

# Motion Session Types for Robotic Interactions

**Rupak Majumdar**

MPI-SWS, Saarbrücken, Germany  
rupak@mpi-sws.org

**Marcus Pirron**

MPI-SWS, Saarbrücken, Germany  
mpirron@mpi-sws.org

**Nobuko Yoshida** 

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

**Damien Zufferey** 

MPI-SWS, Saarbrücken, Germany  
zufferey@mpi-sws.org

---

## Abstract

Robotics applications involve programming concurrent components synchronising through messages while simultaneously executing *motion primitives* that control the state of the physical world. Today, these applications are typically programmed in low-level imperative programming languages which provide little support for abstraction or reasoning.

We present a unifying programming model for concurrent message-passing systems that additionally control the evolution of physical state variables, together with a compositional reasoning framework based on multiparty session types. Our programming model combines *message-passing concurrent processes* with *motion primitives*. Processes represent autonomous components in a robotic assembly, such as a cart or a robotic arm, and they synchronise via discrete messages as well as via motion primitives. Continuous evolution of trajectories under the action of controllers is also modelled by motion primitives, which operate in global, physical time.

We use multiparty session types as specifications to orchestrate discrete message-passing concurrency and continuous flow of trajectories. A global session type specifies the communication protocol among the components with joint motion primitives. A projection from a global type ensures that jointly executed actions at end-points are *communication safe* and *deadlock-free*, i.e., session-typed components do not get stuck. Together, these checks provide a compositional verification methodology for assemblies of robotic components with respect to concurrency invariants such as a progress property of communications as well as dynamic invariants such as absence of collision.

We have implemented our core language and, through initial experiments, have shown how multiparty session types can be used to specify and compositionally verify robotic systems implemented on top of off-the-shelf and custom hardware using standard robotics application libraries.

**2012 ACM Subject Classification** Computer systems organization → Robotics; Software and its engineering → Concurrent programming languages; Theory of computation → Process calculi; Theory of computation → Type theory

**Keywords and phrases** Session Types, Robotics, Concurrent Programming, Motions, Communications, Multiparty Session Types, Deadlock Freedom

**Digital Object Identifier** 10.4230/LIPIcs.ECOOP.2019.28

**Category** Brave New Idea Paper

**Funding** *Rupak Majumdar*: DFG 389792660 TRR 248, ERC Synergy Grant 610150.

*Marcus Pirron*: DFG 389792660 TRR 248, ERC Synergy Grant 610150.

*Nobuko Yoshida*: EPSRC EP/K034413/1, EP/K011715/1, EP/L00058X/1, EP/N027833/1 and EP/N028201/1.

*Damien Zufferey*: DFG 389792660 TRR 248, ERC Synergy Grant 610150.



© Rupak Majumdar, Marcus Pirron, Nobuko Yoshida, and Damien Zufferey;  
licensed under Creative Commons License CC-BY

33rd European Conference on Object-Oriented Programming (ECOOP 2019).

Editor: Alastair F. Donaldson; Article No. 28; pp. 28:1–28:27

Leibniz International Proceedings in Informatics



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

## 1 Introduction

Many cyber-physical systems today involve an interaction among communication-centric components which together control trajectories of physical variables. For example, consider an autonomous robotic system executing in an assembly line. The components in such an example would be robotic manipulators or arms as well as robotic carts onto which one or more arms may be mounted. A global task may involve communication between the carts and the arms – for example, to jointly decide the position of the arms and to jointly plan trajectories – as well as the execution of motion primitives – for example, to follow a trajectory or to grip an object. Today, a programmer developing such an application must manually orchestrate the messaging and the dynamics: errors in either can lead to potentially catastrophic system failures. Typically, programs are written in (untyped) imperative programming language using messaging libraries. Arguments about correctness are informal at best, with no support from the language.

In this paper, we take the first steps towards a uniform programming model for autonomous robotic systems. Our model combines message-based communication with physical dynamics (“motion primitives”) over time. Our starting point is the notion of *multiparty session types* [25, 26, 10], a principled, type-based, discipline to specify and reason about *global communication protocols* in a concurrent system. We enrich a process-based core language for communication with the ability to execute *dynamic motion primitives* over time. Motion primitives encapsulate the actions of dynamic controllers on the physical world and define the continuous evolution of the trajectories of the system. At the same time, we enrich a type system for multiparty session-based communication with motion primitives.

The interaction of communication and dynamics is non-trivial. Since time is global to a physical system, every independently running process must be ready to execute their motion primitives simultaneously. Thus, for example, programs in which one component is blocked waiting for a message while another moves along a trajectory must be ruled out as ill-typed. To keep the complexity of the problem manageable, our semantics keeps, as much as possible, the message exchanges separate from the continuous trajectories. In particular, in our model, message exchanges occur instantaneously and at discrete time steps, à la synchronous reactive programming, while motion primitives execute in global time. System evolution is then organised into rounds; each round consists of a logical time for communication followed by physical time for motion. This assumption is realistic for systems where the speed of the trajectories is comparatively slow compared to the message transmission delay.

Our reasoning principles closely follow the usual type-checking approach of multiparty session types. Specifications are described through *global types*, which constrain both message sequences and motion sequences. Global types are projected to *local types*, which specify the actions in a session from the perspective of a single end-point process. Finally, a verification step checks that each process satisfies its local type. The soundness theorem ensures that in this last case, the composition of the processes satisfy a protocol compliance.

Our type system ensures communication safety and deadlock-freedom for messages, ensuring, for example, that communication is not stuck or time cannot progress. In addition, we verify safety properties of physical trajectories such as non-collision by constraint-based verification of simultaneously executed motion primitives specified in the global type.

Existing session type formalisms such as [9] fall short to model a combination of individual interactions and global synchronisations by motions. To demonstrate our initial step and to observe an effect of new primitives specific to robotics interactions, we start from the simplest multiparty session type system in [15, 18]. The programming model and type

system introduced in this paper provides the foundations for PGCD programs, a practical programming system to develop concurrent robotics applications [4]. We have used our calculus and type system to verify correctness properties of (abstract versions of) multi-robot co-ordination programs written in PGCD, which then execute on real robotics hardware. Our evaluation shows that multiparty session types and choreographies for multi-robot co-ordination and manipulation can lead to statically verified implementations that run on off-the-shelf and custom robotics hardware platforms.

**Outline.** We first give a gentle introduction to motion session types to those who are interested in concurrent robotics programming, but not familiar with session types. Section 3 discusses a core abstract calculus of processes where motions are abstracted by just the passage of time; Section 4 defines a typing system with motion primitives; Section 5 extends our theory to deal with continuous trajectories; Section 6 discusses our implementation; Section 7 gives related work and Section 8 concludes.

## 2 A Gentle Introduction to Motion Session Types

The aim of this section is to give a gentle introduction of motion session types for readers who are interested in robotics programming but who are not familiar with session types nor process calculi.

A key difficulty in robotics programming is that the programmer has to reason about concurrent processes communicating through messages as well as about dynamics evolving in time. The idea of motion session types is to provide a typing framework to only allow programs that follow structured sequences of interactions and motion. A *session* will be a natural unit of structured communication and motion. *Motion session types* abstract the structure of a session. and provide a syntax-driven approach to restricting programs to a well-behaved subclass – for this subclass, one can check processes compositionally and derive properties of the composition.

