



# Kent Academic Repository

**Castro-Perez, David and Yoshida, Nobuko (2023) *Dynamically Updatable Multiparty Session Protocols*. In: European Conference on Object-Oriented Programming. . Dagstuhl, Saarbrücken (In press)**

## Downloaded from

<https://kar.kent.ac.uk/100746/> The University of Kent's Academic Repository KAR

## The version of record is available from

## This document version

Author's Accepted Manuscript

## DOI for this version

## Licence for this version

CC BY (Attribution)

## Additional information

## Versions of research works

### Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

### Author Accepted Manuscripts


If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in **Title of Journal**, Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

### Enquiries

If you have questions about this document contact [ResearchSupport@kent.ac.uk](mailto:ResearchSupport@kent.ac.uk). Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).

# Dynamically Updatable Multiparty Session Protocols

## Generating Concurrent Go Code from Unbounded Protocols

David Castro-Perez ✉ 

University of Kent, UK

Nobuko Yoshida ✉ 

University of Oxford, UK

---

### Abstract

---

Multiparty Session Types (MPST) are a typing disciplines that guarantee the absence of deadlocks and communication errors in concurrent and distributed systems. However, existing MPST frameworks do not support protocols with *dynamic unbounded participants*, and cannot express many common programming patterns that require the introduction of new participants into a protocol. This poses a barrier for the adoption of MPST in languages that favour the creation of new participants (processes, lightweight threads, etc) that communicate via message passing, such as Go or Erlang.

This paper proposes *Dynamically Updatable Multiparty Session Protocols*, a new MPST theory (DMst) that supports protocols with an *unbounded* number of fresh participants, whose communication topologies are *dynamically updatable*. We prove that DMst guarantees deadlock-freedom and liveness. We implement a toolchain, GoScr (Go-Scribble), which generates Go implementations from DMst, ensuring **by construction**, that the different participants will only perform I/O actions that comply with a given protocol specification. We evaluate our toolchain by (1) implementing representative parallel and concurrent algorithms from existing benchmarks, textbooks and literature; (2) showing that GoScr does not introduce significant overheads compared to a naive implementation, for computationally expensive benchmarks; and (3) building three realistic protocols (dynamic task delegation, recursive Domain Name System, and a parallel Min-Max strategy) in GoScr that could not be represented with previous theories of session types.

**2012 ACM Subject Classification** Theory of computation → Program specifications; Computing methodologies → Concurrent programming languages

**Keywords and phrases** Multiparty Session Types, Correctness-by-construction, Concurrency, Golang

**Digital Object Identifier** 10.4230/LIPIcs.CVIT.2016.23

**Acknowledgements** We deeply thank Benito Echarren Serrano for his initial collaborations on implementations of the preliminary version of this work. This work is supported by EPSRC EP/T006544/2, EP/K011715/1, EP/K034413/1, EP/L00058X/1, EP/N027833/2, EP/N028201/1, EP/T014709/2, EP/V000462/1, EP/X015955/1 and NCSS/EPSRC VeTSS.

## 1 Introduction

**Multiparty Session Types.** Multiparty Session Types (MPST) are typing disciplines that can guarantee the absence of deadlocks and communication errors in concurrent and distributed systems [22, 23]. MPST allow the specification of *global communication protocols* (**global types**) among a number of participants. The **projection** operation extracts the *local communication protocols* (**local types**), from the point of view of each participant in the system. Projection only succeeds when the protocol is absent of deadlocks and communication errors. These local types can then be used to typecheck processes [22], generate **correct by construction** code [26, 7], or monitor to detect communication errors at runtime [12].

However, MPST have a severe limitation: they cannot model protocols in which *new participants join the system*. Many important protocols rely on this. For example, Chord [54]



© David Castro-Perez and Nobuko Yoshida;  
licensed under Creative Commons License CC-BY 4.0

42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:29

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

is a popular protocol for distributed hash tables where participants join a ring, and relies on a stabilisation protocol to guarantee that each participant keeps up-to-date channels to their successors and predecessors. To model such scenarios using MPST, it would be necessary to *interleave* different sessions. But arbitrary session interleavings can lead to deadlocks, so it must be restricted [6, 9]. This not only rules out the use of MPST for many realistic scenarios, but also limits the applicability of MPST for languages that favour process creation and message passing, such as **Go**, which is the main motivation of our work.

***Dynamic (Unbounded) Participants in Go.*** Go is a concurrent programming language designed in 2009 by Google, and it is increasingly popular among professional developers. According to a 2020 Stack Overflow survey, Go is used by 9.4% of developers, and it is the “third most wanted language” [4]. Go was also the 4th most active language in GitHub in 2020 [2], and it has been adopted in many large software systems such as Kubernetes [33], gRPC [3] and Docker [1]. Its main features are explicit communication primitives, namely *channels* and *goroutines* (lightweight threads), whose design comes from concurrent process calculi [21, 41, 42]. Unfortunately, a recent empirical study reveals that over 50% of Go concurrent bugs are caused by communication [56, 40, 61] (i.e., more than shared memory bugs). While Go includes a global *runtime* deadlock detector, it is neither adequate to verify applications with complex communication structures, nor can it detect deadlocks involving only a strict subset of a program’s goroutines (partial deadlocks) [38].

Figure 1 illustrates Go’s core concurrency constructs. It shows a server (Master) that processes client requests (Line 4), and sends responses back to the Client (Line 20). The Master breaks down the request into different subtasks and delegates them to different Worker goroutines (Lines 7–10). The Master then aggregates the Worker results (Lines 11–19). If the Master receives an error message, it will forward it to the Client and stop processing any new messages (Lines 16–19). This program uses a common Go computation pattern<sup>1</sup>, the master-worker pattern, and the number of workers depends on a **runtime value**.

Unfortunately, there is a bug in the implementation in Figure 1. The implementation uses synchronous channels. Since the Master goroutine stops processing Worker responses after receiving the first error message, all other goroutines which have not sent their result or error messages will be **deadlocked**, as they will be stuck waiting for the Master to process their message. One might think that this error could be fixed by replacing the synchronous channels in the implementation with asynchronous (buffered) channels. Unfortunately, this approach leaves *orphan messages* which could introduce other concurrency bugs, e.g. the Master may need to clean up resources after receiving a response from the Workers.

This example demonstrates how even in simple programs, message passing can introduce concurrency bugs and channel leakage, violating **deadlock-freedom** and **liveness**. While, in simple programs, these concurrency bugs can be fixed with relative ease, identifying and fixing them is usually done during testing phase, which becomes increasingly harder as the complexity of the program and the number of goroutines increases. Unfortunately, standard MPST cannot model protocols such as Figure 1, since the number of participants is not fixed at the start, and depends on a run-time value.

***Adding dynamic participants to MPST.*** This paper introduces *Dynamically Updatable Multiparty Session Types* (DMst), a new theory of MPST whose **novel feature** is to model protocols in which participants can join an already existing session (*dynamic participants*). DMst can guarantee *deadlock-freedom* and *liveness* (*partial-deadlock-freedom*) **by construction** in such protocols. We implement DMst as a tool, GoScr, which generates

<sup>1</sup> E.g. [https://github.com/tmrts/go-patterns/blob/master/messaging/fan\\_out.md](https://github.com/tmrts/go-patterns/blob/master/messaging/fan_out.md)

```

1 func Worker(n int, resp chan int, err chan error) { ... } // Worker returns either result or error
2 func Master(reqCh chan int, respCh chan []int, cErrCh chan error) {
3   for {
4     ubound := <-reqCh // Receive request from Client
5     workerChs := make([]chan int, ubound) // Array to store worker result channels
6     errCh := make(chan error)
7     for i := 0; i < ubound; i++ { // n_workers depends on runtime value
8       workerChs[i] = make(chan int) // Create worker channel
9       go Worker(i+1, workerChs[i], errCh)
10    }
11    var res []int
12    for i := 0; i < ubound; i++ { // Aggregate worker results
13      select {
14        case sql := <-workerChs[i]: // Aggregate successful result
15          res = append(res, sql)
16        case err := <-errCh: // Some worker failed
17          cErrCh <- err // Propagate error and
18          return // stop processing any further messages
19      }
20    }
21    respCh <- res // Send final result to client
  }
}

```

■ **Figure 1** Dynamic task delegation implementation in Go (unsafe)

**correct by construction Go code**, and we evaluate it on a number of representative algorithms in Go, including a safe version of Figure 1 (see § 5.2(a)). While our target language is Go, DMst is not Go specific and a part of GoScr (GoScr protocols, projection and local protocols in Figure 2) is reusable for any language, as long as it supports (1) the creation of new participants (threads or processes) and (2) communication between participants.

**Contributions.** DMst overcomes several bottlenecks of existing theories on session types (A); and the two main lines of work (B,C) for static deadlock detection in Go:

**(A) Dynamic Participants and Session Types.** There are two main existing theoretical lines of work related to dynamic MPST. *Dynamic Multirole Session Types* (MRST) [13] enable a set of participants which belong to the same group (i.e. role) to join a multiparty session type. The roles are *fixed at the start*, and can only join at specific points in the protocol, e.g at the beginning of each protocol iteration. *Nested MPST* [11] model protocols with unbounded new participants. Neither MRST nor Nested MPST can represent DMst protocols (A-1) where participants join dynamically to recursive protocols, except at fixed points and with fixed roles (see Example 6). In addition, (A-2) our theory provides stronger guarantees than [11], while their global types are more complex, as they must be checked by a complex typing system. Hence a safe version of Figure 1 cannot be represented by [11, 13]. Both of [13, 11] are *only theoretical*, and lack any implementation or practical results. DMst’s global types are not only more expressive than those in [13, 11], but also simpler, thus DMst is more suitable for real language implementations. Other lines of work add session types to calculi that allow dynamic participants, or extend MPST to specify where can participants join in a protocol, e.g. [20, 57, 58, 31]. While these lines of work can add or replace participants to a system, these participants must act according to known, fixed roles. Therefore, these lines of work do not allow the specification of cyclic recursive topologies that change dynamically with the introduction of new participants.

**(B) Inference Approach.** This approach verifies safety and liveness properties of Go programs, by using model-checking on their *inferred* concurrent behavioural types [48, 37, 38, 17]. The major limitations of this approach are: (B-1) there is a gap between properties of types and programs, i.e., there are cases where types satisfy liveness but programs do not, leading to *unsound* verification, and (B-2) it cannot verify infinitely spawning goroutines because either the theory is limited to bounded approximation [37] or a decidable set of types are limited to finite-control (i.e. no parallel processes inside loops) [38, 17].

**(C) Go Code Generation.** Another approach is the generation of Go code from

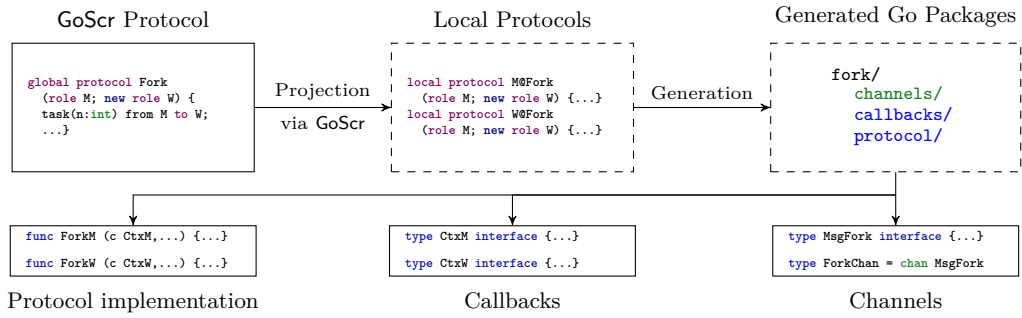


Figure 2 Overview of GoScr toolchain

parameterised multiparty session protocols [7]. However, the major limitation of [7] is that participants in a protocol still need to be fixed at the start of a session, so it cannot express and generate code for typical Go-style programs with goroutines – e.g. a safe version of Figure 1. There is a subtle, but important distinction between dynamic participants and parameterised roles: parameterised roles cannot depend on a run-time value that is exchanged in a message that is part of the protocol, because in parameterised MPST approaches, all of the participants must join the session at session initialisation, and are therefore fixed.

Our challenges are to overcome all these limitations with a scalable (implementable) MPST theory. In summary, this work solves bottlenecks of the existing MPST work by proposing a new theory, DMst, that allows the dynamic generation of an unbounded number of participants in recursive protocols, overcoming expressiveness issues in [13, 11] (A) and (C); unsoundness (B-1), but is not limited to a bounded analysis nor finite-control (B-2).

Outline. § 2 presents an overview of the GoScr toolchain; § 3 presents DMst, multiparty session types extended with the ability to add unbounded participants dynamically during a protocol execution, and proves its deadlock-freedom (Theorem 23), orphan message freedom, and liveness (Theorem 29); § 4 describes the code generation process of GoScr, and how to use it to implement DMst protocols; § 5 first measures the runtime overhead of the GoScr backend, then demonstrate the expressiveness of DMst, comparing the expressiveness of GoScr to (A) [48, 37, 38, 17] and (B) [7] with a number of case studies. We also implement three use cases – dynamic task delegation, a recursive Domain Name System, a noughts and crosses game with Min-Max strategy – to demonstrate the applicability of GoScr; § 6 gives related work, and § 7 concludes with future work.

## 2 Overview of GoScr

GoScr follows a typical Multiparty Session Types workflow (Figure 3). In this workflow, the starting point is the definition of a global protocol (global type in MPST), which describes a structured sequence of interactions between a number of participants. From this global type, we extract automatically a number of local protocols (local types) that describe the interactions (i.e. send or receive actions) from the point of




Figure 3 MPST in a nutshell

view of every participant in the protocol. This is done using the projection operation. If some participant is not projectable, then we raise an error, since the protocol is not well-formed and can lead to deadlocks or other communication errors. If the programmer provides a set of processes that behaves as prescribed by each of the local types, then the whole system

```

1  global protocol UPipe(role M;new role W){
2      rec X {
3          choice at M {
4              (Put:int) from M to W;
5              continue X with W calls UPipe(W);
6          } or {
7              (Quit:int) from M to W; }}}
8  local protocol M@UPipe(role M;new role W){
9      rec X {
10         choice at M {
11             (Put:int) to W;
12             continue X;
13         } or {
14             (Quit:int) to W; }}}

```

```

1  local protocol W@UPipe(role M;new role W){
2      choice at M {
3          (Put:int) from M;
4          invite UPipe(self; new W2);
5          rec X {
6              choice at M {
7                  (Put:int) from M;
8                  (Put:int) to W2;
9                  continue X;
10             } or {
11                 (Quit:int) from M;
12                 (Quit:int) to W2;
13             }}
14         } or { (Quit:int) from M; }}

```

■ **Figure 4** Global and Local protocols for a dynamic recursive pipeline.

is safe. We take a *code generation* approach, where we generate process code from their respective local types, providing safety guarantees *by construction*.

**Adding Participants Dynamically.** GoScr is a code generation tool which extends nuScr [49] with the theory of DMst, targeting the Go language. nuScr is a new implementation of Scribble [52], aimed at experimenting with extensions to core MPST. Figure 2 presents an overview of GoScr. We distinguish toolchain internals (dashed boxes) from tool inputs (solid boxes). Development starts by specifying a global protocol in GoScr [52, 24], a programmer-friendly protocol description language based on MPST [23, 45]. GoScr validates the well-formedness of the protocol, and produces a local type for each participant via *projection*. GoScr generates protocol implementations from these sets of local types. We provide an overview of the GoScr workflow using the dynamic recursive pipeline of Figure 4. Intuitively, this pipeline introduces a new participant after each iteration.

**Global Protocol Specification.** The first key novel feature of GoScr is the ability to *define* and *call* protocols (e.g. Line 1 in Figure 4) that may bring new participants to the protocol *dynamically*, specified by the `new` keyword in the signature. These calls can be recursive, allowing for an *unbounded* number of participants. Lines 1–14 declare `UPipe`, which requires only participant `M`, and introduces a new participant `w` dynamically. Protocol calls create any necessary new participants, as well as any necessary channels, before performing the interactions described by the called protocol. The second key novel feature of GoScr is the ability to *modify* a recursive protocol by *combining* (i.e. *interleaving*) its interactions with those of a protocol call. In our syntax, this is specified by *annotating* recursion variables with protocol calls. We call this *updatable recursion*, and an example of this can be found in Line 5. The meaning of such calls is as follows. Suppose that processes  $r_0$  and  $r_1$  are behaving as `M` and `w` (resp.) in `UPipe`. Just before the protocol jumps back to Line 2, process  $r_1$  calls protocol `UPipe(r1)`. This means that  $r_1$  will create a new participant  $r_2$ , and  $r_1$  will *delegate* to  $r_2$  a session to act as `w` in `UPipe`, with  $r_1$  acting as `M`. But at this point,  $r_1$  should act as both `M` and `w`. To address this,  $r_1$  will combine its interactions acting as `M` and `w`. The fact that  $r_1$  needs to change its behaviour to act as *two* distinct roles in `UPipe` will be reflected in its *local protocol specification* (participant `w` in Figure 4).

