time(NULL));
2752
2753
2754     for (int i = 0; i < s; i++) {
2755         double rand_r = (double)rand() / (double)
2756             RAND_MAX;
2757         double rand_i = (double)rand() / (double)
2758             RAND_MAX;
2759         stages[0][0].elems[i] = rand_r + rand_i * I;
2760     }
2761
2762     for(int j = 1; j < num_stages; j++) {
2763         memcpy(stages[j][0].elems, stages[0][0].
2764             elems, s * sizeof(vec_cplx_t));
2765         stages[j][0].size = s / num_workers;
2766     }
2767
2768     for(int i = 0; i < num_stages; i++) {
2769         for(int j = 1; j < num_workers; j++) {
2770             stages[i][j] = stages[i][j-1];
2771         }
2772     }
2773
2774     return stages[0][0];
2775 }
2776
2777 void usage(const char *nm){
2778     printf("Usage: %s <input_size>\n", nm);
2779     exit(-1);
2780 }
2781
2782 int main(int argc, const char *argv[]) {
2783     setbuf(stdout, NULL);
2784     if (argc <= 1) {
2785         usage(argv[0]);
2786     }
2787     char *endptr = NULL;
2788     errno = 0;
2789     size_t size = strtoumax(argv[1], &endptr, 10);
2790     size = (size_t) 1 << (long)ceil(log2(size));
2791     size = size < 256? 256:size;
2792     if (errno != 0) {
2793         printf("%s", strerror(errno));
2794         usage(argv[0]);
2795     }
2796     if (endptr != NULL && *endptr != 0) {
2797         usage(argv[0]);
2798     }
2799
2800     vec_cplx_t in, out;
2801     /* allocate memory */
2802     in = randvec(size, size);
2803
2804     /* calling generated fft5 */
2805

```

```

out = fft5(in);
show("Result: ", out);

free_fftvec();
}

```

D.3 Automatically Generated C Code

```

#ifndef __FFT__
#define __FFT__

#include<stdio.h>
#include<stdlib.h>
#include<pthread.h>
#include<complex.h>
typedef double _Complex cplx_t;

typedef struct vec_cplx {
    cplx_t * elems; size_t size;
} vec_cplx_t;

typedef struct q_vec_cplx {
    volatile unsigned int q_size;
    int q_head;
    int q_tail;
    pthread_mutex_t q_mutex;
    pthread_cond_t q_full;
    pthread_cond_t q_empty;
    vec_cplx_t q_mem[1];
} q_vec_cplx_t;

void q_vec_cplx_put(q_vec_cplx_t *, vec_cplx_t);
vec_cplx_t q_vec_cplx_get(q_vec_cplx_t *);

typedef enum unit {
    Unit
} unit_t;

typedef struct pair_int_vec_cplx {
    int fst; vec_cplx_t snd;
} pair_int_vec_cplx_t;

typedef struct pair_int_pair_int_vec_cplx {
    int fst; pair_int_vec_cplx_t snd;
} pair_int_pair_int_vec_cplx_t;

vec_cplx_t baseFFT(pair_int_pair_int_vec_cplx_t);
vec_cplx_t fft0(vec_cplx_t);
vec_cplx_t fft1(vec_cplx_t);