Motion session types extend session types, introduced in a series of papers during the 1990s [23, 43, 24], in the context of pure concurrent programming. Session types have since been studied in many contexts over the last decade – see the surveys of the field [27, 17].

We begin by an overview of the key technical ideas of multiparty session types. Then we introduce motion primitives to multiparty session types for specifying actions over time. Finally, we refine the motion primitives to physical motion executed by the robots.

### 2.1 Communication: Multiparty Session Types

We begin with a review of multiparty session types, a methodology to enable compositional reasoning about communication.

As a simple example, consider a scenario in which a cart and arm assembly has to fetch objects. We associate a process with each physical component; thus, we model the scenario using a *cart* (Cart) and an *arm* (Arm) attached to the cart. The task involves synchronisation between the cart and the arm as well as co-ordinated motion. Synchronization is obtained through the exchange of messages. We defer the discussion on motion to Section 2.2.

Specifically, the protocol works as follows.

1. The cart sends the arm a fold command *fold*. On receiving the command, the arm folds itself. When the arm is completely folded, it sends back a message *ok* to the cart. On receipt of this message, the cart moves.

## 28:4 Motion Session Types for Robotic Interactions

2. When the cart reaches the object, it stops and sends a *grab* message to the arm to grab the object. While the cart waits, the arm executes the grabbing operation, followed by a folding operation. Then the arm sends a message *ok* to the cart. This sequence may need to be repeated.
3. When all tasks are finished, the cart sends a message *done* to the arm, and the protocol terminates.

The multiparty session types methodology is as follows. First, define a *global type* that gives a shared contract of the allowed pattern of message exchanges in the system. Second, *project* the global type to each end-point participant to get a *local type*: an obligation on the message sends and receipts for each process that together ensure that the pattern of messages are allowed by the global type. Finally, check that the implementation of each process conforms to its local type.

In our protocol, from a global perspective, we expect to see the following pattern of message exchanges, encoded as a *global type* for the communication:

$$\mu t. \text{Cart} \rightarrow \text{Arm} : \{ \text{fold. Arm} \rightarrow \text{Cart} : \text{ok. Cart} \rightarrow \text{Arm} : \text{grab. Arm} \rightarrow \text{Cart} : \text{ok.t}, \text{done. end} \} \quad (1)$$

The type describes the global pattern of communication between *Cart* and *Arm* using message exchanges, sequencing, choice, and repetition. The basic pattern  $\text{Cart} \rightarrow \text{Arm} : m$  indicates a message  $m$  sent from the *Cart* to the *Arm*. The communication starts with the cart sending either a *fold* or a *done* command to the arm. In case of *done*, the protocol ends (type **end**); otherwise, the communication continues with the sequence *ok. grab. ok* followed by a repetition of the entire pattern. The operator “.” denotes sequencing, and the type  $\mu t. T$  denotes recursion of  $T$ .

The global type states what are the valid message sequences allowed in the system. When we implement *Cart* and *Arm* separately, we would like to check that their composition conforms to the global type. We can perform this check compositionally as follows.

Since there are only two participants, projecting to each participant is simple. From the perspective of the *Cart*, the communication can be described by the type:

$$\mu t. ( ( !\text{fold. ?ok. !grab. ?ok.t} ) \oplus ( !\text{done. end} ) ) \quad (2)$$

where  $!m$  denotes a message  $m$  sent (to the *Arm*) and  $?m$  denotes a message  $m$  received from the *Arm*. and  $\oplus$  denotes an (internal) choice. Thus, the type states that *Cart* repeats actions  $!\text{fold. ?ok. !grab. ?ok}$  until at some point it sends *done* and exits.

Dually, from the viewpoint of the *Arm*, the same global session is described by the dual type

$$\mu t. ( ( ?\text{fold. !ok. ?grab. !ok.t} ) \& ( ?\text{done. end} ) ) \quad (3)$$

in which  $\&$  means that a choice is offered externally.

We can now individually check that the implementations of the cart and the arm conform to these local types.

The global type seems overkill if there are only two participants; indeed, the global type is uniquely determined given the local type (2) or its dual (3). However, for applications involving *multiple parties*, the global type and its projection to each participant are essential to provide a shared contract among all participants.

For example, consider a simple ring protocol, where the Arm process above is divided into two parts, Lower and Upper. Now, Cart sends a message *fold* to the lower arm Lower, which forwards the message to Upper. After receiving the message, Upper sends an acknowledgement *ok* to Cart. We start by specifying the global type as:

$$\text{Cart} \rightarrow \text{Lower} : \text{fold} . \text{Lower} \rightarrow \text{Upper} : \text{fold} . \text{Upper} \rightarrow \text{Cart} : \text{ok} . \text{end} \quad (4)$$

As before, we want to check each process locally against a local type such that if each process conforms to its local type then the composition satisfies the global type.

The global type in (4) is *projected* into the three endpoint session types:

Cart's endpoint type:  $\text{Lower}! \text{fold} . \text{Upper} ? \text{ok} . \text{end}$

Lower's endpoint type:  $\text{Cart} ? \text{fold} . \text{Upper} ! \text{fold} . \text{end}$

Upper's endpoint type:  $\text{Lower} ? \text{fold} . \text{Cart} ! \text{ok} . \text{end}$

where  $\text{Lower}! \text{fold}$  means “send to Lower a *fold* message,” and  $\text{Upper} ? \text{ok}$  means “receive from Upper an *ok* message.” Then each process is type-checked against its own endpoint type. When the three processes are executed, their interactions automatically follow the stipulated scenario.

If instead of a global type, we only used three separate binary session types to describe the message exchanges between Cart and Lower, between Lower and Upper, and between Upper and Cart, respectively, without using a global type, then we lose essential sequencing information in this interaction scenario. Consequently, we can no longer guarantee deadlock-freedom among these three parties. Since the three separate binary sessions can be interleaved freely, an implementation of the Cart that conforms to  $\text{Upper} ? \text{ok} . \text{Lower} ! \text{fold} . \text{end}$  becomes typable. This causes the situation that each of the three parties blocks indefinitely while waiting for a message to be delivered. Thus, we shall use the power of multiparty session types to ensure correct communication patterns.

## 2.2 Motion: Motion Primitives and Trajectories

So far, we focused on the communication pattern and ignored the physical actions of the robots. Our framework of motion session types extends multiparty session types to also reason about *motion primitives*, which model change of state in the physical world effected by the robots. We add motion in two steps: first we treat motion primitives as abstract actions that have associated durations, and second as dynamic trajectories.

Abstractly, we model motion primitives as actions that take physical time. Accordingly, we extend session types with motion primitive  $\text{dt}(\mathbf{p}_i : a_i)$ , which indicates that the participants  $\mathbf{p}_i$  jointly execute motion primitives  $a_i$  for the same duration of time.