**Local Protocol Specification.** GoScr extracts local protocol specifications from global protocols using an operation called *projection*. Local protocols describe the structured sequence of interactions, from the point of view of a single participant. Figure 4 lists local protocols for `M` and `w`. Consider the point of view of the new participant `w` in protocol `UPipe` from Figure 4. `w` first receives an integer, either with label `Put` or `Quit` from `M`. If `W`

```

1 type Put int
2 type Quit int
3 type Ctx_UPipe_W interface {
4     Recv_M_Put(v_2 Put)
5     Init_W_UPipe() Ctx_UPipe_W
6     ...
7     Recv_M_Quit(v_2 Quit)
8     Quit() }
9 func UPipeW(ctx Ctx_UPipe_W,
10  wg *sync.WaitGroup, chMW chan MsgUPipe){
11     defer wg.Done()
12     x_1 := <- chMW
13
14     switch v_2 := x_1.(type) {
15     case Put:
16         ctx.Recv_M_Put(v_2)
17         ch_W_W_1 := make(chan MsgUPipe, 1)
18         ctx_1 := ctx.Init_W_UPipe_Ctx()
19         wg.Add(1)
20         go UPipeW(ctx_1, wg, ch_W_W_1)
21     case Quit:
22         ctx.Recv_M_Quit(v_2)
23         ctx.End() }
24 }

```

■ **Figure 5** Implementation and context of role `W` in `UPipe` in Figure 4.

receives `Quit`, then the protocol finishes. Otherwise, `w` performs a protocol call, bringing in a new participant `w2` to act as `w` in `UPipe`. In the subsequent interactions, from Line 5, `w` acts as both `M` (with respect to `w2`) and `w` (with respect to `M`). These lines (5 – 13) appear as a result of projecting Line 5 onto `w`. Notice that, if we have two participants, one acting as `M` and another one acting as `w`, this will generate a pipeline with an unbounded number of stages, until the first participant acting as `M` sends `Quit`. These kinds of protocols could not be represented in previous MPST theories and frameworks.

**Program Logic.** From local protocol specifications, `GoScr` generates the implementation of each role as a self-contained function. `GoScr` interleaves communication actions and the program logic. Communication actions in Go are a direct translation of those in local protocols: a send is a regular Go send, a receive is a regular Go receive, a choice is a type switch on a label, etc. Programmer inputs at this stage are, therefore, protocol specifications and program logic. We follow a *callback* approach similar to [43, 62] that guarantees correctness of communication *by construction*, unlike other approaches that required runtime *linearity* checks [7]. We discuss this approach in detail in §4. Figure 5 presents the code that `GoScr` generates for `w` in `UPipe`. The generated implementation requires that the programmer implements the **context** interface `Ctx_UPipe_W` (Lines 3 – 8, Figure 5). This interface defines all the necessary callbacks to implement the program logic. Programmers can use any type definition to store a local state for each participant in the protocol, e.g.

```

1 type CtxW int // This type implements Ctx_UPipe_W, and stores the accumulated sum
2 func (c *CtxW) Recv_M_Put(v upipe.Put) { // upipe.Put is also an 'int'
3     *c += CtxW(v) }
4 ...

```

By using `CtxW` for implementing `Ctx_UPipe_W`, workers will store the sum of all the numbers that they receive, and forward their accumulated sum to the next participant. The generated code for `w` will signal when it has terminated (Line 11), and starts by receiving from `M` (Line 12). Depending on whether `w` receives `Put` or `Quit`, `w` continues with the corresponding branch (Line 14). If `M` sends `Put`, then `w` creates a new participant that also acts as `w` (with respect to the previous `w`). To create this participant, first a channel is created (Line 17), then a new context is created (Line 18), the participant count is increased to guarantee that execution does not end before all participants have ended (Line 19), and finally a new goroutine is created (Line 20). Otherwise, if `M` sends `Quit`, the callback for ending is called, a last callback to perform any necessary cleanup is called, and the participant ends (Lines 25 and 26).

As we show in Figure 2, from a global protocol specification `GoScr` produces an implementation of all of its participants. To run this generated implementation, programmers must provide the necessary types to represent protocol contexts and their required callbacks. Our code generation scheme statically ensures that implementations never lead to the errors described in § 1, i.e. there will be no **deadlocks** and **orphan messages**.

### 3 Dynamically Updatable Unbounded Multiparty Session Protocols

This section introduces the theory of *Dynamically Updatable Multiparty Session Types* (DMst) with examples, and proves that DMst satisfies deadlock-freedom and liveness. DMst is the formalism that underlies GoScr. To illustrate our theory, consider the dynamic pipeline of Figure 4. In this protocol, new participants are introduced into the protocol after each iteration. In DMst, we write this dynamic pipeline as follows:

$$\text{Pipe} = \lambda\langle \mathbf{p}; \nu \mathbf{q} \rangle. \mu \mathbf{t}. (\mathbf{p} \rightarrow \mathbf{q}; \text{put}[\text{nat}]. (\mathbf{t} \blacklozenge \mathbf{q} \hookrightarrow \text{Pipe}(\mathbf{q}))) + (\mathbf{p} \rightarrow \mathbf{q}; \text{quit}. \text{end})$$

This protocol definition requires two participants  $\mathbf{p}$  and  $\mathbf{q}$ . Participant  $\mathbf{q}$  is annotated with  $\nu$  to specify that it is introduced *dynamically* (a *dynamic participant*). Participant  $\mathbf{p}$  is called a *parameter participant*. The body of the protocol specifies that it is a recursive protocol ( $\mu \mathbf{t}. \dots$ ), with recursion variable  $\mathbf{t}$ , where  $\mathbf{p}$  sends to  $\mathbf{q}$  either **put** or **quit**. This is a choice (+), where each branch starts with  $\mathbf{p} \rightarrow \mathbf{q}; \text{put}[\text{nat}]$  and  $\mathbf{p} \rightarrow \mathbf{q}; \text{quit}$  respectively. If  $\mathbf{p}$  sends **put**, then both participants enter a new iteration, but  $\mathbf{q}$  *extends* the protocol by performing call to *Pipe* ( $\mathbf{q} \hookrightarrow \text{Pipe}(\mathbf{q})$ ) before entering the new iteration. Note that although the signature mentions two participants  $\mathbf{p}$  and  $\mathbf{q}$ , the call in the global type only needs to list the parameter participants. This protocol call effectively brings in a new participant to the protocol (e.g.  $\mathbf{r}$ ), creates and distributes the necessary additional channels, and extends the interactions of the protocol with those of *Pipe*( $\mathbf{q}; \mathbf{r}$ ).

Introducing new interactions into an existing protocol requires to *interleave* them with the actions of this existing protocol. For example, the interactions of *Pipe*( $\mathbf{q}; \mathbf{r}$ ) need to be interleaved with the remaining interactions of *Pipe*( $\mathbf{p}; \mathbf{q}$ ). Our protocol specification allows two forms of interleavings: (a) sequencing all the interactions of a protocol call with the remaining interactions; and (b) alternating the actions of each iteration of two recursive protocols. We introduce a protocol construct  $\blacklozenge$  to specify the latter.

#### 3.1 Global Types of DMst

The syntax of DMst global types (given in Definition 1) is an extension of the simplest version of MPST [60]. The novel added features are **highlighted**.

► **Definition 1** (DMst Global Types).

$$\gamma ::= \mathbf{p} \rightarrow \mathbf{q}; m[U] \mid \mathbf{p} \hookrightarrow x(\vec{\mathbf{q}}) \qquad G ::= \text{end} \mid \gamma.G \mid \sum_{i \in I} G_i \mid \mu \mathbf{t}. G \mid \mathbf{t} \mid G \blacklozenge \vec{\gamma}$$

**Prefixes** ( $\gamma, \gamma', \dots$ ) represent individual interactions between *participants*, also called *roles*<sup>2</sup>, ( $\mathbf{p}, \mathbf{q}, \mathbf{r}, \dots$ ). There are two prefixes: messages and protocol calls. A *message* between  $\mathbf{p}$  and  $\mathbf{q}$  with *label*  $m$  and *payload type*  $U$  (e.g. `int`, `bool`, `...`) is written  $\mathbf{p} \rightarrow \mathbf{q}; m[U]$ , or  $\mathbf{p} \rightarrow \mathbf{q}; m$  whenever the payload is not relevant, e.g. when  $U$  is unit. We write  $\mathbf{p} \hookrightarrow x(\vec{\mathbf{q}})$  to denote a *call* to protocol  $x$  by  $\mathbf{p}$ , with participants  $\vec{\mathbf{q}}$  ( $= \mathbf{q}_1 \dots \mathbf{q}_n$ ) (see *protocol definitions* below). A protocol call prefix will introduce the new interactions described by  $x$ .

**Global types** ( $G, G', \dots$ ) denote global protocols among participants. The syntax of global types is mostly standard: **end** is *termination* and it is often omitted.  $\mathbf{t}$  denotes a *recursive variable*. **Choice**  $\sum_{i \in I} G_i$  chooses any  $G_i$ , depending on the first action of each  $G_i$  (see Definition 3). **Recursive protocol**  $\mu \mathbf{t}. G$  behaves as  $G$ , binding recursive variable  $\mathbf{t}$  to  $\mu \mathbf{t}. G$ . **Sequencing**  $\gamma.G$  denotes the execution of a prefix  $\gamma$ , and a continuation  $G$ . The new construct  $G \blacklozenge \vec{\gamma}$  denotes an *updatable protocol*, where  $G$  is extended with the interactions

<sup>2</sup> A participant plays a *role* in the protocol, and this role is determined by the structured sequence of interactions that are allowed by the global type.



and participants introduced by  $\vec{\gamma}$  (if any). When  $G$  is a recursive variable  $\mathbf{t}$  ( $\mathbf{t} \blacklozenge \vec{\gamma}$ ), we often call these *updatable recursion*, or *updatable recursion variable*. We use updatable protocols to represent recursive protocols where subsequent iterations are extended with new message exchanges and/or participants. We will show in Example 6 how to use updatable recursion to represent the dynamic recursive pipeline of Figure 4.

**Choice well-formedness.** Standard MPST syntax only allows choices where a participant  $\mathbf{p}$  sends to another participant  $\mathbf{q}$  a distinct label in each branch. This means that  $\mathbf{p}$  and  $\mathbf{q}$  can use the label to distinguish each branch of the choice [22, 60]. DMst’s syntax is more flexible, since branches can also be distinguished by distinct protocol calls. However, we still require that a single participant either sends a distinct label, or performs a distinct protocol call as the first interaction of each branch. We say that the choices that satisfy this condition are *directed*. Checking that choices are directed is necessary for well-formedness, but it is not sufficient. Protocol well-formedness is defined in a standard way later in Definition 15. To refer to the interaction that occurs in a branch, we use the *extended labels*.

► **Definition 2** (Extended Labels). We define extended labels,  $\ell ::= m \mid i @ x(\vec{\mathbf{p}}; \vec{\mathbf{q}})$ , where  $i @ x(\vec{\mathbf{p}}; \vec{\mathbf{q}})$  identifies a protocol call as the  $i$ -th participant of  $x$  with participants  $\vec{\mathbf{p}}; \vec{\mathbf{q}}$ . We use participant index instead of name, since  $x$  may give different names to  $\vec{\mathbf{p}}$  and  $\vec{\mathbf{q}}$ .

► **Definition 3** (Directed Choices). Then, we define  $\text{dc}$  (directed choice):

$$\text{dc}(\mathbf{p}, \{\ell_i\}_{i \in I}, \sum_{i \in I} \gamma_i \cdot G_i) = (\forall i \in I. \text{inter}(\mathbf{p}, \ell_i, \gamma_i)) \text{ with all } \ell_i \neq \ell_j \text{ for } i \neq j$$

The predicate  $\text{inter}(\mathbf{p}, \ell_i, \gamma_i)$  states that  $\gamma_i$  is an interaction initiated by  $\mathbf{p}$  with extended label  $\ell_i$ :  $\text{inter}(\mathbf{p}, i @ x(\vec{\mathbf{p}}; \vec{\mathbf{q}}), \mathbf{p} \hookrightarrow x(\vec{\mathbf{q}}))$ , if  $i \leq \text{size}(\vec{\mathbf{p}}\vec{\mathbf{q}})$ , and  $\text{inter}(\mathbf{p}, m, \mathbf{p} \rightarrow \mathbf{q}; m[U])$ .

**Protocol definitions** ( $x = \lambda(\vec{\mathbf{q}}; \nu \vec{\mathbf{r}}). G$ ) associate a protocol name  $x$  with a global type  $G$ , given a sequence of parameter participants  $\vec{\mathbf{q}}$ , and a sequence of *new* participants  $\vec{\mathbf{r}}$  (where “ $\nu$ ” means “new” [42]) that join the protocol *dynamically* (we call these *dynamic participants*). Any participant occurring in  $G$  must be bound by  $\vec{\mathbf{q}}$  or  $\vec{\mathbf{r}}$ . Protocol call prefixes ( $x(\vec{\mathbf{q}})$ ) only specify the parameter participants, not the dynamic ones. To refer to the global type of a definition, we write  $x(\vec{\mathbf{q}}; \vec{\mathbf{r}})$ , with parameter participants  $\vec{\mathbf{q}}$ , and dynamic participants  $\vec{\mathbf{r}}$ .

► **Example 4** (Fibonacci). The following protocol represents the interactions of an unbounded series of participants, that together compute the Fibonacci sequence:

$$\text{Fib} = \lambda(\mathbf{s}, \mathbf{f}_1, \mathbf{f}_2; \nu \mathbf{f}_3). \mathbf{f}_1 \rightarrow \mathbf{f}_3; \text{F}[\text{int}]. \mathbf{f}_2 \rightarrow \mathbf{f}_3; \text{F}[\text{int}]. \mathbf{f}_3 \rightarrow \mathbf{s}; \text{NF}[\text{int}]. \mathbf{f}_3 \hookrightarrow \text{Fib}(\mathbf{s}, \mathbf{f}_2, \mathbf{f}_3). \text{end}$$

$\text{Fib}$  defines a protocol that recursively creates new participants ( $\mathbf{f}_3$  in the global type) to compute the next element of the Fibonacci sequence after receiving the results from the previous two participants ( $\mathbf{f}_1$  and  $\mathbf{f}_2$ ). Participant  $\mathbf{s}$  receives all the results. Intuitively, the implementation of  $\mathbf{f}_3$  starts by receiving from  $\mathbf{f}_1$  and  $\mathbf{f}_2$ , sends the new Fibonacci number to  $\mathbf{s}$ , and then creates a new participant and continues with  $\mathbf{f}_2$  acting as  $\mathbf{f}_1$ , and  $\mathbf{f}_3$  as  $\mathbf{f}_2$ . The code generated by a similar protocol is shown later in Figure 6.

Protocol calls can also be used to represent recursive protocols that are augmented dynamically with new interactions and/or participants. To represent such protocols we use updatable recursion variables. Intuitively, subsequent iterations of a recursive protocol  $\mu \mathbf{t}. G$  that contains an updatable recursion variable  $\mathbf{t} \blacklozenge \mathbf{p} \hookrightarrow x(\vec{\mathbf{q}})$  will proceed as  $\mu \mathbf{t}. G$  combined with the global type defined by  $x$ . Global types are combined by interleaving their interactions.

► **Definition 5** (Combining Recursive Global Types). Let  $\text{cont}$  be a function that computes the set of final continuations, i.e. recursion variables or  $\text{end}$ , after executing all possible prefixes:  $\text{cont}(\gamma. G) = \text{cont}(G)$ ,  $\text{cont}(\mu \mathbf{t}. G) = \text{cont}(G) \setminus \{\mathbf{t}\}$ ,  $\text{cont}(\sum_{i \in I} G_i) = \cup_{i \in I} \text{cont}(G_i)$ ,  $\text{cont}(\mathbf{t}) = \{\mathbf{t}\}$ ,  $\text{cont}(G \blacklozenge \gamma) = \text{cont}(G)$ ,  $\text{cont}(\text{end}) = \{\text{end}\}$ . We define

$$(\mu t. \sum_{i \in I} G'_i) \diamond (\mu t. \sum_{i \in I} G_i) = \mu t. \sum_{i \in I} (G'_i \diamond_t G_i)$$

where  $G' \diamond_t G = [G/t]G'$  if  $\text{cont}(G') = \text{cont}(G) = \{t\}$ ,  $G' \diamond_t G = [G/\text{end}]G'$  if  $\text{cont}(G') = \text{cont}(G) = \{\text{end}\}$ , and is undefined otherwise.