typedef struct pair_int_int {

```

```

2861         int fst; int snd;
2862     } pair_int_int_t;
2863
2864     typedef struct pair_vec_cplx_vec_cplx {
2865         vec_cplx_t fst; vec_cplx_t snd;
2866     } pair_vec_cplx_vec_cplx_t;
2867
2868     pair_vec_cplx_vec_cplx_t deinterleave(
2869         pair_int_int_t);
2870
2871     vec_cplx_t cat(pair_vec_cplx_vec_cplx_t);
2872
2873     typedef struct
2874     pair_pair_int_int_pair_vec_cplx_vec_cplx {
2875         pair_int_int_t fst;
2876         pair_vec_cplx_vec_cplx_t snd;
2877     }
2878     pair_pair_int_int_pair_vec_cplx_vec_cplx_t
2879     ;
2880
2881     vec_cplx_t zip_add(
2882         pair_pair_int_int_pair_vec_cplx_vec_cplx_t);
2883
2884     vec_cplx_t map_exp(pair_int_pair_int_vec_cplx_t);
2885
2886     vec_cplx_t zip_sub(
2887         pair_pair_int_int_pair_vec_cplx_vec_cplx_t);
2888
2889     vec_cplx_t fft2(vec_cplx_t);
2890
2891     vec_cplx_t fft3(vec_cplx_t);
2892
2893     vec_cplx_t fft4(vec_cplx_t);
2894
2895     vec_cplx_t fft5(vec_cplx_t);
2896
2897     vec_cplx_t fft6(vec_cplx_t);
2898
2899     vec_cplx_t fft7(vec_cplx_t);
2900
2901     vec_cplx_t fft8(vec_cplx_t);
2902
2903     #endif
2904
2905     Listing 2. Generated FFT.h
2906
2907     #include "FFT.h"
2908
2909     q_vec_cplx_t ch0 = { 0, 0, 0, { } };
2910     q_vec_cplx_t ch2 = { 0, 0, 0, { } };
2911
2912     q_vec_cplx_t ch3 = { 0, 0, 0, { } };
2913
2914     vec_cplx_t fft2_part_0(vec_cplx_t v_s)
2915
2916     {
2917         pair_int_int_t v_t;
2918         v_t.fst = 0;
2919         v_t.snd = 1;
2920         pair_vec_cplx_vec_cplx_t v_u;
2921         v_u = deinterleave(v_t);
2922         vec_cplx_t v_v;
2923         v_v = v_u.fst;
2924         q_vec_cplx_put(&ch0, v_v);
2925         vec_cplx_t v_w;
2926         v_w = v_u.snd;
2927         q_vec_cplx_put(&ch2, v_w);
2928         vec_cplx_t v_x;
2929         v_x = q_vec_cplx_get(&ch1);
2930         vec_cplx_t v_y;
2931         v_y = q_vec_cplx_get(&ch3);
2932         pair_vec_cplx_vec_cplx_t v_z;
2933         v_z.fst = v_x;
2934         v_z.snd = v_y;
2935         vec_cplx_t v_aa;
2936         v_aa = cat(v_z);
2937         return v_aa;
2938     }
2939
2940     q_vec_cplx_t ch4 = { 0, 0, 0, { } };
2941
2942     q_vec_cplx_t ch5 = { 0, 0, 0, { } };
2943
2944     unit_t fft2_part_1()
2945     {
2946         vec_cplx_t v_ba;
2947         v_ba = q_vec_cplx_get(&ch0);
2948         pair_int_pair_int_vec_cplx_t v_ca;
2949         v_ca.fst = 1;
2950         pair_int_vec_cplx_t v_da;
2951         v_da.fst = 0;
2952         v_da.snd = v_ba;
2953         v_ca.snd = v_da;
2954         vec_cplx_t v_ea;
2955         v_ea = baseFFT(v_ca);
2956         q_vec_cplx_put(&ch4, v_ea);
2957         vec_cplx_t v_fa;
2958         v_fa = q_vec_cplx_get(&ch5);
2959         pair_pair_int_int_pair_vec_cplx_vec_cplx_t
2960             v_ga;
2961         pair_int_int_t v_ha;
2962         v_ha.fst = 2;
2963         v_ha.snd = 0;
2964         v_ga.fst = v_ha;
2965         pair_vec_cplx_vec_cplx_t v_ia;
2966         v_ia.fst = v_ea;
2967         v_ia.snd = v_fa;
2968         v_ga.snd = v_ia;
2969         vec_cplx_t v_ja;
2970

```