Let us add the motion primitives to the cart and arm example. Recall that on receiving the command *fold*, the arm folds itself; meanwhile, the cart waits. When the arm is completely folded, it sends back a message to the cart, then the cart moves, following a trajectory to the object. This means the time the arm folds and the time the cart is idle (waiting for the arm) should be the same. Similarly, the time cart is moving and the idle time the arm waits for the cart should be synchronised. This explicit synchronisation is represented by the following global type:

$$\begin{aligned} \text{Cart} &\rightarrow \text{Arm} : \text{fold} . \text{dt}(\text{Cart} : \text{idle}, \text{Arm} : \text{fold}) . \\ \text{Arm} &\rightarrow \text{Cart} : \text{ok} . \text{dt}(\text{Cart} : \text{move}, \text{Arm} : \text{idle}) . G \end{aligned}$$

where “ $\text{dt}\langle \text{Cart} : \text{idle}, \text{Arm} : \text{fold} \rangle$ ” specifies the joint motion primitives `idle` executed by the `Cart` and `fold` executed by the `Arm` are synchronised. We extend local types with motion primitives as well. The conformance check ensures that, if each process conforms to its local types, then the composition of the system conforms to the global type – which now includes both message-based synchronization as well as synchronization over time using motion primitives.

Finally, we expand the abstract motion primitives with the underlying dynamic controllers and ensure that the joint execution of motion primitives is possible in the system. This requires refining each motion primitive to its underlying dynamical system and checking that whenever the global type specifies a joint execution of motion primitives, there is in fact a joint trajectory of the system that can be executed.

### 3 Motion Session Calculus

We now introduce the syntax and semantics of a synchronous multiparty motion session calculus. Our starting point is to associate a process with the physical component it controls. This can be either a “complete” robot or parts of a robot (like the cart or arm in the previous section). This makes it possible to model modular robots where parts may be swapped for different tasks. In the following, we simply say “robot” to describe a physical component (which may be a complete robot or part of a larger robot). Our programming model will associate a process with each such robot.

We build our motion session calculus based on a session calculus studied in [15, 18], which simplifies the synchronous multiparty session calculus in [29] by eliminating both shared channels for session initiations and session channels for communications inside sessions.

► **Notation 3.1** (Base sets). *We use the following base sets: values, ranged over by  $v, v', \dots$ ; expressions, ranged over by  $e, e', \dots$ ; expression variables, ranged over by  $x, y, z, \dots$ ; labels, ranged over by  $l, l', \dots$ ; session participants, ranged over by  $p, q, \dots$ ; motion primitives, ranged over by  $a, b, \dots$ ; process variables, ranged over by  $X, Y, \dots$ ; processes, ranged over by  $P, Q, \dots$ ; and multiparty sessions, ranged over by  $M, M', \dots$ .*

#### Motion Primitives

When reasoning about communication and synchronisation, the actual trajectory of the system is not important and only the time taken by a motion is important. Therefore, we first abstract away trajectories by just keeping the name of the motion primitive ( $a, b, \dots$ ) and, for each motion, we assume we know up front how long the action takes. We use the notation  $\text{dt}\langle a \rangle$  to represent that a motion primitive executes and time elapses. Every motion can have a different, a priori known, duration denoted  $\text{duration}(a)$ . We write the tuple  $\text{dt}\langle (p_i : a_i) \rangle$  to denote a group of processes executing their respective motion primitives at the same time. For the sake of simplicity, we sometimes use  $a$  for both single or grouped motions. In Section 5, we look in more details into the trajectories defined by the joint execution of motion primitives.

#### Syntax of Motion Session Calculus

A value  $v$  can be a natural number  $n$ , an integer  $i$ , a Boolean `true` / `false`, or a real number. An expression  $e$  can be a variable, a value, or a term built from expressions by applying (type-correct) computable operators. The processes of the synchronous multiparty session calculus are defined by:



$$\begin{aligned}
P ::= & \mathfrak{p}!\ell\langle e \rangle.P \mid \sum_{i \in I} \mathfrak{p}?\ell_i(x_i).P_i \mid \sum_{i \in I} \mathfrak{p}?\ell_i(x_i).P_i + \mathfrak{dt}\langle a \rangle.P \mid \mathfrak{dt}\langle a \rangle.P \\
& \mid \text{if } e \text{ then } P \text{ else } Q \mid \mu X.P \mid X \mid \mathbf{0}
\end{aligned}$$

The output process  $\mathfrak{p}!\ell\langle e \rangle.Q$  sends the value of expression  $e$  with label  $\ell$  to participant  $\mathfrak{p}$ . The sum of input processes (external choice)  $\sum_{i \in I} \mathfrak{p}?\ell_i(x_i).P_i$  is a process that can accept a value with label  $\ell_i$  from participant  $\mathfrak{p}$  for any  $i \in I$ ;  $\sum_{i \in I} \mathfrak{p}?\ell_i(x_i).P_i + \mathfrak{dt}\langle a \rangle.P$  is an external choice with a *default branch* with a motion action  $\mathfrak{dt}\langle a \rangle.P$  which can always proceed when there is no message to receive. According to the label  $\ell_i$  of the received value, the variable  $x_i$  is instantiated with the value in the continuation process  $P_i$ . We assume that the set  $I$  is always finite and non-empty. The conditional process  $\text{if } e \text{ then } P \text{ else } Q$  represents the internal choice between processes  $P$  and  $Q$ . Which branch of the conditional process will be taken depends on the evaluation of the expression  $e$ . The process  $\mu X.P$  is a recursive process. We assume that the recursive processes are *guarded*. For example,  $\mu X.\mathfrak{p}?\ell(x).X$  is a valid process, while  $\mu X.X$  is not. We often omit  $\mathbf{0}$  from the tail of processes.

We define a *multiparty session* as a parallel composition of pairs (denoted by  $\mathfrak{p} \triangleleft P$ ) of participants and processes:

$$M ::= \mathfrak{p} \triangleleft P \mid M \mid M$$

with the intuition that process  $P$  plays the role of participant  $\mathfrak{p}$ , and can interact with other processes playing other roles in  $M$ . The participants correspond to the physical components in the system and the processes correspond to the code run by that physical component. A multiparty session is *well formed* if all its participants are different. We consider only well-formed multiparty sessions.

## Operational Semantics of Motion Session Calculus

The value  $v$  of expression  $e$  (notation  $e \downarrow v$ ) is computed as expected. We assume that  $e \downarrow v$  is effectively computable and takes logical “zero time.”

We adopt some standard conventions regarding the syntax of processes and sessions. Namely, we will use  $\prod_{i \in I} \mathfrak{p}_i \triangleleft P_i$  as short for  $\mathfrak{p}_1 \triangleleft P_1 \mid \dots \mid \mathfrak{p}_n \triangleleft P_n$ , where  $I = \{1, \dots, n\}$ . We will sometimes use infix notation for external choice process. For example, instead of  $\sum_{i \in \{1,2\}} \mathfrak{p}?\ell_i(x).P_i$ , we will write  $\mathfrak{p}?\ell_1(x).P_1 + \mathfrak{p}?\ell_2(x).P_2$ .

The *computational rules of multiparty sessions* are given in Table 1. They are closed with respect to structural congruence. The structural congruence includes a recursion rule  $\mu X.P \equiv P\{\mu X.P/X\}$ , as well as expected rules for multiparty sessions such as  $P \equiv Q \Rightarrow \mathfrak{p} \triangleleft P \mid M \equiv \mathfrak{p} \triangleleft Q \mid M$ . Other rules are standard from [15, 18]. However, unlike the usual treatment of  $\pi$ -calculi, our structural congruence does not have a rule to simplify inactive processes ( $\mathfrak{p} \triangleleft \mathbf{0}$ ). The reason is that even when a program might be logically terminated, the physical robot continues to exist and may still collide with another robot. Therefore, in our model, all processes need to terminate at the same time, and so we need to keep  $\mathfrak{p} \triangleleft \mathbf{0}$ .