The composition operator takes two recursive protocols with the same branching structure, and combines each of the branches using  $G' \diamond_t G$ . This operator simply appends the interactions of  $G$  after the interactions of  $G'$  by substituting either  $\text{end}$  or  $t$  by  $G$ . Both  $G$  and  $G'$  must finish with the same last continuation, either  $t$  or  $\text{end}$ . For example:

$$((\gamma_1. \text{end}) + (\gamma_2. t)) \diamond_t ((\gamma_3. \text{end}) + (\gamma_4. t)) = (\gamma_1. \gamma_3. \text{end}) + (\gamma_2. \gamma_4. t),$$

but the following case is undefined:  $((\gamma_1. \text{end}) + (\gamma_2. t)) \diamond_t ((\gamma_3. \text{end}) + (\gamma_4. t'))$  (if  $t \neq t'$ )

► **Example 6** (Dynamic Recursive Pipeline). Consider again the dynamic pipeline of Figure 4:

$$\text{Pipe} = \lambda \langle p; \nu q \rangle. \mu t. (p \rightarrow q: \text{put}[\text{nat}]. (t \blacklozenge q \hookrightarrow \text{Pipe}(q))) + (p \rightarrow q: \text{quit}. \text{end})$$

A set of processes that runs according to this specification would proceed as follows. The first iteration is the same as the first iteration of  $\text{Pipe}$ , but without updatable recursion. This is equivalent to the following global type:

$$G_0 = \mu t. (p \rightarrow q: \text{put}[\text{nat}]. t) + (p \rightarrow q: \text{quit}. \text{end})$$

I.e. participant  $p$  would start by sending  $\text{put}$  or  $\text{quit}$  to  $q$ , and  $q$  would receive this message. Subsequent iterations will *combine*  $G_0$ , with the result of the protocol call  $(\text{Pipe}(q))$ . Given a fresh participant  $r$ , this is as follows:

$$\begin{aligned} G_1 &= G_0 \diamond \text{Pipe}(q; r) = G_0 \diamond (\mu t. (q \rightarrow r: \text{put}[\text{nat}]. (t \blacklozenge r \hookrightarrow \text{Pipe}(r))) + (q \rightarrow r: \text{quit}. \text{end})) \\ &= \mu t. (p \rightarrow q: \text{put}[\text{nat}]. q \rightarrow r: \text{put}[\text{nat}]. (t \blacklozenge r \hookrightarrow \text{Pipe}(r))) + (p \rightarrow q: \text{quit}. q \rightarrow r: \text{quit}. \text{end}) \end{aligned}$$

Note that  $\diamond$  plugs in the interactions of the first (second) branch of  $\text{Pipe}(q; r)$  after the first (resp. second) branch of  $G_0$ . This has the effect that, after each iteration of the protocol, a new participant will join the pipeline, until the first participant sends message  $\text{quit}$ . Such protocols could not be represented in previous MPST extensions. See §6 for a discussion.

► **Example 7** (Dynamic Ring). DMst can also be used to model protocols, such a *dynamic ring*, in which participants join a recursive ring protocol. Such dynamic rings are at the core of some well-known protocols, such as Chord and its extensions. The protocol in DMst is as follows, omitting choices and payload types for simplicity:

$$\text{Ring} = \lambda \langle i; p; \nu q \rangle. \mu t. p \rightarrow q: N. t \blacklozenge (q \rightarrow i: N. i \hookrightarrow \text{Ring}(i, q))$$

The entrypoint is  $\text{Ring}(p, p; q)$ . Subsequent iterations would be combined with new protocol calls (e.g.  $\text{Ring}(p, q; r)$ ), producing the following sequences of interactions:

$$G_0 = \mu t. p \rightarrow q: N. q \rightarrow p: N. t \quad G_1 = \mu t. p \rightarrow q: N. q \rightarrow q': N. q' \rightarrow p: N. t \quad G_2 = \dots$$

### 3.2 Asynchronous Semantics of DMst Global Types

We guarantee the processes implementing all roles in a global type  $G$  indeed behave as  $G$ . To characterise the set of behaviours that are allowed by  $G$ , we define the semantics of global types as a Labelled State Transition System. The labels are the *observable actions*:

$$\alpha ::= pq! \ell \mid pq? \ell \mid pq \nu i @ x(\vec{r}; \vec{s})$$

Observable  $pq! \ell$  is a *send* action from  $p$  to  $q$  with an *extended label* (either a label or a protocol call, see Definition 2). Action  $pq? \ell$  is *receive* and action  $pq \nu i @ x(\vec{r}; \vec{s})$  is *participant creation* which brings in  $q$  as a new participant acting as the  $i$ -th role in the protocol specified

by  $x(\vec{r}; \vec{s})$ . For simplicity, we sometimes write  $\mathbf{pq} \nu \ell$  to refer to participant creation, assuming that  $\ell$  is of the form  $i @ x(\vec{r}; \vec{s})$ , for some  $i, x, \vec{r}$  and  $\vec{s}$ .

**Extended Global Types.** We extend the global types (Definition 1) with constructs that capture intermediate states of the execution of a protocol. Note that these intermediate states only appear as a result of applying the rules of the operational semantics, and these will not need to be written by users specifying full protocols. Since extended global types are a superset of Definition 1, we will use the same meta-variable  $G$  for both and, unless we specify otherwise, all global types from now on are considered to be extended.

$$\gamma ::= \dots \mid \mathbf{p} \rightarrow \mathbf{q} : \ell[U] \mid \mathbf{p} \rightsquigarrow \mathbf{q} : \ell[U] \mid \mathbf{p} \nu(\vec{r} : [i, j] @ x(\vec{p}; \vec{q})) \mid \triangleright[G]$$

Sending protocol call labels (e.g.  $\mathbf{p} \rightarrow \mathbf{q} : i @ x(\vec{p}; \vec{q})$ ) is a form of delegation that is used to perform protocol calls (see **Notations** below).  $\mathbf{p} \rightsquigarrow \mathbf{q} : \ell[U]$  means  $\mathbf{p}$  has sent a message to  $\mathbf{q}$ , yet  $\mathbf{q}$  has not received it.  $\mathbf{p} \nu(\vec{r} : [i, j] @ x(\vec{p}; \vec{q}))$  represents that  $\mathbf{p}$  creates new participants  $\vec{r}$ , acting as the  $i$ th to  $j$ th participants in  $x$ , and  $\triangleright[G]$  is the nested protocol with global type  $G$ . Intuitively, a nested protocol prefix  $\triangleright[G_0]. G_1$  is equivalent to sequencing  $G_0$  and  $G_1$ .

**Notations.** We use notations to break down protocol calls into the individual interactions. Suppose that  $\vec{q} = (\mathbf{q}_1, \dots, \mathbf{q}_n)$  and  $\vec{r} = (\mathbf{r}_1, \dots, \mathbf{r}_m)$ . We define  $\text{id}x(\mathbf{p}; \vec{q})$  as  $\{i\}$ , if  $\mathbf{p} = \mathbf{q}_i$  with  $1 \leq i \leq n$ , or the empty set  $\{\}$  otherwise. We define the following shortcuts:

$$\begin{aligned} \mathbf{p} \rightarrow \vec{q} : i @ x(\vec{q}; \vec{r}) &= \mathbf{p} \rightarrow \mathbf{q}_1 : i_1 @ x(\vec{q}; \vec{r}). \dots \mathbf{p} \rightarrow \mathbf{q}_n : i_n @ x(\vec{q}; \vec{r}) \\ \mathbf{p} \text{ call } x(\vec{q}; \vec{r}) &= \mathbf{p} \nu(\vec{r} : [n+1, n+m] @ x(\vec{q}; \vec{r})). \mathbf{p} \rightarrow \vec{q} \setminus \{\mathbf{p}\} : ([1, n] \setminus \text{id}x(\mathbf{p}; \vec{q})) @ x(\vec{q}; \vec{r}) \end{aligned}$$

Notation  $\mathbf{p} \rightarrow \vec{q} : i @ x(\vec{q}; \vec{r})$  represents a sequence of messages from  $\mathbf{p}$  to each of the  $\mathbf{q} \in \vec{q}$  with the respective extended label. These are sometimes called **invitations** to  $x$ . Notation  $\mathbf{p} \text{ call } x(\vec{q}; \vec{r})$  is a sequence of actions, where  $\mathbf{p}$  first creates  $\vec{r}$ , and then sends invitations to  $\vec{q}$ , excluding itself to avoid self-communication.

**Global Type Equivalence and LTS.** We define the erasure of updatable recursive variables as  $|t \blacklozenge \gamma|_{t'} = t$  if  $t = t'$ ; and  $|t \blacklozenge \gamma|_{t'} = t \blacklozenge \gamma$  otherwise (other cases are homomorphic). The LTS is defined up to the equivalence: (1)  $\mathbf{p} \nu(); G \equiv G$ ; (2)  $\triangleright[\text{end}]. G \equiv G$ ; (3)  $\mu t. G \equiv [\mu t. |G|_t / t] G$ , and, assuming  $\vec{r}$  fresh, (4)  $\mathbf{p} \hookrightarrow x(\vec{q}). G \equiv \mathbf{p} \text{ call } x(\vec{q}; \vec{r}). \triangleright[x(\vec{q}; \vec{r})]. G$ , and (5)  $G \blacklozenge (\vec{r}. \mathbf{p} \hookrightarrow x(\vec{q})) \equiv \vec{r}. \mathbf{p} \text{ call } x(\vec{q}; \vec{r}). (G \blacklozenge x(\vec{q}; \vec{r}))$ . Rules (1) and (2) capture that finished prefixes (creating an empty list of participants, or a nested ended global type) can be skipped. Rule (3) is recursion unrolling. Similarly to Example 6, subsequent iterations of the protocol will combine the body of the recursion without updatable recursion variables, with the result of the protocol calls. By this rule, recursion will be updated by protocol calls, and after the first iteration, the protocol can continue as  $\mu t. |G|_t$  (possibly combined with the result of a protocol call). Rule (4) expands the sequence of a protocol call and a global type, and rule (5) *updates* a global type by first executing the specified prefixes and then continuing with  $G$  combined with the result of the protocol call. We guarantee that new roles are globally fresh by adopting a Barendregt convention on all binders, i.e. each time we access a protocol definition  $x$ , we alpha-rename the participants bound by  $\nu \vec{r}$  to avoid participant name clashes. Without it, consecutive protocol calls could incorrectly introduce repeated participant names.

► **Definition 8** (Active Participants). *The active participants of a global type (prefix),  $\text{pt}(G)$  ( $\text{pt}(\gamma)$ ), is the set of participants that can perform an action in the protocol (or prefix).*

$$\begin{aligned} \text{pt}(\mathbf{p} \rightarrow \mathbf{q} : \ell[U]) &= \{\mathbf{p}, \mathbf{q}\} & \text{pt}(\mathbf{p} \hookrightarrow x(\vec{q})) &= \{\mathbf{p}\} \cup \vec{q} & \text{pt}(\mathbf{p} \rightsquigarrow \mathbf{q} : \ell[U]) &= \{\mathbf{q}\} & \text{pt}(\gamma) &= \text{pt}(\gamma) \cup \text{pt}(G) \\ \text{pt}(\mathbf{p} \hookrightarrow x(\vec{q})) &= \{\mathbf{p}\} \cup \vec{q} & \text{pt}(\mathbf{p} \nu(\vec{r} : [i, j] @ x(\vec{p}; \vec{q}))) &= \{\mathbf{p}\} \cup \vec{r} & \text{pt}(\triangleright[G]) &= \text{pt}(G) \\ \text{pt}(\text{end}) &= \text{pt}(t) = \{\} & \text{pt}(\mu t. G) &= \text{pt}(G) & \text{pt}(\sum_{i \in I} G_i) &= \bigcup_{i \in I} \text{pt}(G_i) & \text{pt}(G \blacklozenge \gamma) &= \text{pt}(G) \cup \text{pt}(\gamma) \end{aligned}$$

► **Definition 9** (LTS for Global Types). Let the *subject* of an action denote the role that performs it:  $\mathbf{p} = \text{subj}(\mathbf{pq}! \ell) = \text{subj}(\mathbf{pq}? \ell) = \text{subj}(\mathbf{pq} \nu \ell)$ . The LTS for  $G$ :

$$\begin{array}{c}
\frac{[\text{BR-A}]}{\forall i \in I, G_i \xrightarrow{\alpha} G'_i} \quad \frac{[\text{BR-B}]}{G_j \xrightarrow{\text{pq}! \ell_j} G'_j} \quad \text{dc}(\mathbf{p}, \{\ell_i\}_i, \sum_{i \in I} G_i) \quad \frac{[\text{NEST}]}{G_1 \xrightarrow{\alpha} G_2}}{\sum_{i \in I} G_i \xrightarrow{\alpha} \sum_{i \in I} G'_i} \quad \frac{\sum_{i \in I} G_i \xrightarrow{\text{pq}! \ell_j} G'_j}{\triangleright[G_1]. G \xrightarrow{\alpha} \triangleright[G_2]. G}}{[\text{NEW}] \mathbf{p}(\mathbf{r}, \bar{\mathbf{r}} : [i, j] \otimes x(\bar{\mathbf{q}}; \bar{\mathbf{r}})). G \xrightarrow{\text{pr} \nu i \otimes x(\bar{\mathbf{q}}; \bar{\mathbf{r}})} \mathbf{p}(\bar{\mathbf{r}} : [i + 1, j] \otimes x(\bar{\mathbf{q}}; \bar{\mathbf{r}})). G} \\
\frac{[\text{SEND}]}{\mathbf{p} \rightarrow \mathbf{q} : \ell[U]. G \xrightarrow{\text{pq}! \ell} \mathbf{p} \rightsquigarrow \mathbf{q} : \ell[U]. G} \quad \frac{[\text{RECV}]}{\mathbf{p} \rightsquigarrow \mathbf{q} : \ell[U]. G \xrightarrow{\text{qp} ? \ell} G} \quad \frac{[\text{SEQ}]}{G \xrightarrow{\alpha} G' \quad \text{subj}(\alpha) \notin \text{pt}(\gamma)}}{\gamma.G \xrightarrow{\alpha} \gamma.G'}
\end{array}$$

[BR-A] specifies that if an action can be taken in all branches of a choice, it can be taken before the choice is decided. The reason is that if an action can be taken in all branches, then it must be *independent* of the choice. [BR-B] states that if the sender of a choice does an action that selects branch  $j$ , then the choice transitions to this branch. [SEQ] states that an action can take place in a continuation, if the action does not involve the participants of the prefix. In the prefix transitions, [SEND] (resp. [RECV]) represents a send (resp. receive) action. [NEW] specifies that a new participant  $\mathbf{r}$  of the nested protocol is created, and [NEST] represents the execution of an action in the nested global type.