```

2971     v_ja = zip_add(v_ga);
2972     q_vec_cplx_put(&ch1, v_ja);
2973     return Unit;
2974 }
2975
2976 unit_t fft2_part_2()
2977 {
2978     vec_cplx_t v_ka;
2979     v_ka = q_vec_cplx_get(&ch2);
2980     pair_int_pair_int_vec_cplx_t v_la;
2981     v_la.fst = 1;
2982     pair_int_vec_cplx_t v_ma;
2983     v_ma.fst = 1;
2984     v_ma.snd = v_ka;
2985     v_la.snd = v_ma;
2986     vec_cplx_t v_na;
2987     v_na = baseFFT(v_la);
2988     pair_int_pair_int_vec_cplx_t v_oa;
2989     v_oa.fst = 2;
2990     pair_int_vec_cplx_t v_pa;
2991     v_pa.fst = 0;
2992     v_pa.snd = v_na;
2993     v_oa.snd = v_pa;
2994     vec_cplx_t v_qa;
2995     v_qa = map_exp(v_oa);
2996     q_vec_cplx_put(&ch5, v_qa);
2997     vec_cplx_t v_ra;
2998     v_ra = q_vec_cplx_get(&ch4);
2999     pair_pair_int_int_pair_vec_cplx_vec_cplx_t
3000         v_sa;
3001     pair_int_int_t v_ta;
3002     v_ta.fst = 2;
3003     v_ta.snd = 1;
3004     v_sa.fst = v_ta;
3005     pair_vec_cplx_vec_cplx_t v_ua;
3006     v_ua.fst = v_ra;
3007     v_ua.snd = v_qa;
3008     v_sa.snd = v_ua;
3009     vec_cplx_t v_va;
3010     v_va = zip_sub(v_sa);
3011     q_vec_cplx_put(&ch3, v_va);
3012     return Unit;
3013 }
3014
3015 void * fun_thread_1_1(void * arg)
3016 {
3017     fft2_part_1();
3018     return NULL;
3019 }
3020
3021 void * fun_thread_2(void * arg)
3022 {
3023     fft2_part_2();
3024     return NULL;
3025

```

```

}
vec_cplx_t fft2(vec_cplx_t v_wa)
{
    vec_cplx_t v_xa;
    pthread_t thread1;
    pthread_t thread2;
    pthread_create(&thread1, NULL, fun_thread_1_1,
        NULL);
    pthread_create(&thread2, NULL, fun_thread_2,
        NULL);
    v_xa = fft2_part_0(v_wa);
    pthread_join(thread1, NULL);
    pthread_join(thread2, NULL);
    return v_xa;
}

```

Listing 3. Fragment of generated FFT.c

E Artifact Appendix

E.1 Abstract

This artifact provides a prototype implementation of PAI_g, embedded in Haskell, along with a number of benchmarks used to test the scalability of our approach. We provide scripts to regenerate the execution time measurements that we used in our paper. This will allow to evaluate our results on any multi-core shared-memory architecture.

We also provide a small tutorial that is meant to guide a programmer, step-by-step, in the implementation of a message-passing parallel algorithm using our library. The tutorial includes a guide on how to visualise the global types that correspond to the achieved parallelisations, as well as any asynchronous optimisations applicable to the generated message-passing code.

E.2 Artifact check-list (meta-information)

- **Algorithm:** Message-passing C code generation from first-order Haskell functions. Global type inference of the communication protocol followed by the parallelisation.
- **Program:** Haskell libraries Language.CA_{lg}, noindent Language.CA_{lg}.CSyn and dependencies, as well as session-arrc, to compile to C Haskell functions built using such libraries.
- **Compilation:** GHC ≥ 8.6 && < 8.8 , and C compiler that supports C11.
- **Transformations:** Compilation to C, and asynchronous optimisation pass.
- **Binary:** Source code and scripts included to generate the binaries from the sources.
- **Data set:** Included original run-time measurements for comparison.
- **Hardware:** We used a 12-core Intel Xeon CPU E5-2650 v4 @ 2.20GHz. We recommend a shared-memory architecture, with uniform access times, to measure the overheads of our approach, not message latencies.

- 3081 • **Execution:** We include a script to run the benchmarks.
- 3082 • **Output:** Benchmark execution times.
- 3083 • **Experiments:** Small, representative benchmarks of com-
3084 mon parallel algorithms.
- 3085 • **How much memory required (approximately)?:** 64GB
3086 for using the maximum benchmark input size.
- 3087 • **How much time is needed to complete experiments**
3088 **(approximately)?:** 5 days on the hardware stated in §E.3.2.
- 3089 • **Publicly available?:** Yes.
- 3090 • **Code licenses (if publicly available)?:** BSD-3.

3091 E.3 Description

3092 E.3.1 How delivered

3093 We provide a docker image with the necessary dependencies: <https://imperialcollegelondon.box.com/v/cc20-artifact-p43>. After down-
3094 loading, the image can be loaded using:

```
3095 $ sudo docker load -i cc20-artifact-p43.docker
```

3097 To run the image, run the command:

```
3098 $ sudo docker run -ti cc20-artifact-p43
```

3100 File README.md inside the docker image contains additional in-
3101 structions. Our benchmarks, source code and scripts are also pub-
3102 licly available on Github, in <https://github.com/session-arr/session-arr>.
3103

3104 E.3.2 Hardware dependencies

3105 We used a 12-core Intel Xeon CPU E5-2650 v4 @ 2.20GHz. We
3106 recommend using a shared-memory architecture, with uniform ac-
3107 cess times, to measure the overheads of our approach, not message
3108 latencies.

3109 E.3.3 Software dependencies

3111 All our dependencies are listed in the Dockerfile in our public
3112 repository. We list them below. To compile our tool:

- 3113 1. GHC >= 8.6 (not tested with GHC >= 8.8)
- 3114 2. stack Version 1.9.1

3115 To run our experiments:

- 3116 1. C compiler that supports C11 (tested with GCC >= 4.8 && <
3117 8.3)
- 3118 2. glibc (tested with versions >= 2.17 && < 2.29)
- 3119 3. numactl

3120 To generate the graphs:

- 3121 1. python (== 2.7)
- 3122 2. python-matplotlib (== 2)
- 3123 3. python-pint (== 0.7)

3124 E.3.4 Data sets

3125 We include as part of the artifact the raw data that we obtained for
3126 our benchmarks. These are included under
3127 benchmarks/<bench_name>/data/t_<num_cores>, where
3128 <num_cores> is either 12 or 24. There is additionally a file t_48,
3129 that uses all full 24 cores + hyperthreading. The structure of the
3130 files is:

```

size: <size> 3136
      K: seq 3137
          mean: <avg_execution_time> 3138
          stddev: <std_dev> 3139
      K: 1 3140
          mean: ... 3141
          stddev: ... 3142
      ... 3143
    
```

3144 Keyword size denotes the size of the inputs for the particular
3145 benchmark. Keyword mean is the average execution time. Keyword
3146 stddev is the standard deviation. We write K: to denote the number
3147 of recursion unfoldings used to produce the parallel version.

3148 Examples of global types for each benchmark are under
3149 benchmarks/<bench_name>/protocol/
3150 <bench_name>_<fun_name>.mpst, where <fun_name> is the func-
3151 tion name in <bench_name>.hs that corresponds to this protocol.

3152 E.4 Installation

3153 **Note:** this section can be omitted if using our docker image.

3154 We recommend using Stack ([https://docs.haskellstack.org/en/stable/](https://docs.haskellstack.org/en/stable/README/#how-to-install)
3155 README/#how-to-install). To build our tool:

```

3156 $ git clone \ 3157
3158     https://github.com/session-arr/session-arr 3159
3160 $ cd session-arr 3161
3162 $ stack build 3163
    
```

3164 There is no need to install the tool. However, to install it, run:

```
3165 $ stack install 3166
```

3167 This will copy the binary session-arrc to a local directory,
3168 usually \${HOME}/.local/bin.

3169 Manual compilation and installation using GHC is also possible,
3170 but we discourage it. Read session-arr/package.yaml to find
3171 out which haskell packages are required.

3172 E.5 Experiment workflow

3173 E.5.1 Automatic

3174 We included script session-arr/benchmark.sh to compile and
3175 run all the benchmarks used in the paper. To customise the amount
3176 of cores, the number of repetitions per experiment and the maxi-
3177 mum input size, run:

```

3178 $ CORES=<ncores> REPETITIONS=<nreps> \ 3179
3180     MAXSIZE=<nsize> ./benchmark.sh 3181
    
```

3182 The defaults are:

- 3183 1. CORES: number of physical cores on your machine
- 3184 2. REPETITIONS: 50
- 3185 3. MAXSIZE: 30

3186 The script requires that MAXSIZE ≥ 15.

3187 **Note:** using MAXSIZE= 30 requires a machine with a large amount
3188 of memory. We run our experiments on a machine with 64GB of
3189 memory.