In rule [COMM], the participant  $\mathfrak{q}$  sends the value  $v$  choosing the label  $\ell_j$  to participant  $\mathfrak{p}$ , who offers inputs on all labels  $\ell_i$  with  $i \in I$ . In rules [T-CONDITIONAL] and [F-CONDITIONAL], the participant  $\mathfrak{p}$  chooses to continue as  $P$  if the condition  $e$  evaluates to **true** and as  $Q$  if  $e$  evaluates to **false**. Rule [R-STRUCT] states that the reduction relation is closed with respect to structural congruence. We use  $\longrightarrow^*$  for the reflexive transitive closure of  $\longrightarrow$ .

The motion primitives are handled with [MOTION] and [M-PAR]. Here, we need to label transitions with the time taken by the action and propagate these labels with the parallel composition. This ensures that when (physical) time elapses for one process, it elapses

■ **Table 1** Reduction rules. The communication between an output and an external choice (without the default motion action) is formalised similarly to [COMM].

$$\begin{array}{c}
 \text{[COMM]} \\
 \frac{j \in I \quad e \downarrow v}{\mathfrak{p} \triangleleft \sum_{i \in I} \mathfrak{q} ? \ell_i(x). P_i + \text{dt}\langle a \rangle . P \mid \mathfrak{q} \triangleleft \mathfrak{p} ! \ell_j \langle e \rangle . Q \longrightarrow \mathfrak{p} \triangleleft P_j \{v/x\} \mid \mathfrak{q} \triangleleft Q} \\
 \\
 \begin{array}{cc}
 \text{[DEFAULT]} & \text{[MOTION]} \\
 \mathfrak{p} \triangleleft \sum_{i \in I} \mathfrak{q} ? \ell_i(x). P_i + \text{dt}\langle a \rangle . P \xrightarrow{\text{dt}\langle a \rangle} \mathfrak{p} \triangleleft P & \mathfrak{p} \triangleleft \text{dt}\langle a \rangle . P \xrightarrow{\text{dt}\langle a \rangle} \mathfrak{p} \triangleleft P
 \end{array} \\
 \\
 \begin{array}{cc}
 \text{[T-CONDITIONAL]} & \text{[F-CONDITIONAL]} \\
 \frac{e \downarrow \text{true}}{\mathfrak{p} \triangleleft \text{if } e \text{ then } P \text{ else } Q \longrightarrow \mathfrak{p} \triangleleft P} & \frac{e \downarrow \text{false}}{\mathfrak{p} \triangleleft \text{if } e \text{ then } P \text{ else } Q \longrightarrow \mathfrak{p} \triangleleft Q}
 \end{array} \\
 \\
 \begin{array}{cc}
 \text{[R-PAR]} & \text{[M-PAR]} \\
 \frac{\mathfrak{p} \triangleleft Q \longrightarrow \mathfrak{p} \triangleleft Q'}{\mathfrak{p} \triangleleft Q \mid M \longrightarrow \mathfrak{p} \triangleleft Q' \mid M} & \frac{\mathfrak{p}_i \triangleleft P_i \xrightarrow{\text{dt}\langle a_i \rangle} \mathfrak{p}_i \triangleleft P'_i \quad \forall i, j. \text{duration}(a_i) = \text{duration}(a_j)}{\prod_i \mathfrak{p}_i \triangleleft P_i \xrightarrow{\text{dt}\langle (\mathfrak{p}_i : a_i) \rangle} \prod_i \mathfrak{p}_i \triangleleft P'_i}
 \end{array} \\
 \\
 \begin{array}{cc}
 \text{[R-STRUCT]} & \text{[M-STRUCT]} \\
 \frac{M'_1 \equiv M_1 \quad M_1 \longrightarrow M_2 \quad M_2 \equiv M'_2}{M'_1 \longrightarrow M'_2} & \frac{M'_1 \equiv M_1 \quad M_1 \xrightarrow{\text{dt}\langle a \rangle} M_2 \quad M_2 \equiv M'_2}{M'_1 \xrightarrow{\text{dt}\langle a \rangle} M'_2}
 \end{array}
 \end{array}$$

equally for all processes; every process has to spend the same amount of time. This style of synchronisation is reminiscent of broadcast calculi [39]. Instead of broadcast messages, we broadcast *time*.

In order to state that communications can always make progress, we formalise when a multiparty session contains communications or motion actions that will never be executed.

► **Definition 3.2.** *A multiparty motion session  $M$  is stuck if  $M \not\equiv \prod_{i \in I} \mathfrak{p}_i \triangleleft \mathbf{0}$  and there is no multiparty session  $M'$  such that  $M \longrightarrow M'$ . A multiparty session  $M$  gets stuck, notation  $\text{stuck}(M)$ , if it reduces to a stuck motion multiparty session.*

We finish this section with some examples of multi-party sessions.

► **Example 3.3 (A Simple Fetch Scenario).** Recall the scenario from Section 2 in which a cart and arm assembly has to fetch an object. There are two processes: a cart and an arm; the arm is attached to the cart. The task involves synchronization between the cart and the arm. Specifically, the protocol works as follows. Initially, the cart sends the arm a command to fold. On receiving the command, the arm folds itself. Meanwhile, the cart waits. When the arm is completely folded, it sends back a message to the cart. On receipt of this message, the cart moves, following a trajectory to the object. When it reaches the object, it stops and sends a message back to the arm to grab the object. While the cart waits, the arm executes the grabbing operation, followed by a folding operation. When the arm is done, it again synchronises with the cart. At this point, the cart moves back to its original position. (We simplify the example from Section 2 so that the sequence is not repeated.)



<pre> Cart&lt;   Arm!fold().   wait (dt&lt;idle&gt;){     Arm?ok().     dt&lt;move&gt;.     Arm!grab.     wait (dt&lt;idle&gt;){       Arm?ok().       dt&lt;move&gt;.       Arm!done().0     }   } </pre>	<pre> Arm&lt;   μX.wait (dt&lt;idle&gt;){     Cart?fold().dt&lt;fold&gt;.Cart!ok().X   }   + Cart?grab().   dt&lt;grip&gt;.   Cart!ok().X   + Cart?done().0 } </pre>
--	--

■ **Figure 1** A cart and arm example.

Figure 1 shows how the cart and arm processes can be encoded in our core language. We introduce some syntactic sugar for readability. We write  $\text{wait (dt}\langle a \rangle) \{ \sum_{i \in I} \mathbf{p}^? \ell_i(x_i).P_i \}$  as shorthand for the process  $\mu X. \sum_{i \in I} \mathbf{p}^? \ell_i(x_i).P_i + \text{dt}\langle a \rangle.X$ , which keeps running the default motion  $a$  until it receives a message.

The motion primitive `idle` keeps the cart or the arm stationary. The primitive `move` moves the cart, the primitives `grip` and `fold` respectively move the arm to grab an object or to fold the arm. At this point, we focus on the communication pattern and therefore abstract away the actual trajectories traced by the motion primitives. We come back to the trajectories in Section 5.

Finally, the multiparty session is the parallel composition of the participants `Cart` and `Arm` with the corresponding processes.