► **Example 10** (DMst Semantics). Consider the following protocol, cf. Example 6:

$$\text{Pipe} = \lambda(\mathbf{p}; \nu \mathbf{q}). \mu \mathbf{t}. G_0 \quad \text{with } G_0 = (\mathbf{p} \rightarrow \mathbf{q} : \text{put}[\text{nat}]. (\mathbf{t} \blacklozenge \mathbf{q} \hookrightarrow \text{Pipe}(\mathbf{q}))) + (\mathbf{p} \rightarrow \mathbf{q} : \text{quit}. \text{end})$$

First, assuming two initial participants ( $\mathbf{p}$  and  $\mathbf{q}$ ), we unfold recursion using  $\equiv$ :

$$\mu \mathbf{t}. G_0 \equiv [\mu \mathbf{t}. |G_0|_{\mathbf{t}} / \mathbf{t}] G_0 = (\mathbf{p} \rightarrow \mathbf{q} : \text{put}[\text{nat}]. (\mu \mathbf{t}. |G_0|_{\mathbf{t}} \blacklozenge \mathbf{q} \hookrightarrow \text{Pipe}(\mathbf{q}))) + (\mathbf{p} \rightarrow \mathbf{q} : \text{quit}. \text{end})$$

There are two allowed actions: sending **put** and sending **quit**. By [BR-B] and [SEND],

$$[\mu \mathbf{t}. |G_0|_{\mathbf{t}} / \mathbf{t}] G_0 \xrightarrow{\text{pq}! \text{put}} \mathbf{p} \rightsquigarrow \mathbf{q} : \text{put}[\text{nat}]. (\mu \mathbf{t}. |G_0|_{\mathbf{t}} \blacklozenge \mathbf{q} \hookrightarrow \text{Pipe}(\mathbf{q}))$$

There are now two actions accepted. First, we can use [RECV]:

$$\mathbf{p} \rightsquigarrow \mathbf{q} : \text{put}[\text{nat}]. (\mu \mathbf{t}. |G_0|_{\mathbf{t}} \blacklozenge \mathbf{q} \hookrightarrow \text{Pipe}(\mathbf{q})) \xrightarrow{\text{qp} ? \text{put}} \mu \mathbf{t}. |G_0|_{\mathbf{t}} \blacklozenge \mathbf{q} \hookrightarrow \text{Pipe}(\mathbf{q})$$

To enable the second action, we use equivalences to unfold the updatable global type:

$$\begin{aligned}
G_1 &= \mathbf{p} \rightsquigarrow \mathbf{q} : \text{put}[\text{nat}]. (\mu \mathbf{t}. |G_0|_{\mathbf{t}} \blacklozenge \mathbf{q} \hookrightarrow \text{Pipe}(\mathbf{q})) \\
&\equiv \mathbf{p} \rightsquigarrow \mathbf{q} : \text{put}[\text{nat}]. \mathbf{q} \nu(\mathbf{r} : 2 \otimes \text{Pipe}(\mathbf{q}; \mathbf{r})). (\mu \mathbf{t}. |G_0|_{\mathbf{t}} \blacklozenge \text{Pipe}(\mathbf{q}; \mathbf{r}))
\end{aligned}$$

Note that  $\mathbf{p}$  is not in the set of active participants of the prefix, so  $\mathbf{p}$  can take a step, using repeated applications of [SEQ], in  $(\mu \mathbf{t}. |G_0|_{\mathbf{t}} \blacklozenge \text{Pipe}(\mathbf{q}; \mathbf{r}))$ .

$$\begin{aligned}
\mu \mathbf{t}. G_2 &= (\mu \mathbf{t}. |G_0|_{\mathbf{t}} \blacklozenge \text{Pipe}(\mathbf{q}; \mathbf{r})) \\
&= \mu \mathbf{t}. (\mathbf{p} \rightarrow \mathbf{q} : \text{put}[\text{nat}]. \mathbf{q} \rightarrow \mathbf{r} : \text{put}[\text{nat}]. (\mathbf{t} \blacklozenge \mathbf{r} \hookrightarrow \text{Pipe}(\mathbf{r}))) + (\mathbf{p} \rightarrow \mathbf{q} : \text{quit}. \mathbf{q} \rightarrow \mathbf{r} : \text{quit}. \text{end}) \\
&\equiv (\mathbf{p} \rightarrow \mathbf{q} : \text{put}[\text{nat}]. \mathbf{q} \rightarrow \mathbf{r} : \text{put}[\text{nat}]. (\mu \mathbf{t}. |G_2|_{\mathbf{t}} \blacklozenge \mathbf{r} \hookrightarrow \text{Pipe}(\mathbf{r}))) + (\mathbf{p} \rightarrow \mathbf{q} : \text{quit}. \mathbf{q} \rightarrow \mathbf{r} : \text{quit}. \text{end})
\end{aligned}$$

Suppose that  $G_2$  proceeds by  $\mathbf{p}$  sending **quit**:  $\mu \mathbf{t}. G_2 \xrightarrow{\text{pq}! \text{quit}} (\mathbf{p} \rightsquigarrow \mathbf{q} : \text{quit}. \mathbf{q} \rightarrow \mathbf{r} : \text{quit}. \text{end})$ .

Then,  $G_1$  transitions to the following global type:

$$\mu \mathbf{t}. G_1 \xrightarrow{\text{pq}! \text{quit}} \mathbf{p} \rightsquigarrow \mathbf{q} : \text{put}[\text{nat}]. \mathbf{q} \nu(\mathbf{r} : 2 \otimes \text{Pipe}(\mathbf{q}; \mathbf{r})). \mathbf{p} \rightsquigarrow \mathbf{q} : \text{quit}. \mathbf{q} \rightarrow \mathbf{r} : \text{quit}. \text{end}$$

After [SQ-A] and [RECV], the global type transitions to:

$$G_3 = \mathbf{q} \nu(\mathbf{r} : 2 \otimes \text{Pipe}(\mathbf{q}; \mathbf{r})). \mathbf{p} \rightsquigarrow \mathbf{q} : \text{quit}. \mathbf{q} \rightarrow \mathbf{r} : \text{quit}. \text{end}$$

With [SQ-A] and [NEW], the protocol transitions as follows:

$$G_3 \xrightarrow{\text{qr} \nu 2 \otimes \text{Pipe}(\mathbf{q}; \mathbf{r})} \mathbf{p} \rightsquigarrow \mathbf{q} : \text{quit}. \mathbf{q} \rightarrow \mathbf{r} : \text{quit}. \text{end}$$

At this stage,  $\mathbf{r}$  is a new active participant of the protocol. The remaining global type can run to completion via a sequence of [RECV], [SEND], and finally [RECV].

### 3.3 Local Types

Local types describe the interactions of a protocol from the point of view of a single participant.

## 23:12 Dynamically Updatable Multiparty Session Protocols

► **Definition 11** (DMst Local Types). Let  $M ::= l[U] \mid L$ . The syntax of local types is:  
 $\pi ::= \mathbf{p}!M \mid \mathbf{p}?M \mid \nu(\vec{\mathbf{p}} : \vec{L}) \mid \triangleright[L] \quad L ::= \text{end} \mid \pi.L \mid \sum_{i \in I} L_i \mid \mu\mathbf{t}.L \mid \mathbf{t} \mid L \blacklozenge \vec{\pi}$

Local type syntax differs from that of global types in the prefixes ( $\pi$  instead of  $\gamma$ ). Local type prefixes are as follows: *send*  $\mathbf{p}!M$ , *receive*  $\mathbf{p}?M$ , *new participant creation*  $\nu(\mathbf{p}_1 : L_1) \cdots (\mathbf{p}_n : L_n)$ , and the *nested local type*  $\triangleright[L]$ . We lift the definitions of directed choices, updatable recursion erasure, and the composition operator from global types to local types. *Endpoint projection* takes a global type  $G$  and a participant  $\mathbf{r}$ , and produces the local type (the local interactions) of  $\mathbf{r}$  in  $G$ .

Similarly to global types, we introduce the notations for protocol calls. Assuming  $\vec{\mathbf{q}} = \mathbf{q}_1, \dots, \mathbf{q}_n$  and  $\vec{\mathbf{r}} = \mathbf{r}_1, \dots, \mathbf{r}_m$ , we define these notations as follows:

$$\begin{aligned} \vec{\mathbf{q}}!i @ x(\vec{\mathbf{q}}; \vec{\mathbf{r}}) &= \mathbf{q}_1!i_1 @ x(\vec{\mathbf{q}}; \vec{\mathbf{r}}). \dots \mathbf{q}_n!i_n @ x(\vec{\mathbf{q}}; \vec{\mathbf{r}}) \\ \mathbf{p} \text{ call } x(\vec{\mathbf{q}}; \vec{\mathbf{r}}) &= \nu(\vec{\mathbf{r}} : [n+1, n+m] @ x(\vec{\mathbf{q}}; \vec{\mathbf{r}})). \vec{\mathbf{q}} \setminus \mathbf{p}!([1, n] \setminus \text{id}x(\mathbf{p}; \vec{\mathbf{q}})) @ x(\vec{\mathbf{q}}; \vec{\mathbf{r}}) \end{aligned}$$

► **Definition 12** (Prefix Projection). Global type projection is defined in terms of prefix projection. Prefix projection is a partial function that takes a global prefix, and produces a possibly empty ( $\varepsilon$ ) sequence of local prefixes. We give the two main rules:

$$\begin{aligned} \mathbf{p} \rightarrow \mathbf{q}:l[U] \upharpoonright \mathbf{r} &= \begin{cases} \mathbf{q}!l[U] & \mathbf{p} = \mathbf{r} \neq \mathbf{q} \\ \mathbf{p}?l[U] & \mathbf{p} \neq \mathbf{r} = \mathbf{q} \\ \varepsilon & \mathbf{p}, \mathbf{r}, \mathbf{q} \text{ distinct} \end{cases} & \mathbf{p} \hookrightarrow x(\vec{\mathbf{q}}) \upharpoonright \mathbf{r} \quad (\vec{\mathbf{r}} \text{ fresh}) &= \begin{cases} \mathbf{p} \text{ call } x(\vec{\mathbf{q}}; \vec{\mathbf{r}}). \triangleright[i @ x(\vec{\mathbf{q}}; \vec{\mathbf{r}})] & \mathbf{p} = \mathbf{r} \in_i \vec{\mathbf{q}} \\ \mathbf{p} \text{ call } x(\vec{\mathbf{q}}; \vec{\mathbf{r}}) & \mathbf{p} = \mathbf{r} \notin \vec{\mathbf{q}} \\ \mathbf{p}?i @ x(\vec{\mathbf{q}}; \vec{\mathbf{r}}). \triangleright[i @ x(\vec{\mathbf{q}}; \vec{\mathbf{r}})] & \mathbf{p} \neq \mathbf{r} \in \vec{\mathbf{q}} \\ \varepsilon & \mathbf{p} \neq \mathbf{r} \notin \vec{\mathbf{q}} \end{cases} \end{aligned}$$

The projection of  $\mathbf{p} \rightarrow \mathbf{q}:l[U]$  onto  $\mathbf{r}$  is a *send* if  $\mathbf{r}$  is  $\mathbf{p}$ , and a *receive* if  $\mathbf{r}$  is  $\mathbf{q}$ , an empty prefix if all roles are distinct, or undefined if  $\mathbf{r} = \mathbf{p} = \mathbf{q}$ . The projection of  $\mathbf{p} \hookrightarrow x(\vec{\mathbf{q}})$  follows a similar pattern. If  $\mathbf{r}$  is  $\mathbf{p}$ , then the projected sequence of prefixes is the one that corresponds to making the protocol call, i.e. delegating channels and creating new participants. If  $\mathbf{r}$  is the  $i$ th participant in  $\vec{\mathbf{q}}$ , then  $\mathbf{r}$  also takes part in the protocol, so the prefixes correspond to the reception of the channel for acting as the  $i$ th participant in  $x$ , followed by the execution of the nested local type for this  $i$ th participant in  $x$ ,  $i @ x(\vec{\mathbf{q}}; \vec{\mathbf{r}})$ . If  $\mathbf{r}$  is both the protocol caller  $\mathbf{p}$ , and also takes part in it, then the prefix sequence is the sequence of prefixes for making the protocol call, followed by the nested local type for  $i @ x$ . Note that a participant may call a protocol, and not take part in it. When this happens, the protocol caller simply distributes the necessary channels for executing the nested protocol, and then proceeds to the continuation, without entering the nested protocol.

► **Definition 13** (Projection and Merging). Projection is defined as follows:

$$\begin{aligned} \gamma.G \upharpoonright \mathbf{r} &= \gamma \upharpoonright \mathbf{r}. G \upharpoonright \mathbf{r} & \mathbf{t} \upharpoonright \mathbf{r} &= \mathbf{t} & G \blacklozenge \vec{\gamma} \upharpoonright \mathbf{r} &= \begin{cases} G \upharpoonright \mathbf{r} & (\text{if } \mathbf{r} \notin \vec{\gamma}) \\ G \upharpoonright \mathbf{r} \blacklozenge (\vec{\gamma} \upharpoonright \mathbf{r}) & (\text{if } \mathbf{r} \in \vec{\gamma}) \end{cases} \\ \mu\mathbf{t}.G \upharpoonright \mathbf{r} &= \begin{cases} \mu\mathbf{t}.G \upharpoonright \mathbf{r} & \mathbf{r} \in \text{pt}(G) \text{ or } \\ & \text{fv}(\mu\mathbf{t}.G) \neq \emptyset \\ \text{end} & \text{otherwise} \end{cases} & \sum_{i \in I} G_i \upharpoonright \mathbf{r} &= \begin{cases} \sum_{i \in I} (G_i \upharpoonright \mathbf{r}) & \text{dc}(\mathbf{p}, \vec{\ell}, \sum_{i \in I} G_i), \mathbf{r} = \mathbf{p} \\ \prod_{i \in I} (G_i \upharpoonright \mathbf{r}) & \text{dc}(\mathbf{p}, \vec{\ell}, \sum_{i \in I} G_i), \mathbf{r} \neq \mathbf{p} \end{cases} \end{aligned}$$

Projection is a partial function from global to local types. We lift the definition of directed choices (Definition 3, dc) to local types. We define  $\prod_{i \in I} L_i$  as the *merging* operator:

$$\begin{aligned} (1) \quad L \sqcap L &= L & (2) \quad \mu\mathbf{t}.L_1 \sqcap \mu\mathbf{t}.L_2 &= \mu\mathbf{t}.(L_1 \sqcap L_2) \\ (3) \quad \sum_{i \in I} \mathbf{p}?\ell_i[U_i]. L_i \sqcap \sum_{j \in J} \mathbf{p}?\ell_j[U_j]. L'_j &= \\ & \sum_{k \in I \cap J} (\mathbf{p}?\ell_k[U_k]. L_k \sqcap L'_k) + \sum_{i \in I \setminus J} (\mathbf{p}?\ell_i[U_i]. L_i) + \sum_{j \in J \setminus I} (\mathbf{p}?\ell_j[U_j]. L'_j) \end{aligned}$$

The projection rules are standard [60], except the choice. A choice is only defined if it is directed. The projection of the participant that makes the choice is a local type choice of the

projection of the branches. The projection for all other participants is the *merging* of the projection of the branches. Local types can be merged in three cases: (1) they are the same, (2) they are recursive local types whose bodies can be merged, or (3) they become aware of which branch of the choice was taken (if necessary), by receive actions with distinct labels. Case (3) implies that both local types are choices with a receive prefix as the first action, where the continuations for the branches with the same labels can be merged.

It is standard in MPST to define *well-formedness* in terms of *projectability* [22]. This means that if a global type is projectable onto all of its roles, then it is well-formed and therefore live and deadlock-free. Unfortunately, the use of  $\blacklozenge$  means that this is not possible with DMst. E.g., the following global type is projectable, but it will get stuck:

$$\begin{aligned} \text{Proto1} &= \lambda\langle \mathbf{p}; \nu \mathbf{r} \rangle. \mu \mathbf{t}. (\mathbf{r} \rightarrow \mathbf{p}:m_1. \text{end}) + (\mathbf{r} \rightarrow \mathbf{p}:m_2. \mathbf{t}) \\ \text{IllFormed} &= \lambda\langle \mathbf{p}; \nu \mathbf{q} \rangle. \mu \mathbf{t}. (\mathbf{p} \rightarrow \mathbf{q}:m_1. \text{end}) + (\mathbf{p} \rightarrow \mathbf{q}:m_2. \mathbf{t} \blacklozenge (\mathbf{p} \leftrightarrow \text{Proto1}\langle \mathbf{p} \rangle)) \end{aligned}$$

Specifically,  $\mathbf{r}$  in *Proto1* will not become aware of the branch taken by  $\mathbf{p}$  in *IllFormed*, so after unfolding *IllFormed* once, we will obtain the following global type:

$$\mu \mathbf{t}. (\mathbf{p} \rightarrow \mathbf{q}:m_1. \mathbf{r} \rightarrow \mathbf{p}:m_1. \text{end}) + (\mathbf{p} \rightarrow \mathbf{q}:m_2. \mathbf{r} \rightarrow \mathbf{p}:m_2. \mathbf{t})$$

But this protocol would not be projectable. To avoid such cases, we define a necessary condition for well-formedness, the *safe protocol update* condition.