```

3191 E.5.2 Manual
3192 Clone and build the repository:
3193 $ git clone \
3194     https://github.com/session-arr/session-arr
3195 $ cd session-arr
3196 $ stack build
3197
3198 Navigate to one of the benchmarks
3199 $ cd examples/FFT
3200 Here, there should be two files: FFT.hs and main.c.
3201 $ ls
3202 FFT.hs main.c run.sh
3203 To run our tool, run session-arrc using stack, with the .hs file
3204 as input.
3205 $ stack exec session-arrc -- FFT.hs
3206
3207 The tool should output the list of functions found in module FFT.hs
3208 that are going to be compiled to C, and produce two files FFT.c and
3209 FFT.h. The interface file contains the type definitions and function
3210 signatures of the functions in FFT.c. Finally, compile main.c:
3211 $ gcc FFT.c main.c -o bench -lpthread -lm
3212 To configure the number of repetitions, recompile the benchmark
3213 as follows:
3214 $ gcc FFT.c main.c -DREPETITIONS=<num_reps> \
3215     -o bench -lpthread -lm
3216
3217 You may use run.sh to run the benchmark on a range of inputs.
3218 The usage is:
3219 $ ./run.sh <num_cores> <max_size>
3220 For example, ./run.sh 2 10 will run the benchmark with sizes
3221 29 and 210. The maximum size must be > 9. To generate the
3222 graphs, you need measurements using at least 7 different sizes, i.e.
3223 size must be > 14.
3224 Running each benchmark manually Pass a valid input size to
3225 bench, the output looks as follows (run in a 4-core machine):
3226 $ ./bench $((2**17))
3227 K: seq
3228     mean: 0.039446
3229     stddev: 0.000713
3230     ...
3231 K: 4
3232     mean: 0.011952
3233     stddev: 0.000636
3234     ...
3235 Save all execution times to files with the format described in §E.3.4,
3236 as follows:
3237 $ mkdir data
3238 $ echo "size: <size1>" >> data/t_<num_cores>
3239 $ ./bench <size1> >> data/t_<num_cores>
3240 $ ...
3241 $ echo "size: <sizeN>" >> data/t_<num_cores>
3242 $ ./bench <sizeN> >> data/t_<num_cores>
3243
3244
3245

```

Ensure that there are measurements with at least $N > 14$ sizes. 3246

Plotting the speedups: Navigate to examples/. The speedups can 3247
be plotted using scripts plotall.sh and plot.py, these will re- 3248
generate the graphs used in our paper. The usage is 3249
./plotall.sh BENCHMARK_DIR CORES, where CORES is the number 3250
of cores used for the experimental workflow. For example: 3251

```
$ ./plotall.sh FFT 4 3252
```

This will generate the graphs for FFT run on 4 cores under 3253
examples/plots. 3254
3255

E.6 Evaluation and expected result 3256

If you followed the experiment workflow, you should find under 3257
examples/plots a series of graphs with the speedups for each 3258
benchmark. To visualise them, we recommend copying them to a 3259
local directory, by running docker cp from **outside** the docker 3260
container: 3261

```
$ docker cp \ 3262
<NM>:/home/cc20-artifact/session-arr/examples/plots \ 3263
<DIR> 3264
```

Here, <NM> is the container name obtained via docker ps -a, and 3265
<DIR> is the destination path. 3266

Outcome When run on similar hardware to the one that we de- 3267
scribe in the paper, following our workflow, comparable speedups 3268
and scalability to the ones that we reported in the paper should be 3269
observed. 3270
3271

Note: for more reliable results, execution should be done **outside** 3272
the docker container. Use the container to generate all C code, then 3273
copy it running docker cp from outside the container, as well as 3274
the necessary scripts run.sh, and proceed locally. If you decide to 3275
run the experiments locally, please check §E.3.3 and ensure that 3276
you have all required software. 3277

E.7 Experiment customization 3278

Several aspects can be customised in the benchmark source code, 3279
execution scripts and compilation options: 3280

- Annotations to functions in the .hs files should produce 3281
different parallelisations. 3282
- Files main.c can be compiled using different numbers of 3283
repetitions using -DREPETITIONS=<num_reps>. 3284
- The script benchmark.sh can be run with such number of 3285
repetitions, to reduce execution times. The maximum input 3286
size for the benchmarks, and the number of cores can also 3287
be customised: 3288

```
$ cd session-arr 3289
$ REPETITIONS=<num_reps> CORES=<cores> \ 3290
  MAXSIZE=<max_size> ./benchmark.sh 3291
```

E.8 Methodology 3292

Submission, reviewing and badging methodology: 3293

- <http://cTuning.org/ae/submission-20190109.html> 3294
- <http://cTuning.org/ae/reviewing-20190109.html> 3295
- [https://www.acm.org/publications/policies/ 3296](https://www.acm.org/publications/policies/artifact-review-badging)
artifact-review-badging 3297