The processes in our calculus closely follow the syntax of PGCD programs [4]. In Figure 2, we show a side by side comparison of a PGCD program and the corresponding process expressed in the motion session calculus.

► **Example 3.4** (Multi-party Co-ordination: Handover). We describe a more complex *handover* example in which a cart and arm assembly transfers an object to a second cart, called the carrier. The process for the arm is identical to Figure 1, but the cart now co-ordinates with the carrier as well. Figure 3 shows all the processes. Note that the cart now synchronises both with the arm and with the carrier.

The protocol is as follows. As before, the cart moves to a target position, having ensured that the arm is folded, and then waits for the carrier to be ready. When the carrier is ready, the arm is instructed to grab an object on the carrier. Once the object is grabbed, the arm synchronises with the cart, which then informs the carrier that the handover is complete. The cart and the carrier move back to their locations and the protocol is complete. The multiparty session is the parallel composition of the participants `Cart`, `Arm`, and `Carrier`, with the corresponding processes.

PGCD: pseudo code for the Arm	Arm $\triangleleft$
1 <b>while</b> <i>true</i> <b>do</b>	$\mu X.$ wait (dt(idle)){
2   <b>receive</b> (idle)	Cart? <i>fold</i> ().
3       <i>fold</i> $\Rightarrow$	dt( <i>fold</i> ).
4           <b>fold</b> ();	Cart! <i>ok</i> () $\cdot X$
5           <b>send</b> (Cart, <i>ok</i> )	+ Cart? <i>grab</i> ().
6           <i>grab</i> $\Rightarrow$	dt( <i>grip</i> ).
7               <b>grip</b> ();	Cart! <i>ok</i> () $\cdot X$
8               <b>send</b> (Cart, <i>ok</i> )	+ Cart? <i>done</i> ().
9               <i>done</i> $\Rightarrow$	<b>0</b>
10               <b>break</b>	}
	}

■ **Figure 2** Comparison of a PGCD code and the corresponding motion session calculus process.

## 4 Multiparty Motion Session Types

This section introduces motion session types for the calculus presented in Section 3. The formulation is based on [29, 30, 14], with adaptations to account for our motion calculus.

### 4.1 Motion Session Types and Projections

Global types act as specifications for the message exchanges among robotic components.

► **Definition 4.1** (Sorts and global motion session types). Sorts, ranged over by  $S$ , are used to define base types:

$$S ::= \text{unit} \mid \text{nat} \mid \text{int} \mid \text{bool} \mid \text{real}$$

Global types, ranged over by  $G$ , are terms generated by the following grammar:

$$G ::= \text{dt}(\langle p_i : a_i \rangle).G \mid p \rightarrow q : \{\ell_i(S_i).G_i\}_{i \in I} \mid \mathbf{t} \mid \mu \mathbf{t}.G \mid \mathbf{end}$$

We require that  $p \neq q$ ,  $I \neq \emptyset$ ,  $\ell_i \neq \ell_j$ , and  $\text{duration}(a_i) = \text{duration}(a_j)$  whenever  $i \neq j$ , for all  $i, j \in I$ . We postulate that recursion is guarded and recursive types with the same regular tree are considered equal [37, Chapter 20, Section 2].

In Definition 4.1, the type  $\text{dt}(\langle p_i : a_i \rangle).G$  is a *motion global type* which explicitly declares a *synchronisation* by a motion action among all the participants  $p_i$ . The rest is the standard definition of global types in multiparty session types [29, 30, 14]. The *branching* type  $p \rightarrow q : \{\ell_i(S_i).G_i\}_{i \in I}$  formalises a protocol where participant  $p$  must send to  $q$  one message with label  $\ell_i$  and a value of type  $S_i$  as payload, for some  $i \in I$ ; then, depending on which  $\ell_i$  was sent by  $p$ , the protocol continues as  $G_i$ . Value types are restricted to sorts. The type  $\mathbf{end}$  represents a terminated protocol. A recursive protocol is modelled as  $\mu \mathbf{t}.G$ , where recursion variable  $\mathbf{t}$  is bound and guarded in  $G$ , e.g.,  $\mu \mathbf{t}.\mathbf{t}$  is not a valid type. The notation  $\text{pt}\{G\}$  denotes a set of participants of a global type  $G$ .

```

Cart⊢
  Arm!fold().
  wait (dt⟨idle⟩){
    Arm?ok().Carrier!ok().
    dt⟨move⟩.
    wait (dt⟨idle⟩){
      Carrier?ok().Arm!grab().
      wait (dt⟨idle⟩){
        Arm?ok().Carrier!ok().
        dt⟨move⟩.
        Arm!done().Carrier!done().0
      }
    }
  }

Carrier⊢
  wait (dt⟨idle⟩){
    Cart?ok().dt⟨move⟩.
    Cart!ok().
    wait (dt⟨idle⟩){
      Cart?ok().
      dt⟨move⟩.
      wait (dt⟨idle⟩){Cart?done().0}
    }
  }

Arm⊢
  μX.wait (dt⟨idle⟩){
    Cart?fold().dt⟨fold⟩.Cart!ok().X
  + Cart?grab().dt⟨grip⟩.Cart!ok().X
  + Cart?done().0
  }

```

■ **Figure 3** A multi-party handover example.

► **Example 4.2** (Global session types). The global session type for the fetch example (Example 3.3) is:

```

Cart → Arm : fold(unit).dt⟨Cart : idle, Arm : fold⟩.
Arm → Cart : ok(unit).dt⟨Cart : move, Arm : idle⟩.
Cart → Arm : grab(unit).dt⟨Cart : idle, Arm : grip⟩.
Arm → Cart : ok(unit).dt⟨Cart : move, Arm : idle⟩.
Cart → Arm : done(unit).end

```

and the global session type for the handover example (Example 3.4) is:

```

Cart → Arm : fold(unit).dt⟨Cart : idle, Carrier : idle, Arm : fold⟩.
Arm → Cart : ok(unit).Cart → Carrier : ok(unit).
dt⟨Cart : move, Carrier : move, Arm : idle⟩.
Carrier → Cart : ok(unit).Cart → Arm : grab(unit).
dt⟨Cart : idle, Carrier : idle, Arm : grip⟩.
Arm → Cart : ok(unit).Cart → Carrier : ok(unit).
dt⟨Cart : move, Carrier : move, Arm : idle⟩.
Cart → Arm : done(unit).Cart → Carrier : done(unit).end

```

## 28:12 Motion Session Types for Robotic Interactions

A (local) motion session type describes the behaviour of a single participant in a multiparty motion session.

► **Definition 4.3** (Local motion session types). *The grammar of local types, ranged over by  $T$ , is:*

$$T ::= \text{dt}\langle a \rangle.T \mid \&\{\mathfrak{p}?\ell_i(S_i).T_i\}_{i \in I} \mid \&\{\mathfrak{p}!\ell_i(S_i).T_i\}_{i \in I} \& \text{dt}\langle a \rangle.T \mid \oplus\{\mathfrak{q}!\ell_i(S_i).T_i\}_{i \in I} \\ \mid \mathfrak{t} \mid \mu\mathfrak{t}.T \mid \text{end}$$

We require that  $\ell_i \neq \ell_j$  whenever  $i \neq j$ , for all  $i, j \in I$ . We postulate that recursion is always guarded. Unless otherwise noted, session types are closed.