► **Definition 14** (Safe Protocol Update). *Suppose that  $C[\ ]$  and  $C'[\ ]$  are 1-hole global type contexts. A global type  $\mu \mathbf{t}. C[\mathbf{t} \blacklozenge (\vec{\gamma}. \mathbf{p} \leftrightarrow x(\vec{\mathbf{q}}))]$  contains a safe update if its 1-unfolding is some  $C'[G \blacklozenge (\vec{\gamma}. \mathbf{p} \leftrightarrow x(\vec{\mathbf{q}}))]$ , such that given a sequence of fresh roles  $\vec{\mathbf{r}}$ ,  $G \blacklozenge x(\vec{\mathbf{q}}; \vec{\mathbf{r}})$  is projectable.*

► **Definition 15** (Projection and Well-Formed Global Types). *A global type  $G$  is projectable if its projection  $G \upharpoonright \mathbf{r}$  is defined on all roles  $\mathbf{r} \in G$ . A global type is well formed iff it is projectable, and contains only safe protocol updates.*

► **Definition 16** (Projections of Protocol Definitions). *Assume a definition  $x = \lambda\langle \vec{\mathbf{p}}; \nu \vec{\mathbf{p}}' \rangle. G$ , with participants  $\vec{\mathbf{p}} = (\mathbf{p}_1, \dots, \mathbf{p}_n)$  and with participants  $\vec{\mathbf{p}}' = (\mathbf{p}_{n+1}, \dots, \mathbf{p}_m)$ . The projections of  $x$  are the local protocol definitions that correspond to each of the participants in the protocol:*

$$1\textcircled{x} = \lambda\langle \vec{\mathbf{p}}; \nu \vec{\mathbf{p}}' \rangle. G \upharpoonright \mathbf{p}_1 \quad \dots \quad m\textcircled{x} = \lambda\langle \vec{\mathbf{p}}; \nu \vec{\mathbf{p}}' \rangle. G \upharpoonright \mathbf{p}_m$$

► **Example 17** (Directed Choices and Merging).  *$BFib$  computes the  $n$ -th Fibonacci number:*

$$\begin{aligned} BFib &= \lambda\langle \mathbf{r}, \mathbf{f}_1, \mathbf{f}_2; \nu \mathbf{f}_3 \rangle. \mathbf{f}_1 \rightarrow \mathbf{f}_3: F[\text{int}]. \\ &\quad \mathbf{f}_2 \rightarrow \mathbf{f}_3: F[\text{int}]. ((\mathbf{f}_3 \rightarrow \mathbf{r}: NF[\text{int}]. \mathbf{f}_3 \rightarrow \mathbf{f}_2: \text{quit}. \text{end}) + (\mathbf{f}_3 \leftrightarrow BFib(\mathbf{r}, \mathbf{f}_2, \mathbf{f}_3). \text{end})) \end{aligned}$$

This protocol is similar to that of Example 4, but instead of calling *BFib* indefinitely, the protocol offers a choice:  $\mathbf{f}_3$  will either reply to  $\mathbf{r}$  with its Fibonacci number, or call *BFib* recursively to compute the next number. Participant  $\mathbf{f}_3$  *selects* the branch of the protocol that is taken, and  $\mathbf{r}$  *offers* the two branches. The choice has a single sender, and both branches can be distinguished by the labels or protocol calls, so the choice is *directed* by  $\mathbf{f}_3$ , with extended labels  $\vec{\ell} = NF, i\textcircled{BFib}(\mathbf{r}, \mathbf{f}_2, \mathbf{f}_3; \mathbf{f}_4)$ . In a directed choice, one participant decides the branch. But how do the remaining participants *know* which branch was taken? Consider  $\mathbf{f}_1$  in *BFib*. Its part in both branches of the protocol is the same, *end*, so we can project  $\mathbf{f}_1$  in the choice as *end*. This is one of the cases of Definition 13: two local types can be merged if they are the same. But  $\mathbf{f}_2$ 's behaviour is different in each branch:  $\mathbf{f}_3: \text{quit}. \text{end}$  and  $\mathbf{f}_3: 2\textcircled{BFib}(\mathbf{r}, \mathbf{f}_2, \mathbf{f}_3; \mathbf{f}_4). \text{end}$  respectively. However,  $\mathbf{f}_2$  is *aware* of the branch that was taken by receiving either label *quit* or protocol call label  $2\textcircled{BFib}(\mathbf{r}, \mathbf{f}_2, \mathbf{f}_3; \mathbf{f}_4)$ . This is case (3), as

explained after Definition 13:

$$(\mathbf{f}_3?quit. end) \sqcap (\mathbf{f}_3?2@BFib(\mathbf{r}, \mathbf{f}_2, \mathbf{f}_3; \mathbf{f}_4). end) = (\mathbf{f}_3?quit. end) + (\mathbf{f}_3?2@BFib(\mathbf{r}, \mathbf{f}_2, \mathbf{f}_3; \mathbf{f}_4). end)$$

► **Example 18** (Projecting Pipeline). Consider again Example 6. We are projecting the first and second participants of  $x$ . The result of the syntactic projection is as follows:

$$\begin{aligned} x(\mathbf{p}; \mathbf{q}) &= \mu\mathbf{t}.(\mathbf{p} \rightarrow \mathbf{q}; \text{put}[\text{nat}]. (\mathbf{t} \blacklozenge \mathbf{q} \hookrightarrow x(\mathbf{q}))) + (\mathbf{p} \rightarrow \mathbf{q}; \text{quit}. end) \\ 1@x(\mathbf{p}; \mathbf{q}) &= \mu\mathbf{t}.(\mathbf{q}! \text{put}[\text{nat}]. \mathbf{t}) + (\mathbf{q}! \text{put}[\text{nat}]. end) \\ 2@x(\mathbf{p}; \mathbf{q}) &= \mu\mathbf{t}.(\mathbf{p}? \text{put}[\text{nat}]. (\mathbf{t} \blacklozenge (\text{call } x(\mathbf{q}; \mathbf{r}). \mathbf{q}? \mathbf{q}1@x(\mathbf{q}; \mathbf{r}). \triangleright[1@x(\mathbf{q}; \mathbf{r}]))) + (\mathbf{p}? \text{quit}. end) \end{aligned}$$

### 3.4 Semantics of DMst Local Types and Correctness

The semantics for local types is defined for *local type configurations*. A configuration is a pair of *channel* and *participant* environments,  $\langle \Delta ; \Theta \rangle$ . The channel environment  $\Delta$  contains the shared channels used for the asynchronous communication between each pair of participants, and the participant environment  $\Theta$  is a set of the local types of all participants:

$$\Delta = \mathbf{p}_i \mathbf{q}_j :: \vec{\mathbf{w}}_1, \dots, \mathbf{p}_k \mathbf{q}_l :: \vec{\mathbf{w}}_n \quad \mathbf{w} ::= \ell[U] \quad \Theta = \{\mathbf{p}_1 :: L_1, \dots, \mathbf{q}_m :: L_m\}$$

$\mathbf{w}$  denotes a *payload* of a message. We consider the channel and participant environments up to commutativity and associativity, since all entries must be disjoint. Channels  $\mathbf{p}\mathbf{q}$  are channels of messages to  $\mathbf{p}$  from  $\mathbf{q}$ . We use  $\Delta(\mathbf{p}\mathbf{q})$  as notation for retrieving channel  $\mathbf{p}\mathbf{q}$ , and  $\Delta[\mathbf{p}\mathbf{q} :: \vec{\mathbf{w}}]$  for updating channel  $\mathbf{p}\mathbf{q}$  with  $\vec{\mathbf{w}}$ .  $\Theta$  does not impose the ordering between the entries (like a set). We update the entry by writing  $\Theta[\mathbf{p} :: L] = \mathbf{p} :: L, (\Theta \setminus \mathbf{p})$ .

The semantics of configurations is defined by the LTS of local types and given in Definition 19, and it is defined up to local type equivalences, analogous to those of global types: (1)  $\mu\mathbf{t}.L \equiv \vec{\pi}. end$  if  $L = \mathbf{t} \blacklozenge \vec{\pi}$ ; (2)  $[\mu\mathbf{t}.|L|_{\mathbf{t}}/\mathbf{t}]L$  if  $L \neq \mathbf{t} \blacklozenge \vec{\pi}$ , (3)  $L \blacklozenge (\vec{\pi}. \pi) \equiv \vec{\pi}. \pi. L$ , if  $\pi \neq \triangleright[L']$ , and (4)  $L \blacklozenge (\vec{\pi}. \pi) \equiv \vec{\pi}. (L \blacklozenge L')$ , if  $\pi = \triangleright[L']$ . The semantics of choices requires that they are directed. At the local type, all branches start with a send/receive prefix to/from the same participant  $\mathbf{p}$ . We use the predicate  $dc(\mathbf{p}, \{\ell_i\}_i, \sum_{i \in I} \pi_i. L_i)$ , and define it analogously to the predicate for global types.

► **Definition 19** (LTS for Local Types). The LTS for local types is defined as follows:

$$\begin{aligned} \text{[L-CONG]} \quad & \frac{\langle \Delta ; \mathbf{p} :: L \rangle \xrightarrow{\alpha} \langle \Delta' ; \Theta' \rangle}{\langle \Delta ; \mathbf{p} :: L, \Theta \rangle \xrightarrow{\alpha} \langle \Delta' ; \Theta', \Theta \rangle} \quad \text{[L-NEST]} \quad \frac{\langle \Delta ; \mathbf{p} :: L' \rangle \xrightarrow{\alpha} \langle \Delta' ; \mathbf{p} :: L'', \Theta \rangle}{\langle \Delta ; \mathbf{p} :: \triangleright[L']. L \rangle \xrightarrow{\alpha} \langle \Delta' ; \mathbf{p} :: \triangleright[L'']. L, \Theta \rangle} \\ \text{[L-CHOICE]} \quad & \frac{j \in I \quad \langle \Delta ; \mathbf{p} :: L_j \rangle \xrightarrow{\alpha} \langle \Delta' ; \Theta \rangle \quad dc(\mathbf{q}, \vec{\ell}, \sum_{i \in I} L_i)}{\langle \Delta ; \mathbf{p} :: \sum_{i \in I} L_i \rangle \xrightarrow{\alpha} \langle \Delta' ; \Theta \rangle} \\ \text{[L-SEND]} \quad & \langle \Delta, \mathbf{q}\mathbf{p} :: \vec{\mathbf{w}} ; \mathbf{p} :: \mathbf{q}! \ell[U]. L \rangle \xrightarrow{\mathbf{p}\mathbf{q}! \ell} \langle \Delta, \mathbf{q}\mathbf{p} :: \vec{\mathbf{w}} \cdot \ell[U] ; \mathbf{p} :: L \rangle \\ \text{[L-RECV]} \quad & \langle \Delta, \mathbf{p}\mathbf{q} :: \ell[U] \cdot \vec{\mathbf{w}} ; \mathbf{p} :: \mathbf{q}? \ell[U]. L \rangle \xrightarrow{\mathbf{p}\mathbf{q}? \ell} \langle \Delta, \mathbf{p}\mathbf{q} :: \vec{\mathbf{w}} ; \mathbf{p} :: L \rangle \\ \text{[L-NEW]} \quad & \langle \Delta ; \mathbf{p} :: \nu(\mathbf{q}_i : L_i) \cdots (\mathbf{q}_j : L_j) \rangle \xrightarrow{\mathbf{p}\mathbf{q}_i \nu L_i} \langle \Delta ; \mathbf{q}_i :: L_i, \mathbf{p} :: \nu(\mathbf{q}_{i+1} : L_{i+1}) \cdots (\mathbf{q}_j : L_j) \rangle \end{aligned}$$

[L-CONG] specifies a step by a participant in the configuration. [L-RECUR] unfolds recursion, and [L-CHOICE] selects one branch of a choice by performing a step into one of the continuations. In [L-CHOICE], only one action can take place in one branch, because the labels of all branches must be distinct for the choice to be directed. [L-SEND] executes a send prefix by enqueueing the label and the payload type into the channel of the receiver, [L-RECV] executes a receive prefix by dequeuing the label and payload type from the corresponding channel, [L-NEW] creates a new participant by composing its associated local type in parallel with the remainder of the local type environment, [L-NEST] performs a step into a nested local type. We allow the renaming of participants introduced by [L-NEW] to avoid participant name clashes. For

simplicity, we assume that  $\Delta$  always contains a (possibly empty) sequence of payloads for every pair of roles. For example, if  $\mathbf{pq}$  is not in  $\Delta$ , we allow to match  $\Delta$  with  $\Delta, \mathbf{pq} :: \epsilon$ .

We prove the correctness of DMst: (1) the global type semantics coincides with behaviours of local endpoints, a well-formed global type is (2) deadlock-free and (3) live. (1) together with (2) and (3) imply that the programs generated from local types projected from well-formed global types are deadlock-free and live.

We define *the projection of  $G$*  as  $\llbracket G \rrbracket = \langle \square ; \mathbf{p} :: G \upharpoonright \mathbf{p}, \dots, \mathbf{q} :: G \upharpoonright \mathbf{q} \rangle$ , for all  $\mathbf{p}, \dots, \mathbf{q} \in \text{pt}(G)$ . A configuration is a subtype of another if it contains the same participants and their local types are related under the standard subtyping relation [60], i.e.,  $\langle \Delta ; \Theta \rangle \leq \langle \Delta' ; \Theta' \rangle$  implies that  $\Theta(\mathbf{p}) \leq \Theta'(\mathbf{p})$  for all  $\mathbf{p}$ .

► **Theorem 20** (Trace Equivalence). *If  $\langle \Delta ; \Theta \rangle \leq \llbracket G \rrbracket$ , then  $\Gamma \vdash G \xrightarrow{\alpha^*} G'$  if and only if there exists  $\langle \Delta' ; \Theta' \rangle$  such that  $\langle \Delta ; \Theta \rangle \xrightarrow{\alpha^*} \langle \Delta' ; \Theta' \rangle$  and  $\langle \Delta' ; \Theta' \rangle \leq \llbracket G' \rrbracket$ .*

**Proof.** The full proof uses the extended projection, that produces both local types and the queue contents implicit in the intermediate forms. The core part of the proof is completed by induction on the derivations for the global and local type LTS, using the fact that if  $G \equiv G'$ , then  $G \upharpoonright \mathbf{r} \equiv G' \upharpoonright \mathbf{r}$ . ◀

A configuration  $\langle \Delta ; \Theta \rangle$  is *final* if for all  $\mathbf{pq} \in \text{dom}(\Delta)$ ,  $\Delta(\mathbf{pq}) = \epsilon$ , and for all  $\mathbf{p} \in \text{dom}(\Theta)$ ,  $\Theta(\mathbf{p}) = \text{end}$ . The configuration is in a *deadlock* if it cannot make progress and it is not final, i.e. the protocol has not ended, and all participants are stuck.

► **Definition 21** (Deadlock).  $\langle \Delta ; \Theta \rangle$  is a *deadlock configuration* if there exists a sequence of actions  $\alpha^*$  such that  $\langle \Delta ; \Theta \rangle \xrightarrow{\alpha^*} \langle \Delta' ; \Theta' \rangle$ , with  $\langle \Delta' ; \Theta' \rangle$  not final and for all action  $\alpha$ ,  $\langle \Delta' ; \Theta' \rangle \not\xrightarrow{\alpha}$ .

► **Example 22** (Deadlock Configuration). A deadlock configuration is one in which the whole system can get stuck and cannot progress. A usual example of this is a configuration where all participants need to receive, but their messages have not been sent. We show below such configuration, where after one action, it reaches a receive cycle:

$$\begin{array}{l} \langle [ \quad ] ; \mathbf{p} :: \mathbf{q}!l[U]. \mathbf{q}?l[U]. L_1, \mathbf{q} :: \mathbf{r}?l[U]. L_2, \mathbf{r} :: \mathbf{p}?l[U]. L_3 \rangle \xrightarrow{\mathbf{pq}!l} \\ \langle [\mathbf{qp} :: l[U]] ; \mathbf{p} :: \mathbf{q}?l[U]. L_1, \quad, \mathbf{q} :: \mathbf{r}?l[U]. L_2, \mathbf{r} :: \mathbf{p}?l[U]. L_3 \rangle \not\xrightarrow{\alpha} \end{array}$$

► **Theorem 23** (Deadlock-Freedom). *If  $\langle \Delta ; \Theta \rangle \leq \llbracket G \rrbracket$ , then  $\langle \Delta ; \Theta \rangle$  is deadlock-free.*

**Proof.** We show that either  $G$  is ended, or there is a step available for  $G$ , and use trace equivalence to conclude this for  $\langle \Delta ; \Theta \rangle \leq \llbracket G \rrbracket$ . ◀

Theorem 23 refers exclusively to the absence of *global* deadlocks, i.e. the whole system will never get stuck. But DMst also guarantees the absence of *local* deadlocks, i.e. that no participant in the system gets stuck. An example of such partial deadlocks is the usual *receive-cycle*, where a subset of participants are waiting forever, and can never make progress. DMst guarantees that this situation cannot happen. To prove this, we first show that DMst guarantees *orphan message freedom* [14], which means that all messages are eventually consumed without a type mismatch.

► **Definition 24** (Orphan Message).  $\langle \Delta ; \Theta \rangle$  has an *orphan message* if there exists  $\mathbf{w} \in \Delta(\mathbf{pq})$  but there exists no transition such that consumes it, i.e. there is no transition  $\langle \Delta ; \Theta \rangle \xrightarrow{\alpha^*} \langle \Delta' ; \Theta' \rangle$  with  $\mathbf{pq}?\mathbf{w} \in \alpha^*$ .



## 23:16 Dynamically Updatable Multiparty Session Protocols

► **Example 25** (Orphan Message). Orphan messages can occur whenever a send prefix is not coupled with the corresponding receive, thus leaving a message *hanging* in the corresponding buffer. For example, the following situation contains an orphan message:

$$\langle [ \quad ] ; \mathbf{p} :: \mathbf{q}!l[U]. \mathbf{q}?l[U]. \text{end}, \mathbf{q} :: \mathbf{p}!l[U]. \text{end} \rangle \xrightarrow{\mathbf{pq}!l} \langle [\mathbf{qp} :: l[U]] ; \mathbf{p} :: \mathbf{q}?l[U]. \text{end}, \mathbf{q} :: \mathbf{p}!l[U]. \text{end} \rangle \xrightarrow{\mathbf{qp}!l \cdot \mathbf{pq}?l} \langle [\mathbf{qp} :: l[U]] ; \mathbf{p} :: \text{end}, \mathbf{q} :: \text{end} \rangle$$

At the end of the execution, the configuration contains a non-empty buffer:  $\mathbf{qp} :: l[U]$ .

Another example of orphan messages is one in which the reduction gets stuck because of receiving a message of the wrong type or label, i.e. there is a reception error.

$$\langle [ \quad ] ; \mathbf{p} :: \mathbf{q}!l[\text{int}]. \text{end}, \mathbf{q} :: \mathbf{p}?l[\text{bool}]. \text{end} \rangle \xrightarrow{\mathbf{pq}!l} \langle [\mathbf{qp} :: l[\text{int}]] ; \mathbf{p} :: \text{end}, \mathbf{q} :: \mathbf{p}?l[\text{bool}]. \text{end} \rangle \not\rightarrow$$

In this case, reduction cannot continue, and the message  $\mathbf{qp} :: l[\text{int}]$  cannot be consumed, because  $\mathbf{q}$  is expecting payload type `bool`.

Proving that DMst guarantees the absence of orphan messages relies on the absence of *blocked local types*. A blocked local type is a local type that contains a nested session that cannot terminate, followed by a non-empty continuation. For example, if  $L = \triangleright[\mu\mathbf{t}.\mathbf{q}_2!l'[U']. \mathbf{t}]. \mathbf{p}?l[U]$ , then  $L$  is blocked, because it will enter the nested protocol (with local type  $\mu\mathbf{t}.\mathbf{q}_2!l'[U']. \mathbf{t}$ ), but it will never be able to continue executing  $\mathbf{p}?l[U]$ .

► **Definition 26** (Blocked Participant). A *blocked* local type is one that contains a continuation of the form  $\triangleright[L_1]. L_2$ , where: (a)  $L_1$  is blocked, or (b)  $L_2 \neq \text{end}$  and  $\text{end}$  is not reachable from  $L_1$  (see Def. C.8).

► **Definition 27** (Liveness). We say that  $\langle \Delta ; \Theta \rangle$  is *live*, if no participant is stuck. A participant  $\mathbf{p}$  is stuck in a configuration whenever it cannot progress, i.e. if  $\Theta(\mathbf{p}) = L$  with  $L \neq \text{end}$ , but there is no trace  $\langle \Delta ; \Theta \rangle \xrightarrow{\alpha^*} \langle \Delta' ; \Theta' \rangle$  with  $\mathbf{p} = \text{subj}(\alpha)$  and  $\alpha \in \alpha^*$ .

► **Example 28** (Stuck Participant). The following configuration is *not live*, because even if  $\mathbf{p}$  and  $\mathbf{q}$  can continue interacting,  $\mathbf{r}$  and  $\mathbf{s}$  are stuck in a local receive cycle:

$$\langle [ \quad ] ; \mathbf{p} :: \mu\mathbf{t}.\mathbf{q}!l[U]. \mathbf{t}, \mathbf{q} :: \mu\mathbf{t}.\mathbf{p}?l[U]. \mathbf{t}, \mathbf{r} :: \mathbf{s}?l[U]. L_3, \mathbf{s} :: \mathbf{r}?l[U]. L_4 \rangle \xrightarrow{\mathbf{pq}!l \cdot \mathbf{qp}?l} \langle [ \quad ] ; \mathbf{p} :: \mu\mathbf{t}.\mathbf{q}!l[U]. \mathbf{t}, \mathbf{q} :: \mu\mathbf{t}.\mathbf{p}?l[U]. \mathbf{t}, \mathbf{r} :: \mathbf{s}?l[U]. L_3, \mathbf{s} :: \mathbf{r}?l[U]. L_4 \rangle \xrightarrow{\alpha^*} \dots$$

No possible trace can contain  $\mathbf{rs}?l$  or  $\mathbf{sr}?l$ . Participants  $\mathbf{r}$  and  $\mathbf{s}$  are stuck. Note that, from Definition 19, only receive prefixes can get stuck, since send prefixes will always succeed.

► **Theorem 29** (Orphan Message Freedom and Liveness).

*Suppose  $\langle \Delta ; \Theta \rangle \leq \llbracket G \rrbracket$ , such that  $\Theta$  contains no blocked participants. Then  $\langle \Delta ; \Theta \rangle$  is free of orphan messages and live.*

**Proof.** Liveness is a straightforward consequence of orphan message freedom. The prefix of a local type can have two kinds of actions: outputs (sending data or invitations), or inputs (receiving data, or accepting invitations). Every input is coupled with an output by another participant (see Definition 13). Hence outputs can always be performed, in any state. To prove that inputs can always be consumed, we use trace equivalence. We show that any pending message can be received, since a step can only happen in a continuation if its subject is not in any of the previous prefixes, and it is always possible to end nested protocols, because they cannot be blocked. ◀

► **Proposition 30.** *As a consequence of Theorems 23, 29 and 20, the global types of Example 10 are live and deadlock-free.*

## 4 GoScr Code Generation

This section describes the GoScr toolchain. GoScr is an extension of nuScr [49], which is a new implementation of Scribble in OCaml. nuScr is designed with modularity and extensibility in mind, so that extensions of the core MPST theory [60] can be easily integrated.

**GoScr Global Protocols.** The syntax of GoScr global protocols is given in Definition 31.

► **Definition 31** (GoScr syntax).

$$\begin{aligned}
 P &::= \text{global protocol } x(\text{role } \mathbf{p}_1, \dots, \text{role } \mathbf{p}_n; \text{new role } \mathbf{q}_1, \dots, \text{role } \mathbf{q}_m) \{P^* G\} \\
 g &::= m[U] \text{ from } \mathbf{p} \text{ to } \mathbf{q} \mid \mathbf{p} \text{ calls } x(\mathbf{p}_1, \dots, \mathbf{p}_n) \\
 G &::= \text{choice at } \mathbf{p} \{G_1\} \text{ or } \dots \text{ or } \{G_n\} \mid g; G \mid \text{rec } \mathbf{t} \{G\} \mid \text{continue } \mathbf{t} \mid \text{end} \\
 &\quad \mid \text{continue } \mathbf{t} \text{ with } \{g_1; \dots g_n; \mathbf{p} \text{ calls } x(\mathbf{p}_1, \dots, \mathbf{p}_n)\} \mid \text{do } x(\mathbf{p}_1, \dots, \mathbf{p}_n); G
 \end{aligned}$$

A GoScr module is a sequence of one or more `global` protocols. The last protocol definition is the entry point. The constructs of GoScr were chosen to mirror those defined in Definition 1: `global protocol` are protocol definitions ( $x = \lambda\langle\vec{p}; \nu\vec{q}\rangle.G \dots$ ; `global protocol` can be used for protocol declarations with no new participants; `choice at p` defines directed choices from `p` to the receiver of the first interaction in the  $G_i$ ; `do` is a protocol call to a global protocol; and the rest of the constructs correspond to those of DMst. Protocol definitions in GoScr can start by defining other nested protocols, but this is simply a syntactic convenience, since we require every role in a nested protocol to be bound by the protocol signature.

**Steps for Code Generation.** The steps of code generation in GoScr are: (1) lifting all nested protocol definitions to the top-level; (2) obtaining the projections of all roles in all protocol definitions; (3) preprocessing local types to deal with instances of  $\blacklozenge$  (or `continue...with...` in GoScr); and, (4) translating the local types to Go functions, where communication is implemented using Go channels, interleaved with *callbacks* that will be used to implement the program logic.

Step (1) is straightforward. Step (2) is an implementation of Definition 13. Step (3) requires applying local type equivalences to unfold any updatable recursion. Step (4) traverses the local types, and generates on demand the necessary channels and callback interfaces. The type of the Go channels is an interface that represents the allowed payload types. Then, for each labelled message exchange: (1) we add a new type declaration for the label and payload type that implements the interface of allowed messages; (2) we search for a channel for the required endpoints, creating it if necessary; (3) we create the necessary callbacks before or after the interaction. The channels can be created either synchronous, or buffered with a user-specified size. Choosing synchronous channels is safe, since the traces accepted by using synchronous semantics is a subset of those accepted by our asynchronous semantics, which implies that the same safety properties will hold.

**Go Code Generation.** The Go code for each role and protocol is generated in `protocol/`. Communication is implemented using regular Go send/receive statements. There is no need to explicitly send message labels, since labels are encoded as type declarations. Protocol choices are encoded as type switches, either on the value returned from a previous callback (internal choices), or on the received value (external choices). We only generate implementations for branching choices that start with an explicit interaction.

Calling a nested protocol is implemented as regular Go function calls. `rec` constructs are generated as labelled `for` loops, where the body of the recursion is used to generate the body of the `for` loop, and recursive variables are translated as `continue` to the label of the corresponding variable. It is also possible to represent recursion using protocol calls.

## 23:18 Dynamically Updatable Multiparty Session Protocols

```
1 func BFib_F2(ctx Ctx_BFib_F2, wg *sync.WaitGroup, ch_F2_F3, ch_F3_F2 chan MsgBFib) {
2     x := ctx.Send_F3_BFib_Fib2() // Callback to generate payload
3     ch_F3_F2 <- x // Send payload to F3
4     x_1 := <- ch_F2_F3 // External choice by from F3
5     switch v := x_1.(type) {
6     case End: // F3 chooses to finish the protocol
7         ctx.Recv_F3_BFib_End(v) // Callback for processing label End
8         ctx.End()
9         return
10    case Call_F1_BFib: // F3 sends the channel for acting as F1 in BFib
11        ctx_1 := ctx.Init_F1_BFib_Ctx() // Initialise context for F1 in BFib
12        BFib_F1(ctx_1, wg, v) // Run code for F1 in BFib with channel [v]
13        ctx.End_F1_BFib_Ctx(ctx_1) // Close context for F1 in BFib
14        ctx.End()
15        return
16    } }
```

■ **Figure 6** Implementation of role **F2** in protocol **BFib**

However, protocol calls would need to create the necessary channels and send them to any participant in the protocol, thus being less efficient than using `rec` and `for`.

### 4.1 Linearity and CFSM Code Generation

Program logic is defined through *callbacks*, similar to [43, 62], to avoid the *linearity* problem of previous Communicating Finite State Machine approaches (e.g. [7]). In a CFSM approach, code generation from a local protocol produces a series of interfaces that encode the protocol states. Each protocol state exposes *only* the permitted actions (e.g. send/receive), and returns the next state in the protocol. Programmers must use such states to implement their program logic. The linearity problem arises from the fact that nothing prevents programmers from mistakenly using the same protocol state again. For example, suppose that `st0`, `st1`, `...`, are protocol states that expose different send/recv actions. A programmer might (mistakenly) save state `st1` and perform its action twice in the implementation. In the Go code snippet below, `st1` is used both in Line 2 and Line 4, violating linearity:

```
1 st1 := st0.send_Msg_to_p(x)
2 st2 := st1.recv_Lbl_from_p(&z)
3 ...
4 stn := st1.recv_Lbl_from_p(&buffer) /* linearity error at st1 */
```

If participant `p` does not send any other message, then this implementation will **deadlock**. If `p` does send another message, this might cause a run-time error. A callback-based approach solves this problem *by construction*, since channels are not exposed to programmers [43, 62].

### 4.2 Example of Generated Go Code

Consider the following GoScr global type:

```
1 global protocol BFib(role Res, role F1, role F2; new role F3) {
2     Fib1(v:int) from F1 to F3; Fib2(v:int) from F2 to F3;
3     choice at F3 {
4         F3 calls BFib(Res, F2, F3);
5     } or {
6         Result(fib:int) from F3 to Res; End() from F3 to F2; }}
```

This is a *bounded* version of Example 4, that computes the Fibonacci sequence up to an upper bound. `F1` and `F2` send their respective  $n - 2$  and  $n - 1$  Fibonacci numbers to `F3`. Then, `F3` computes the  $n$ -th number, and makes a choice: compute the  $n + 1$  number, or end the protocol. If `F3` decides to continue, then a recursive call to `BFib` happens. Otherwise, it sends

the result to `Res`, and notifies `F2` that the protocol is ending. `F3` needs to notify `F2`, because depending on `F3`'s decision, `F2` may need to forward its  $n - 1$  number.

Figure 6 shows the code for `F2` in `BFib`. The parameters of `BFib_F2` are: `ctx` is the local state for `F2`; `wg` is used to ensure that the main thread does not resume execution until all participants have finished executing; `ch_A_B` is the channel for communicating from `B` to `A`. The first interaction of `F2` is a message to `F3`. The payload for this message is generated in Line 2, it is sent in Line 3. Then, `F3` makes a choice: either it sends the result back to `Res` and sends `End` to `F2` to communicate the end of the protocol, or it calls `BFib` recursively. `F2` performs a type switch to check which branch it needs to take (Line 4). If the label it receives is `End` (Line 6), then `F2` processes this label and ends the protocol. Otherwise, `F2` receives an invitation as `F1` in `BFib` (Line 10); then `F2` initialises a new context for `F1` using the callback on Line 11; it calls `BFib_F1` with this new context, the waitgroup, and the received channel (Line 12); `F2` performs cleanup on the context for `F1`, gathering any necessary results from the call (Line 13); and, finally `F2` finishes.

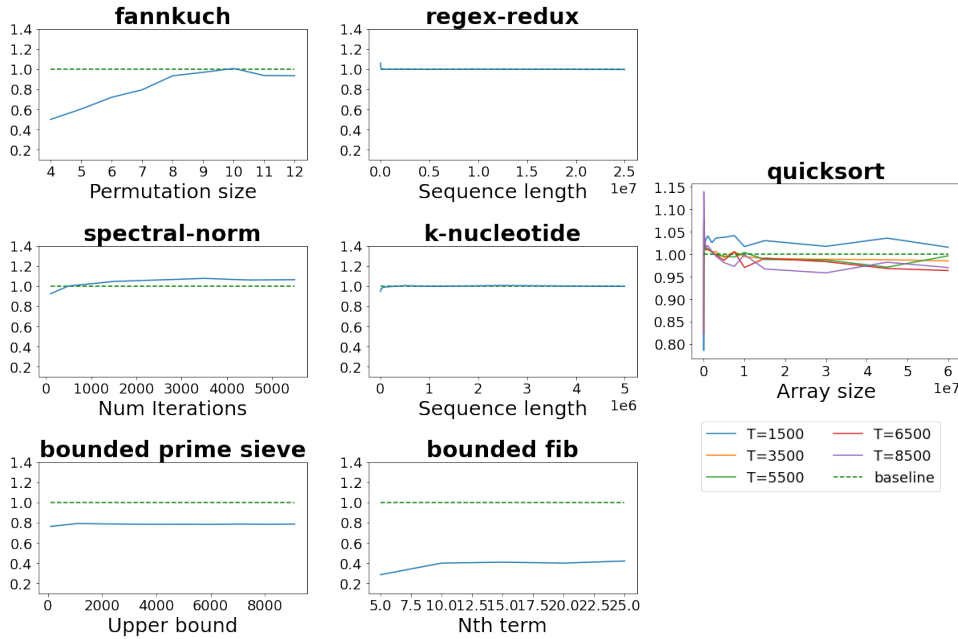
Finally, `GoScr` generates the main protocol entrypoint, which creates the goroutines for `F1`, `F2` and `F3`, all the needed channels, and waits for the completion of the protocol.

**Usability and GoScr Front-end.** The tool requires the user to instantiate a large number of callbacks and interfaces to allow running a protocol. Since the `GoScr` methodology is top-down, the user must start by specifying a protocol. Therefore we expect an end-user to be aware of the callbacks and contexts that need to be instantiated. However, many of such instantiations are tedious, but straightforward, and can be automated in future work. We discuss this improvement in Section 7.

**Deadlock Freedom and Liveness.** Since the generated code follows the behaviour of the local types, it will satisfy both **deadlock freedom** and **liveness** (Theorems 23 and 29). Although the generated code satisfies these properties, whether the final code that is run also satisfies them depends on three requirements on the callbacks. These requirements are not checked by `GoScr`, and must be guaranteed by `GoScr` users. The three requirements that the callbacks must satisfy are: **1.** callbacks must not have side-effects that interfere with other participants (e.g. using channels to add communication that is not accounted for in the protocol) **2.** callbacks must be terminating, otherwise a participant may block before a necessary interaction, in a non-terminating callback; and, **3.** callbacks must ensure that nested protocol calls that are not in tail position are terminating. Requirement **1** is to guarantee that programmers do not use local synchronisation mechanisms that are not accounted for in the protocol, and can cause blocking. Requirement **3** is to guarantee that any interaction after a nested protocol call is eventually performed. `GoScr` checks that local types are not blocked (Definition 26), so the code for nested calls that are not in tail position will always contain a path that ends the protocol. However, whether the actual code is terminating depends on the callback implementation that the users need to provide satisfying Requirement **3**. Provided that these requirements are met, and assuming a fair scheduler, `GoScr` implementations will be deadlock free and live by construction. These requirements are not unique to our implementation. Similar requirements must be satisfied in other MPST code generation approaches.