Labels in a type need to be pairwise different, e.g.,  $\mathfrak{p}!\ell(\text{int}).\text{end}\&\mathfrak{p}!\ell(\text{nat}).\text{end}$  is not a type. The *motion local type*  $\text{dt}\langle a \rangle.T$  represents a motion action followed by the type  $T$ ; the *external choice* or *branching type*  $\&\{\mathfrak{p}?\ell_i(S_i).T_i\}_{i \in I}$  requires to wait to receive a value of sort  $S_i$  (for some  $i \in I$ ) from the participant  $\mathfrak{p}$ , via a message with label  $\ell_i$ ; if the received message has label  $\ell_i$ , the protocol will continue as prescribed by  $T_i$ . The *motion branching choice* is equipped with a default motion type  $\text{dt}\langle a \rangle.T$ . The *internal choice* or *selection type*  $\oplus\{\mathfrak{q}!\ell_i(S_i).T_i\}_{i \in I}$  says that the participant implementing the type must choose a labelled message to send to  $\mathfrak{q}$ ; if the participant chooses the message  $\ell_i$ , for some  $i \in I$ , it must include in the message to  $\mathfrak{q}$  a payload value of sort  $S_i$ , and continue as prescribed by  $T_i$ . Recursion is modelled by the session type  $\mu\mathfrak{t}.T$ . The session type  $\text{end}$  says that no further communication is possible and the protocol is completed. We adopt the following conventions: we do not write branch/selection symbols in case of a singleton choice, we do not write unnecessary parentheses, and we often omit trailing  $\text{ends}$ . The notation  $\mathfrak{pt}\{T\}$  denotes a set of participants of a session type  $T$ .

In Definition 4.4 below, we define the *global type projection* as a relation  $G \upharpoonright_r T$  between global and local types. Our definition extends the one originally proposed by [25, 26], along the lines of [12] and [13] with motion types: i.e., it uses a *merging operator*  $\sqcap$  to combine multiple session types into a single type.

► **Definition 4.4.** *The projection of a global type onto a participant  $r$  is the largest relation  $\upharpoonright_r$  between global and session types such that, whenever  $G \upharpoonright_r T$ :*

- $G = \text{end}$  implies  $T = \text{end}$ ; [PROJ-END]
- $G = \text{dt}\langle (p_i : a_i) \rangle.G'$  implies  $T = \text{dt}\langle a_j \rangle.T'$  with  $r = p_j$  and  $G' \upharpoonright_r T'$ ; [PROJ-MOTION]
- $G = \mathfrak{p} \rightarrow r : \{\ell_i(S_i).G_i\}_{i \in I}$  implies  $T = \&\{\mathfrak{p}?\ell_i(S_i).T_i\}_{i \in I}$  with  $G_i \upharpoonright_r T_i$ ; [PROJ-IN]
- $G = r \rightarrow \mathfrak{q} : \{\ell_i(S_i).G_i\}_{i \in I}$  implies  $T = \oplus\{\mathfrak{q}!\ell_i(S_i).T_i\}_{i \in I}$  and  $G_i \upharpoonright_r T_i, \forall i \in I$ ; [PROJ-OUT]
- $G = \mathfrak{p} \rightarrow \mathfrak{q} : \{\ell_i(S_i).G_i\}_{i \in I}$  and  $r \notin \{\mathfrak{p}, \mathfrak{q}\}$  implies that there are  $T_i, i \in I$  s.t. [PROJ-CONT]  
 $T = \sqcap_{i \in I} T_i$ , and  $G_i \upharpoonright_r T_i$ , for every  $i \in I$ .
- $G = \mu\mathfrak{t}.G$  implies  $T = \mu\mathfrak{t}.T'$  with  $G \upharpoonright_r T'$  if  $r$  occurs in  $G$ , otherwise  $T = \text{end}$ . [PROJ-REC]

Above,  $\sqcap$  is the merging operator, that is a partial operation over session types defined as:

$$T_1 \sqcap T_2 = \begin{cases} T_1 & \text{if } T_1 = T_2 & \text{[MRG-ID]} \\ T_3 & \text{if } \exists I, J : \begin{cases} T_1 = \&\{\mathbf{p}'?l_i(S_i).T_i\}_{i \in I} & \text{and} \\ T_2 = \&\{\mathbf{p}'?l_j(S_j).T_j\}_{j \in J} & \text{and} \\ T_3 = \&\{\mathbf{p}'?l_k(S_k).T_k\}_{k \in I \cup J} \end{cases} & \text{[MRG-BRA1]} \\ T_3 & \text{if } \exists I, J : \begin{cases} T_1 = \&\{\mathbf{p}'?l_i(S_i).T_i\}_{i \in I} \& \text{dt}\langle a \rangle.T' & \text{and} \\ T_2 = \&\{\mathbf{p}'?l_j(S_j).T_j\}_{j \in J} \& \text{dt}\langle a \rangle.T' & \text{and} \\ T_3 = \&\{\mathbf{p}'?l_k(S_k).T_k\}_{k \in I \cup J} \& \text{dt}\langle a \rangle.T' \end{cases} & \text{[MRG-BRA2]} \\ T_3 & \text{if } \exists I, J : \begin{cases} T_1 = \&\{\mathbf{p}'?l_i(S_i).T_i\}_{i \in I} & \text{and} \\ T_2 = \&\{\mathbf{p}'?l_j(S_j).T_j\}_{j \in J} \& \text{dt}\langle a \rangle.T' & \text{and} \\ T_3 = \&\{\mathbf{p}'?l_k(S_k).T_k\}_{k \in I \cup J} \& \text{dt}\langle a \rangle.T' \end{cases} & \text{[MRG-BRA3]} \\ T_3 & \text{if } \exists I, J : \begin{cases} T_1 = \text{dt}\langle a \rangle.T' & \text{and} \\ T_2 = \&\{\mathbf{p}'?l_i(S_i).T_i\}_{i \in I} \& \text{dt}\langle a \rangle.T' & \text{and} \\ T_3 = \&\{\mathbf{p}'?l_i(S_i).T_i\}_{i \in I} \& \text{dt}\langle a \rangle.T' \end{cases} & \text{[MRG-BRA4]} \\ T_3 & \text{if } \exists I, J : \begin{cases} T_1 = \text{dt}\langle a \rangle.T' & \text{and} \\ T_2 = \&\{\mathbf{p}'?l_j(S_j).T_j\}_{j \in I} & \text{and} \\ T_3 = \&\{\mathbf{p}'?l_i(S_i).T_i\}_{i \in I} \& \text{dt}\langle a \rangle.T' \end{cases} & \text{[MRG-BRA5]} \\ T_2 \sqcap T_1 & \text{if } T_2 \sqcap T_1 \text{ is defined,} \\ \text{undefined} & \text{otherwise.} \end{cases}$$

We omit the cases for recursions and selections (defined as [42, S 3]).

Note that our definition is slightly simplified w.r.t. the one of [12] and [13]. Instead of this mergeability operator, one might use more general approach from [42]. This definition is sufficient for our purposes (i.e., to demonstrate an application of session types to robotics communications).

► **Example 4.5.** The projection of the global session type for the fetch example on the cart gives the following local session type:

```

Arm!fold<unit>.dt<idle>.Arm?ok<unit>.dt<move>.Arm!grab<unit>.
dt<idle>.Arm?ok<unit>.dt<move>.Arm!done<unit>.end

```

The local motion session type for the arm is:

```

Cart?fold<unit>.dt<fold>.Cart!ok<unit>.dt<idle>.Cart?grab<unit>.
dt<grip>.Cart!ok<unit>.dt<idle>.Cart!done<unit>.end

```

**On Progress of Time.** Our model assumes that the computation and message transmission time is much faster than the dynamics of the system and, therefore, the messages can be seen as instantaneous. This assumption depends on parameters of the system, like the speed of the network and the dynamics of the physical system, and also on the program being executed. While we cannot directly change the physical system, we can at least check the program is well behaved w.r.t. to time.

If a program can send an unbounded number of messages without executing a motion then this assumption, obviously, does not hold. From the perspective of using the motion calculus to verify a system, this may lead to situation where an unsafe program is deemed safe because time does not progress. For instance, a robot driving straight into a wall could “avoid” crashing into the wall by sending messages in a loop and, therefore, stopping the progress of time.

This problem is not unique to our system but a more general problem in defining the semantics of hybrid systems [20, 21]. In general, one needs to assume that time always *diverges* for infinite executions. In this work, we take a pragmatic solution and simply *disallow*

*0-time recursion.* When recursion is used, all the paths between a  $\mu\mathbf{t}$  and the corresponding  $\mathbf{t}$  must contain at least one motion primitive. This is a simple check which can be done at the syntactic level of global types and it is a sufficient condition for forcing the progress of time.

## 4.2 Motion Session Typing

We now introduce a type system for the multiparty session calculus presented in Section 3. We distinguish three kinds of typing judgments:

$$\Gamma \vdash e : S \qquad \Gamma \vdash P : T \qquad \vdash M : G$$

where  $\Gamma$  is the *typing environment* defined as:  $\Gamma ::= \emptyset \mid \Gamma, x : S \mid \Gamma, X : T$ , i.e., a mapping that associates expression variables with sorts, and process variables with session types.

We use the subtyping relation  $\leq$  to augment the flexibility of the type system by determining when a type  $T$  is “smaller” than  $T'$ , it allows to use a process typed by the former whenever a process typed by the latter is required.

► **Definition 4.6** (Subsorting and subtyping). Subsorting  $\leq$ : is the least reflexive binary relation such that  $\mathbf{nat} \leq \mathbf{int} \leq \mathbf{real}$ . Subtyping  $\leq$  is the largest relation between session types coinductively defined by the following rules:

$$\begin{array}{c} \text{[SUB-END]} \\ \text{end} \leq \text{end} \end{array} \quad \frac{\text{[SUB-IN1]} \quad \forall i \in I : S'_i \leq S_i \quad T_i \leq T'_i \quad T \leq T'}{\&\{p^?\ell_i(S_i).T_i\}_{i \in I \cup J} \& \text{dt}\langle a \rangle.T \leq \&\{p^?\ell_i(S'_i).T'_i\}_{i \in I} \& \text{dt}\langle a \rangle.T'}$$

$$\frac{\text{[SUB-MOTION]} \quad T \leq T'}{\text{dt}\langle a \rangle.T \leq \text{dt}\langle a \rangle.T'} \quad \frac{\text{[SUB-IN2]} \quad \forall i \in I : S'_i \leq S_i \quad T_i \leq T'_i}{\&\{p^?\ell_i(S_i).T_i\}_{i \in I \cup J} \& \text{dt}\langle a \rangle.T \leq \&\{p^?\ell_i(S'_i).T'_i\}_{i \in I}}$$

$$\frac{\text{[SUB-IN3]} \quad T \leq T'}{\&\{p^?\ell_i(S_i).T_i\}_{i \in I} \& \text{dt}\langle a \rangle.T \leq \text{dt}\langle a \rangle.T'} \quad \frac{\text{[SUB-OUT]} \quad \forall i \in I : S_i \leq S'_i \quad T_i \leq T'_i}{\oplus\{p^!\ell_i(S_i).T_i\}_{i \in I} \leq \oplus\{p^!\ell_i(S'_i).T'_i\}_{i \in I \cup J}}$$

The double line in the subtyping rules indicates that the rules are interpreted *coinductively* [37, Chapter 21].

The typing rules for expressions are given as expected and omitted. The typing rules for processes and multiparty sessions are the content of Table 2:

- [T-SUB] is the *subsumption rule*: a process with type  $T$  is also typed by the supertype  $T'$ ;
- [T-0] says that a terminated process implements the terminated session type;
- [T-REC] types a recursive process  $\mu X.P$  with  $T$  if  $P$  can be typed as  $T$ , too, by extending the typing environment with the assumption that  $X$  has type  $T$ ;
- [T-VAR] uses the typing environment assumption that process  $X$  has type  $T$ ;
- [T-MOTION] types a motion process as a motion local type;
- [T-INPUT-CHOICE] types a summation of input prefixes as a branching type and a default branch as a motion type. It requires that each input prefix targets the same participant  $q$ , and that, for all  $i \in I$ , each continuation process  $P_i$  is typed by the continuation type  $T_i$ , having the bound variable  $x_i$  in the typing environment with sort  $S_i$ . Note that the rule implicitly requires the process labels  $\ell_i$  to be pairwise distinct (as per Definition 4.3);

■ **Table 2** Typing rules for motion processes.

$$\begin{array}{c}
\begin{array}{c}
\text{[T-0]} \\
\Gamma \vdash \mathbf{0} : \text{end}
\end{array}
\quad
\begin{array}{c}
\text{[T-REC]} \\
\frac{\Gamma, X : T \vdash P : T}{\Gamma \vdash \mu X.P : T}
\end{array}
\quad
\begin{array}{c}
\text{[T-VAR]} \\
\Gamma, X : T \vdash X : T
\end{array}
\quad
\begin{array}{c}
\text{[T-MOTION]} \\
\frac{\Gamma \vdash Q : T}{\Gamma \vdash \text{dt}\langle a \rangle.Q : \text{dt}\langle a \rangle.T}
\end{array}
\\
\\
\begin{array}{c}
\text{[T-OUT]} \\
\frac{\Gamma \vdash e : S \quad \Gamma \vdash P : T}{\Gamma \vdash \mathbf{q!}\ell(e).P : \mathbf{q!}\ell(S).T}
\end{array}
\quad
\begin{array}{c}
\text{[T-INPUT-CHOICE1]} \\
\frac{\forall i \in I \quad \Gamma, x_i : S_i \vdash P_i : T_i}{\Gamma \vdash \sum_{i \in I} \mathbf{q?}\ell_i(x_i).P_i : \&\{\mathbf{q?}\ell_i(S_i).T_i\}_{i \in I}}
\end{array}
\\
\\
\begin{array}{c}
\text{[T-INPUT-CHOICE2]} \\
\frac{\forall i \in I \quad \Gamma, x_i : S_i \vdash P_i : T_i \quad \Gamma \vdash \text{dt}\langle a \rangle.Q : T}{\Gamma \vdash \sum_{i \in I} \mathbf{q?}\ell_i(x_i).P_i + \text{dt}\langle a \rangle.Q : \&\{\mathbf{q?}\ell_i(S_i).T_i\}_{i \in I} \& T}
\end{array}
\\
\\
\begin{array}{c}
\text{[T-CHOICE]} \\
\frac{\Gamma \vdash e : \text{bool} \quad \exists k \in I \quad \Gamma \vdash P_1 : T_k \quad \Gamma \vdash P_2 : \oplus\{T_i\}_{i \in I \setminus \{k\}}}{\Gamma \vdash \text{if } e \text{ then } P_1 \text{ else } P_2 : \oplus\{T_i\}_{i \in I}}
\end{array}
\\
\\
\begin{array}{c}
\text{[T-SUB]} \\
\frac{\Gamma \vdash P : T \quad T \leq T'}{\Gamma \vdash P : T'}
\end{array}
\quad
\begin{array}{c}
\text{[T-SESS]} \\
\frac{\forall i \in I \quad \vdash P_i : G \upharpoonright p_i \quad \text{pt}\{G\} = \{p_i \mid i \in I\}}{\vdash \prod_{i \in I} p_i \triangleleft P_i : G}
\end{array}
\end{array}$$