## 5 Evaluation

We evaluate three aspects of `GoScr`: **1.** the runtime overhead of the `GoScr` backend (§ 5.1); **2.** the increased expressiveness with respect to related approaches (§ 5.3); and **3.** the applicability of `GoScr` for building realistic protocols, by implementing dynamic task delegation, a



■ **Figure 7** Execution time comparison ( $t_{base} / t_{GoScr}$ ), CLBG and Quicksort.

Domain Name System, and a parallel Min-Max strategy. We show that for computation-intensive protocols, the runtime overhead of GoScr is negligible.

## 5.1 Runtime Overhead of GoScr

We use an Intel(R) Core(TM) i7-4770 CPU @ 3.40GHz processor with 4 physical cores, 16GB RAM, running Ubuntu 16.04.7 and Go version `go1.15.11`. We use Golang’s time package to measure execution times. There are two main sources of run-time overheads: (1) callbacks; and (2) type switches and assertions. Our approach is to compare GoScr implementations against baseline Go code. Baselines are taken from benchmarking repositories, and follow similar communication patterns to the GoScr implementations. The measured time includes session initialisation. We execute each benchmark for a minimum of 20 iterations and a minimum of 20 seconds. The standard deviation for computationally expensive benchmarks is less than 5%. Only the standard deviation of fibonacci and prime sieve with small inputs ( $< 10$ th term, bound  $< 2000$ ) remain high, at 70%. This is because these benchmarks with very short execution times (in the order of nanoseconds) are highly dependent on the system (e.g. channel creation, goroutine scheduling, etc). Our benchmarks are mainly taken from the *Computer Language Benchmarks Game* [19], and we include a parallel Quicksort that showcases the handling unbalanced workloads. Figure 7 shows the execution time of the Go baseline relative to GoScr:  $t_{base} / t_{GoScr}$  (below  $y = 1$  is a slowdown, above is a speedup).

**Computer Language Benchmarks Game (CLBG).** CLBG [19] is a repository of programs used to compare the performance of different languages. We use four concurrent Go programs: (a) `fannkuch` counts the maximum number of flips for a permutation of length  $n$ ; (b) `regex` matches regex patterns in a DNA string; (c) `spect` (spectral-norm) calculates the greatest eigenvalue of a matrix; and (d) `k-nuc` (k-nucleotide) counts the occurrences of a molecule sequence in a DNA string. We selected these benchmarks out of [19] because they parallelise the work using goroutines and channels, following a similar scatter/gather approach that depends on runtime values, and they could not be accurately captured by previous MPST approaches. We use the CLBG implementations [19] as the Go

```

1  global protocol DynTaskGen(role S;
2     new role W) {
3     choice at S {
4         Req(req: string) from S to W;
5         S calls DynTaskGen(S);
6         choice at W {
7             Resp(resp: string) from W to S;
8         } or {
9             Error(err: string) from W to S; }
10    } or {
11        LastReq(req: string) from S to W;
12        choice at W {
13            Resp(resp: string) from W to S;
14        } or {
15            Error(err: string) from W to S;
16        }}}
11 global protocol ClientServer(role C,
12     role S) {
13     rec REPEAT {
14         Req(req: string) from C to S;
15         S calls DynTaskGen(S);
16         choice at S {
17             Resp(resp: string) from S to C;
18             continue REPEAT;
19         } or {
20             Error(err: string) from S to C;
21             continue REPEAT;
22         }}}

```

■ **Figure 8** GoScr protocol for Dynamic Task Generation

baseline implementations, and we extracted the communication structure of the baseline implementations as GoScr protocols. A single execution for each of these protocols takes between 1 millisecond–10 seconds depending on the input size. Smaller input sizes imply smaller local computation times, and therefore, the overhead introduced by GoScr will be more significant. We can observe a slowdown of up to 50%, in `fannkuch`, for executions in the order of magnitude of milliseconds. However, as the workload increases, the difference in the execution time shrinks to the point of becoming negligible, as we can observe in Figure 7. The `regex` baseline has a high standard deviation, which explains the small peak for the first result of `regex`, since when the execution time is in the order of hundreds of microseconds, the non-deterministic scheduling of the goroutines can significantly affect the results. `spect` seems to show that for large enough values, the GoScr implementation performs better than its naively implemented counterpart. However, the real difference in the execution time is negligible, and it is explained by differences in the program structure, e.g. the baseline uses a single shared channel, whereas GoScr generates different channels for every new goroutine.

**Microbenchmarks.** Bounded fibonacci (`fibonacci`) shows, as expected, that the overhead of performing type switches and callbacks is relatively high when compared with a simple addition. The baseline runs in 40% of the execution time of GoScr. Bounded prime sieve (`prime`) shows that, when the computation complexity increases slightly (modulus operation on a stream of values), then the GoScr version performs in about 80% the execution time of the baseline. In both cases, when we add more participants and interactions to the protocol (larger values on the x-axis) the overhead *remains constant, and does not increase*.

**Unbalanced Workload.** In *Parallel QuickSort* (`qsort`), workers either partition the array and spawn two new workers, or apply a sequential Quicksort, depending on a threshold size ( $T$ ). The execution times are similar to the CLBG benchmarks (50 microseconds–2 seconds). We observe a negligible difference in the execution time for different threshold sizes, and a spike for small arrays due to the high standard deviation for array sizes under the threshold. GoScr execution times are in the range of 1.05 and 0.95 times the baseline.

## 5.2 Use Cases

We demonstrate the expressiveness of GoScr using three applications, all of which require dynamic participants, and could not be expressed by previous work [7, 38].

- (a) **Dynamic Task Generation:** We present a correct implementation of the program in Figure 1 using GoScr. It is a master-worker pattern with dynamic participants.

## 23:22 Dynamically Updatable Multiparty Session Protocols

Protocol	Dyn	Unb	Inv	DMst	[7]	[38]	[11]	[13]
1. Dynamic Ring	●		●	✓	X	X	✓	X
2. Dynamic Pipeline	●		●	✓	X	X	✓	X
3. Dynamic Recursive Pipeline	●	●		✓	X	X	X	X
4. Dynamic Recursive Tree	●	●		✓	X	X	X	X
5. Dynamic Recursive Task Gen.	●	●		✓	X	X	X	X
6. Dynamic Fork-Join	●			✓	X	X	✓	X
7. Recursive Fork-Join	●			✓	X	X	✓	X
8. Bounded Fibonacci [7]	○		○	✓	△	△	✓	X
9. Unbounded Fibonacci	●	●		✓	X	X	✓	X
10. Fannkuch-Redux [19]	○		○	✓	△	△	✓	✓
11. Spectral-Norm [19]	○			✓	△	△	✓	✓
12. Regex-Redux [19]	○			✓	✓	✓	✓	✓
13. K-Nucleotide [19]	○			✓	✓	✓	✓	✓
14. Bounded Prime Sieve	●			✓	X	X	✓	X
15. Dynamic Task Generation	●	●		✓	X	X	✓	X
16. Domain Name System [29]	●		●	✓	X	X	✓	X
17. Noughts and Crosses [43, 50]	●		●	✓	X	X	✓	X

Dyn: Dynamic participants; Unb: Unbounded participants; Inv: Choice through invitations

■ **Table 1** Comparison of Expressiveness

- (b) **Domain Name System (DNS) protocol** (§ D.0.2): We demonstrate how GoScr can be used to specify one of the core Internet protocols, modelling as dynamic participants the different DNS servers which may need to be contacted in order to resolve a host’s IP address.
- (c) **Noughts and Crosses with Min-Max [50]** (§ D.0.3): We implement a Min-Max strategy for the well-known two-player game of Noughts and Crosses to demonstrate the suitability of DMst to model a parallel Divide and Conquer paradigm.

**Dynamic Task Generation.** The aim of this program is to generate the first  $n$  square numbers by delegating the calculation of each square number to a different worker goroutine. The program uses a common computation in Go, the *master-worker pattern*, where goroutines dynamically divide and delegate part of their tasks to other goroutines, aggregating their partial results to produce the complete result. We highlighted in § 1 (Figure 1) how even in such a simple example, incorrect management of channels can lead to orphan messages and deadlocks. Figure 8 shows a GoScr protocol specification whose behaviour is a safe version of the program in Figure 1. Notice how the behaviour of the `select` statement in Figure 1 is represented as a choice. In Figure 8, the `clientServer` protocol models the behaviour of the main loop of the program, where two roles, a client and a server, repeatedly exchange requests (Line 14) and responses (Line 17). The server may also communicate an error in the computation of the request to the client (Line 20). We model the master-worker pattern as a call to protocol `DynTaskGen` (Line 15). Every call to the protocol introduces a new worker (`w`), and the master (`s`) will delegate a task to each new worker (Lines 4,11). If there are more tasks to assign, it will assign those tasks to new workers through recursive calls to `DynTaskGen` (Line 5). Once it has assigned the final task (Line 11), it will traverse the protocol stack, aggregating the results from the different workers in reverse order (Lines 7,13). While computing their subtask, the workers may encounter an error which they will communicate back to the server (Lines 9, 15). As opposed to the original program in Figure 1, the server will continue aggregating all the results from the workers even after encountering an error in order to ensure that there are no orphan messages.

### 5.3 Expressiveness

We compare the expressiveness of GoScr against the parameterised `Scribble` [7] and the static analysis framework of Go [38]. For a reference purpose, we also list comparisons with **theory-only work** in [11, 13] (i.e., they are *not* implemented). See § 6 for more detailed

comparison with [11, 13]. In Table 1, we present the protocols that we implemented and whether or how closely other approaches [38, 7] can represent them. All our DMst-based implementations introduce dynamic, possibly unbounded participants. All representable protocols (✓) by DMst in Table 1 are deadlock-free and live. For protocols which can be modified and re-implemented with [38] or [7], we use ○. Protocol 3 cannot be captured by any of the previous work, since it requires the dynamic introduction of participants to a recursive protocol. [7, 38] could only precisely model Protocols 12 and 13, as they create all the participants at the start. In Protocols 8, 10 and 11, the goroutines are spawned and assigned tasks dynamically, but [7, 38] can model them by initialising all goroutines at the start. We write  $\Delta$  to represent such changes to protocol structure. Three use cases (Protocols 13–15) discussed in § 5.2, could not be expressed by [7, 38]. In summary, DMst is more expressive than [7, 38], and capture more closely the typical Go programming style.

## 6 Related Work

There are a vast amount of studies of session types [28, 18, 5]. Due to the space limitations, we only compare with the most closely related work on multiparty session types (MPST).

**Binary Session Types.** While Scalas et al. [51] prove that the MPST processes can be mimicked by linearly typed processes with a continuation-passing style translation, in general, it is not possible to guarantee deadlock-freedom for more than two interleaved binary session processes unless one uses additional sophisticated means such as a global causal analysis on channels (e.g. [16, 8, 9]), graph-connectivity analysis with extensions on fork primitives [30], and event-driven constructs [57, 25, 35]. GV, a linear functional language with binary session types, can guarantee deadlock freedom by relying on linear typing [58]. However, linear typing prevents cyclic topologies that change dynamically, since this would require a participant to drop their communication channels when new participants join, as in Example 7. There are further substantial differences with our approach. First, GV is an end-point calculus, whereas DMst’s global types are *global specifications*, from which we can extract endpoint Go code (GoScr). Secondly, while both GV and DMst support similar programming patterns (e.g. pipeline and tree-like topologies, and channel passing), there are two major differences. Both GV and GoScr support sending effectful functions over channels (e.g. using `chan func()` type in Go, and passing a generated protocol implementation), GV’s type system would guarantee deadlock-freedom, but in Go, it would depend on how the function is used (requirements 1-3 in Section 4).

**Code Generation and Multiparty Session Types.** We follow the standard MPST top-down *specification-guided* methodology to guarantee **safety and liveness properties by construction using code generation**, extending an extensible toolchain, nuScr [49]. Safety by construction via code generation is a common approach in MPST. Scribble is a language/tool [52, 49] used for generating APIs for safely implementing distributed systems written in the end-point programming language that are guaranteed to conform to a protocol, and are therefore deadlock-free [26]. This approach has been applied to several languages, e.g. Scala [51, 57], Java [34], F# [44], Go [7], TypeScript [43], F★ [62] and Rust [36, 10]. A later extension of [26] proposed *explicit connection actions* as part of the Scribble protocol [27], which is also recently applied to domain-specific language in [20]. This construction specifies the point in the protocol where the different participants join, but the role of these participants must be statically known. Hence it does not allow the unbounded participants to change the protocol topology, as DMst does. *Parameterised multiparty session types* extend MPST with a parametric number of participants [15]. One example is the work by



Castro-Perez et al. [7], discussed it in § 1. *Pabble* [47, 46] is another parameterised extension of Scribble used for generating safe by construction C+MPI code. Zhou et al. [62] formalised and implemented an extension of MPST with *refinement types*, which can specify constraints in the messages. Their backend targets  $F^*$ , and follows a similar callback approach to the one in this paper. Miu et al. [43] define an extension of MPST for web programming in TypeScript that uses the callback approach. Unlike DMst, the participants in all these approaches are fixed from the start of the protocol. Viering et al. [57] present a theory and implementation of MPST aimed at programming correct fault-tolerant distributed systems that supports the dynamic replacement of participants in a protocol. In their work, the replacement of participants must happen within some known roles, and their global types do not allow to extend the current protocol interactions with those of new participants. Viering et al. [57] use event handlers in their code generation, which allows safe session interleaving. Instead, we use an operator to *combine* global types in a way that does not introduce deadlocks. All previous work, unlike DMst, does not support dynamically growing protocols with an unbounded number of participants such as Example 6. Jacobs et al. [31] extend GV, a binary session typed calculus with multiparty session types. The calculus allows the introduction of new participants, but the protocols themselves are restricted to a fixed set of participants. Their use of linearity prevents the definition of recursive dynamic topologies, unlike DMst.

**Dynamic Multiparty Session Types.** Dynamic multirole session types (MRST) enable a set of participants which belong to the same group (i.e. role) to join a multiparty session type [13]. The major limitations are: (a) all the roles are fixed at the start (b) participants can only join at specific points in the protocol: (1) at the beginning of each iteration of a recursive protocol; or (2) at particular points marked with explicit barriers and locks. We list a number of protocols that cannot be represented using MRST in Table 1. In contrast, DMst allows any arbitrary role to join at any nested session call. A nested session call is a form of delegation, which is not supported by MRST. Therefore, a protocol such as a dynamically growing pipeline (e.g. Fibonacci in Example 4) cannot be represented by [13] either, since it would require participants to evolve their behaviour through channel passing. Nested multiparty session types [11] allow multiparty protocols with unbounded, dynamic participants. However, [11] cannot represent *recursive* protocols that are *updated* with new dynamic participants. Hence the main example of this paper, Example 6, is not representable in [11]. Moreover, nested multiparty session types cannot prove liveness (our Theorem 29), except for non-recursive protocols. Arbitrary session interleaving in [11] can introduce orphan messages. DMst has proven deadlock-freedom and liveness clearly identifying the conditions (Definition 26). This limitation is stated in [11, Proposition 3], i.e. a protocol that violates liveness will be accepted in [11], but not in DMst. Additionally, the theory of DMst has a number of differences that make it better suited for implementing than nested MPST: (a) DMst’s choices are more flexible than those in nested MPST, since DMst can also depend on protocol calls; (b) the semantics of nested MPST is synchronous, while DMst is asynchronous; (c) nested MPST does not prove trace equivalence between global types and local type configurations; (d) The syntax of DMst’s global types are simpler than those in nested MPST, but more expressive – this is because in nested MPST, protocol definitions are part of the global type syntax, which requires the use of a kinding relation for checking well-formedness. Nested MPST protocols do not allow the occurrence of free roles, and are therefore equivalent to DMst’s global types with just top-level protocol definitions, which avoids the kinding relation for checking well-formedness. Due to our simpler but more expressive treatment, DMst is more suitable for real language implementations.

**Verification of Go Programs.** Our work aims at providing *correctness by construction*.

The comparison with the previous code generation approach in Go [7] can be found in (C) in § 1 and *Expressiveness* in § 5.3. All of the previous work is limited to *bounded* participants. The following are several recent lines of work on *a posteriori* verification of message passing in existing Go programs. All of them use *whole-program* techniques, and support only the built-in Go channel *primitives* (i.e., intra-process messaging); none of them, however, support a dynamic, unbounded number of participants. Gobra is an automated tool for the modular verification of Go programs, based on separation logic [59]. Gobra is aimed at the functional verification of Go programs, whereas our approach focuses on communication safety. GoScr is fully automated, and aimed at building live and deadlock-free communicating systems by construction. In contrast, Gobra is aimed at the verification of annotated Go code, and it requires a high amount of invariant annotations.

Ng and Yoshida [48] extract graph-based protocol specifications [39] from Go programs that are checked for deadlock-freedom; Stadtmüller et al. [53] extract regex-based protocol specifications [55], checked for deadlock-freedom. Both approaches work only for programs restricted to *synchronous* Go channels; the former also requires all goroutines to be spawned before any communication among them occurs, and the latter has limited support for branching behaviours. Lange et al. [37, 38] (already compared in (B) in § 1 and *Expressiveness* in § 5.3) statically infer channel communication patterns from Go programs as *behavioural types*, that are checked for liveness properties. This was recently extended to analyse shared memory concurrency [17]. Like previous work, their tool is also limited to verify finite controlled programs, it is best-effort only due to the imprecision of the inference, and the verification times (and timeouts) preclude practical checking on the fly during programming. Liu et al. [40] present a tool that detects blocking misuse-of-channel bugs in Go and produces bug fixes for Go programs. Unlike DMst, Liu et al. [40] focuses only on practice, and does not formalise nor guarantee communication safety, deadlock-freedom nor liveness. Moreover their tool produces both false positive and false negative errors.

## 7 Conclusion and Future Work

GoScr is the first implementation of multiparty session types with dynamic, unbounded participants, from which we generate Go code with unbounded participants that is, *by construction*, deadlock-free and live. GoScr focuses on correctness (Theorems 20, 23 and 29), and it is strictly more expressive than previous Go verification frameworks (see § 1, Table 1, § 6). Furthermore, we observe that whenever the computation time is large with respect to the communication time, the performance overhead becomes negligible. GoScr is therefore suitable for implementing systems where correctness is prioritised, or systems where the computation times dominate over communication. Currently, DMst does not allow a participant to communicate with an unbounded number of participants during protocol execution. This is a limitation of the Go code generation, which we plan to address in future work. We are also considering extending our back-end to use event-handlers in the style of Viering et al. [57], and allow the arbitrary parallel composition of global types instead of our combination operator. We are also planning to extend the back-end to disparate transports (e.g. using TCP instead of Go channels), thus allowing the implementation of distributed systems. The main challenge of this is integrating delegation, as it is required by protocol invitations, in these disparate transports. Finally, to simplify usability, we plan to extend the protocol specification with annotations to guide code generation, so we can automatically generate trivial callback/context instantiation. We plan to draw inspiration for such annotations from *choreographies*, e.g [32].

## References

- 1 Docker: Empowering app development for developers. <https://www.docker.com/>, November 2020.
- 2 Github 2.0: A small place to discover languages in github. [https://madnight.github.io/github/#/pull\\_requests/2020/3](https://madnight.github.io/github/#/pull_requests/2020/3), 2020.
- 3 gRPC - a high-performance, open source universal rpc framework. <https://grpc.io/>, November 2020.
- 4 Stack overflow developer survey 2020. <https://insights.stackoverflow.com/survey/2020>, 2020.
- 5 Davide Ancona, Viviana Bono, Mario Bravetti, Joana Campos, Giuseppe Castagna, Pierre-Malo Deniélou, Simon J. Gay, Nils Gesbert, Elena Giachino, Raymond Hu, Einar Broch Johnsen, Francisco Martins, Viviana Mascardi, Fabrizio Montesi, Romyana Neykova, Nicholas Ng, Luca Padovani, Vasco T. Vasconcelos, and Nobuko Yoshida. Behavioral types in programming languages. *Foundations and Trends in Programming Languages*, 3(2-3):95–230, 2016. doi:10.1561/25000000031.
- 6 Lorenzo Bettini, Mario Coppo, Loris D’Antoni, Marco De Luca, Mariangiola Dezani-Ciancaglini, and Nobuko Yoshida. Global progress in dynamically interleaved multiparty sessions. In *CONCUR 2008 - Concurrency Theory*, pages 418–433. Springer, 2008.
- 7 David Castro-Perez, Raymond Hu, Sung-Shik Jongmans, Nicholas Ng, and Nobuko Yoshida. Distributed programming using role-parametric session types in go. *POPL’19*. ACM, 2019. doi:10.1145/3290342.
- 8 Mario Coppo, Mariangiola Dezani-Ciancaglini, Luca Padovani, and Nobuko Yoshida. Inference of Global Progress Properties for Dynamically Interleaved Multiparty Sessions. In *15th International Conference on Coordination Models and Languages*, volume 7890 of *LNCS*, pages 45–59. Springer, 2013.
- 9 Mario Coppo, Mariangiola Dezani-Ciancaglini, Nobuko Yoshida, and Luca Padovani. Global Progress for Dynamically Interleaved Multiparty Sessions. *MSCS*, 26:238–302, 2015.
- 10 Zak Cutner, Nobuko Yoshida, and Martin Vassor. Deadlock-Free Asynchronous Message Reordering in Rust with Multiparty Session Types. In *27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, volume abs/2112.12693. ACM, 2022.
- 11 Romain Demangeon and Kohei Honda. Nested protocols in session types. In *CONCUR 2012 - Concurrency Theory*, pages 272–286. Springer, 2012.
- 12 Romain Demangeon, Kohei Honda, Raymond Hu, Romyana Neykova, and Nobuko Yoshida. Practical interruptible conversations: distributed dynamic verification with multiparty session types and python. *Formal Methods in System Design*, 46(3):197–225, 2015. doi:10.1007/s10703-014-0218-8.
- 13 Pierre-Malo Deniélou and Nobuko Yoshida. Dynamic multirole session types. *POPL’11*, pages 435–446. ACM, 2011.
- 14 Pierre-Malo Deniélou and Nobuko Yoshida. Multiparty compatibility in communicating automata: Characterisation and synthesis of global session types. In *ICALP 2013*, volume 7966 of *LNCS*, pages 174–186. Springer, 2013. doi:10.1007/978-3-642-39212-2\_18.
- 15 Pierre-Malo Deniélou, Nobuko Yoshida, Andi Bejleri, and Raymond Hu. Parameterised multiparty session types. *Logical Methods in Computer Science*, 8(4), 2012. doi:10.2168/LMCS-8(4:6)2012.
- 16 Mariangiola Dezani-Ciancaglini, Ugo de’Liguoro, and Nobuko Yoshida. On progress for structured communications. In *TGC 2007*, volume 4912 of *LNCS*, pages 257–275. Springer, 2007. doi:10.1007/978-3-540-78663-4\_18.
- 17 Julia Gabet and Nobuko Yoshida. Static Race Detection and Mutex Safety and Liveness for Go Programs. *LIPICs*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. ECOOP’20.
- 18 Simon Gay and Antonio Ravara, editors. *Behavioural Types: from Theory to Tools*. River Publishers series in automation, control and robotics. River Publishers, June 2017.

- 19 Issac Gouy. Computer language benchmark game. <http://benchmarksgame.alioth.debian.org>, 2017.
- 20 Paul Harvey, Simon Fowler, Ornela Dardha, and Simon J. Gay. Multiparty Session Types for Safe Runtime Adaptation in an Actor Language. In *ECOOP 2021*, volume 194 of *LIPICs*, pages 10:1–10:30. Schloss Dagstuhl, 2021. doi:10.4230/LIPICs.ECOOP.2021.10.
- 21 C. A. R. Hoare. *Communicating Sequential Processes*. Prentice-Hall, Inc., 1985.
- 22 Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. In *Proc. of 35th Symp. on Princ. of Prog. Lang.*, POPL '08, pages 273–284. ACM, 2008. doi:10.1145/1328438.1328472.
- 23 Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. *J. ACM*, 63(1):9:1–9:67, 2016. doi:10.1145/2827695.
- 24 Raymond Hu. Distributed programming using Java APIs generated from Session Types. *Behavioural Types: from Theory to Tools*, pages 287–308, 2017.
- 25 Raymond Hu, Dimitrios Kouzapas, Olivier Pernet, Nobuko Yoshida, and Kohei Honda. Type-Safe Eventful Sessions in Java. In *ECOOP 2010*, volume 6183 of *LNCS*, pages 329–353. Springer, 2010.
- 26 Raymond Hu and Nobuko Yoshida. Hybrid session verification through endpoint API generation. In *FASE 2016*, volume 9633 of *LNCS*, pages 401–418. Springer, 2016. doi:10.1007/978-3-662-49665-7\_24.
- 27 Raymond Hu and Nobuko Yoshida. Explicit connection actions in multiparty session types. In *FASE 2017*, volume 10202 of *LNCS*, pages 116–133. Springer, 2017. doi:10.1007/978-3-662-54494-5\_7.
- 28 Hans Hüttel, Ivan Lanese, Vasco T. Vasconcelos, Luís Caires, Marco Carbone, Pierre-Malo Deniérou, Dimitris Mostrous, Luca Padovani, António Ravara, Emilio Tuosto, Hugo Torres Vieira, and Gianluigi Zavattaro. Foundations of session types and behavioural contracts. *ACM Comput. Surv.*, 49(1):3:1–3:36, 2016. doi:10.1145/2873052.
- 29 Keigo Imai, Rumyana Neykova, Nobuko Yoshida, and Shoji Yuen. Multiparty session programming with global protocol combinators. In *ECOOP 2020*, volume 166 of *LIPICs*, pages 9:1–9:30. Schloss Dagstuhl, 2020. doi:10.4230/LIPICs.ECOOP.2020.9.
- 30 Jules Jacobs, Stephanie Balzer, and Robbert Krebbers. Connectivity graphs: a method for proving deadlock freedom based on separation logic. *Proc. ACM Program. Lang.*, 6(POPL):1–33, 2022. doi:10.1145/3498662.
- 31 Jules Jacobs, Stephanie Balzer, and Robbert Krebbers. Multiparty GV: Functional Multiparty Session Types with Certified Deadlock Freedom. *Proc. ACM Program. Lang.*, 6(ICFP), aug 2022. doi:10.1145/3547638.
- 32 Sung-Shik Jongmans and Petra van den Bos. A Predicate Transformer for Choreographies - Computing Preconditions in Choreographic Programming. In *ESOP 2022*, volume 13240 of *LNCS*, pages 520–547. Springer, 2022. doi:10.1007/978-3-030-99336-8\_19.
- 33 Kubernetes: Production-grade container orchestration. <https://kubernetes.io/>, June 2017.
- 34 Dimitrios Kouzapas, Ornela Dardha, Roly Perera, and Simon J. Gay. Typechecking protocols with Mungo and StMungo. In *PPDP*, pages 146–159, 2016. doi:10.1145/2967973.2968595.
- 35 Dimitrios Kouzapas, Nobuko Yoshida, Raymond Hu, and Kohei Honda. On Asynchronous Eventful Session Semantics. *MSCS*, 29:1–62, 2015.
- 36 Nicolas Lagailardie, Rumyana Neykova, and Nobuko Yoshida. Implementing Multiparty Session Types in Rust. In *Coordination Models and Languages*, volume 12134, pages 127–136. Springer, 2020. doi:10.1007/978-3-030-50029-08.
- 37 Julien Lange, Nicholas Ng, Bernardo Toninho, and Nobuko Yoshida. Fencing off Go: Liveness and safety for channel-based programming. In Giuseppe Castagna and Andrew D. Gordon, editors, *POPL 2017*, pages 748–761. ACM, 2017.
- 38 Julien Lange, Nicholas Ng, Bernardo Toninho, and Nobuko Yoshida. A Static Verification Framework for Message Passing in Go using Behavioural Types. In *40th International Conference on Software Engineering*, pages 1137–1148. ACM, 2018.

- 39 Julien Lange, Emilio Tuosto, and Nobuko Yoshida. From communicating machines to graphical choreographies. In *POPL 2015*, pages 221–232. ACM, 2015. doi:10.1145/2676726.2676964.
- 40 Ziheng Liu, Shuofei Zhu, Boqin Qin, Hao Chen, and Linhai Song. Automatically detecting and fixing concurrency bugs in Go software systems. In *ASPLOS '21*, April 2021.
- 41 Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, i. *Information and Computation*, 100(1):1 – 40, 1992. doi:[https://doi.org/10.1016/0890-5401\(92\)90008-4](https://doi.org/10.1016/0890-5401(92)90008-4).
- 42 Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, ii. *Information and Computation*, 100(1):41 – 77, 1992. doi:[https://doi.org/10.1016/0890-5401\(92\)90009-5](https://doi.org/10.1016/0890-5401(92)90009-5).
- 43 Anson Miu, Francisco Ferreira, Nobuko Yoshida, and Fangyi Zhou. Communication-safe web programming in TypeScript with Routed Multiparty Session Types. In *CC 2021*, 2021.
- 44 Rumyana Neykova, Raymond Hu, Nobuko Yoshida, and Fahd Abdeljallal. A session type provider: Compile-time API generation of distributed protocols with refinements in F#. In *CC 2018*, page 128–138. ACM, 2018. doi:10.1145/3178372.3179495.
- 45 Rumyana Neykova and Nobuko Yoshida. Featherweight Scribble. *Models, Languages, and Tools for Concurrent and Distributed Programming*, 11665:236–259, 2019. doi:10.1007/978-3-030-21485-214.
- 46 Nicholas Ng, José Gabriel de Figueiredo Coutinho, and Nobuko Yoshida. Protocols by default - safe MPI code generation based on session types. In *CC 2015*, volume 9031 of *LNCS*, pages 212–232. Springer, 2015. doi:10.1007/978-3-662-46663-6\_11.
- 47 Nicholas Ng and Nobuko Yoshida. Pabble: parameterised scribble. *Service Oriented Computing and Applications*, 9(3-4):269–284, 2015. doi:10.1007/s11761-014-0172-8.
- 48 Nicholas Ng and Nobuko Yoshida. Static deadlock detection for concurrent Go by global session graph synthesis. In *CC 2016*, pages 174–184. ACM, 2016. doi:10.1145/2892208.2892232.
- 49 The nuScr authors. nuscr homepage. <https://nuscr.github.io/nuscr/>, 2019.
- 50 Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall series in artificial intelligence. Pearson Education, 2016.
- 51 Alceste Scalas, Ornela Dardha, Raymond Hu, and Nobuko Yoshida. A linear decomposition of multiparty sessions for safe distributed programming. In *ECOOP 2017*, volume 74 of *LIPICs*, pages 24:1–24:31. Schloss Dagstuhl, 2017. doi:10.4230/LIPICs.ECOOP.2017.24.
- 52 Scribble Authors. Scribble: Describing multiparty protocols. <http://www.scribble.org/>, 2015. Accessed in Nov. 2020.
- 53 Kai Stadtmüller, Martin Sulzmann, and Peter Thiemann. Static trace-based deadlock analysis for synchronous mini-go. In *APLAS 2016*, volume 10017 of *LNCS*, pages 116–136, 2016. doi:10.1007/978-3-319-47958-3\_7.
- 54 I. Stoica, R. Morris, D. Liben-Nowell, D.R. Karger, M.F. Kaashoek, F. Dabek, and H. Balakrishnan. Chord: a Scalable Peer-to-peer Lookup Protocol for Internet Applications. *IEEE/ACM Transactions on Networking*, 11(1):17–32, 2003. doi:10.1109/TNET.2002.808407.
- 55 Martin Sulzmann and Peter Thiemann. Forkable regular expressions. In *LATA 2016*, volume 9618 of *LNCS*, pages 194–206. Springer, 2016. doi:10.1007/978-3-319-30000-9\_15.
- 56 Tengfei Tu, Xiaoyu Liu, Linhai Song, and Yiyang Zhang. Understanding real-world concurrency bugs in go. In *ASPLOS 2019*, page 865–878. ACM, 2019. doi:10.1145/3297858.3304069.
- 57 Malte Viering, Raymond Hu, Patrick Eugster, and Lukasz Ziarek. A multiparty session typing discipline for fault-tolerant event-driven distributed programming. *PACMPL*, 5(OOPSLA):1–30, 2021. doi:10.1145/3485501.
- 58 Philip Wadler. Propositions as sessions. In *Proc. of the 17th International Conference on Functional Programming*, volume 47, pages 273–286, 2012.
- 59 Felix A. Wolf, Linard Arquint, Martin Clochard, Wytse Oortwijn, João C. Pereira, and Peter Müller. Gobra: Modular specification and verification of Go programs. In *CAV*, pages 367–379. Springer, 2021.

- 60 Nobuko Yoshida and Lorenzo Gheri. A very gentle introduction to Multiparty Session Types. In *16th International Conference on Distributed Computing and Internet Technology*, volume 11969 of *LNCS*, pages 73–93. Springer, 2020.
- 61 T. Yuan, G. Li, J. Lu, C. Liu, L. Li, and J. Xue. GoBench: A benchmark suite of real-world Go concurrency bugs. In *CGO 2021*. ACM/IEEE, 2021.
- 62 Fangyi Zhou, Francisco Ferreira, Raymond Hu, Rumyana Neykova, and Nobuko Yoshida. Statically Verified Refinements for Multiparty Protocols. volume 4. ACM, November 2020. doi:10.1145/3428216.