- [T-OUT] types an output prefix with a singleton selection type, provided that the expression in the message payload has the correct sort  $S$ , and the process continuation matches the type continuation;
- [T-CHOICE] types a conditional process by matching the branches of the types to branches of the sub-processes;
- [T-SESS] types multiparty sessions, by associating typed processes to participants. It requires that the processes being composed in parallel can play as participants of a global communication protocol: hence, their types must be projections of a single global type  $G$ . As the temporal evolution (motion) synchronises all the processes condition  $\text{pt}\{G\} = \{p_i \mid i \in I\}$  guarantees that motions are defined for every participant.

► **Example 4.7.** We sketch the main steps to show that the Arm process is typed by the local type from Example 4.5. The type derivation uses the subtyping rules. This is because the process for the arm makes an external choice between the messages *fold*, *grab*, *done*, and the default motion primitive *idle*, and the type fixes a specific sequence of messages. The usual subtyping rules [SUB-IN1] and [SUB-IN2] allow typing the process against the local type, by “expanding” the local type with the other possible choices. The interesting subtyping rule is [SUB-IN3], which states that an external choice with a default motion type refines only the default motion type. This is needed to type the process against the local type

$$\text{dt}\langle \text{idle} \rangle. \text{Cart?grab}(\text{unit}).T$$

This subtyping rule is sound, because the local type ensures that the other message choices cannot arise.



The proposed motion session type system satisfies two fundamental properties: typed sessions only reduce to typed sessions (subject reduction), and typed sessions never get stuck.

In order to state subject reduction, we need to formalise how global types are reduced when local session types reduce and evolve. Note that since the same motion actions always synchronise among all participants, they always make progress (hence they are always consumed).

► **Definition 4.8** (Global types consumption and reduction). *The consumption of the communication  $p \xrightarrow{\ell} q$  and motion  $\text{dt}\langle a \rangle$  for the global type  $G$  (notation  $G \setminus p \xrightarrow{\ell} q$  and  $G \setminus \text{dt}\langle a \rangle$ ) is the global type defined (up to unfolding of recursive types) as follows:*

$$\begin{aligned} \text{dt}\langle a \rangle.G \setminus \text{dt}\langle a \rangle &= G \\ (p \rightarrow q : \{\ell_i(S_i).G_i\}_{i \in I}) \setminus p \xrightarrow{\ell} q &= G_k && \text{if } \exists k \in I : \ell = \ell_k \\ (r \rightarrow s : \{\ell_i(S_i).G_i\}_{i \in I}) \setminus p \xrightarrow{\ell} q &= r \rightarrow s : \{\ell_i(S_i).G_i \setminus p \xrightarrow{\ell} q\}_{i \in I} \\ &&& \text{if } \{r, s\} \cap \{p, q\} = \emptyset \wedge \forall i \in I : \{p, q\} \subseteq G_i \end{aligned}$$

The reduction of global types is the smallest pre-order relation closed under the rule:  $G \Longrightarrow G \setminus p \xrightarrow{\ell} q$  and  $G \Longrightarrow G \setminus \text{dt}\langle a \rangle$ .

We can now state the main results.

► **Theorem 4.9** (Subject Reduction). *Let  $\vdash M : G$ . For all  $M'$ , if  $M \longrightarrow M'$ , then  $\vdash M' : G'$  for some  $G'$  such that  $G \Longrightarrow G'$ .*

► **Corollary 4.10**. *Let  $\vdash M : G$ . If  $M \longrightarrow^* M'$ , then  $\vdash M' : G'$  for some  $G'$  such that  $G \Longrightarrow G'$ .*

► **Theorem 4.11** (Progress). *If  $\vdash M : G$ , then either  $M \equiv \prod_{i \in I} p_i \triangleleft \mathbf{0}$  or there is  $M'$  such that  $M \longrightarrow M'$ .*

As a consequence of subject reduction and progress, we get the safety property stating that a typed multiparty session will never get stuck.

► **Theorem 4.12** (Type Safety). *If  $\vdash M : G$ , then it does not hold  $\text{stuck}(M)$ .*

**Proof.** Direct consequence of Corollary 4.10, Theorem 4.11, and Definition 3.2. ◀

## 5 Motion Primitives: Trajectories and Resources

So far, our motion calculus abstracted the trajectories of the robots and only considered the time it takes to execute motion primitives. This is sufficient to show that the synchronisation and communication protocol between the robots executes correctly. However, it is too abstract to prove more complex properties about executions of the system. In particular, for an execution to proceed correctly we need to check the existence of trajectories for all the robots. A joint trajectory may not exist, for example, if the motion primitives cause a collision in the physical world.

In this section, we explain how to make our model more detailed and how to look inside the motion primitives for the continuous evolution of trajectories. To accomplish this, first, we give a semantics that includes trajectories. Then, we refine our calculus to replace internal choice with guarded choice. Finally, we explain how to use session types to prove properties over the trajectories.

## 5.1 Model for the Robots and Motion Primitives

We proceed following the formalisation of trajectories in the PGCD language for robotics [4].

**Robots.** Each participant  $(p, q, \dots)$  maintains a state in the physical world. This state is updated when its own motion primitives execute as well as on potential physical interactions with other processes.

We model the physical state of a process as a tuple  $(Var, \rho, rsrc)$  where  $Var$  is a set of variables, with two distinguished disjoint subsets  $X$  and  $W$  of *physical state* and *external input* variables,  $\rho : Var \rightarrow \mathbb{R}$  is a *store* mapping variables to values, and  $rsrc$  is a *resource function* mapping a store to a subset of  $\mathbb{R}^3$ . The resource function represents the geometric footprint in space occupied by the robot. We shall use this function to check the absence of collisions between robots.

When two robots  $p_1$  and  $p_2$  are in the same environment, we may connect some state variables of one process to the external inputs of the other. This represents physical coupling between these robots. A *connection*  $\theta$  between  $p_1$  and  $p_2$  is a finite set of pairs of variables,  $\theta = \{(x_i, w_i) \mid i = 1, \dots, m\}$ , such that: (1) for each  $(x, w) \in \theta$ , we have  $x \in p_1.X$  and  $w \in p_2.W$  or  $x \in p_2.X$  and  $w \in p_1.W$ , and (2) there does not exist  $(x, w), (x', w) \in \theta$  such that  $x$  and  $x'$  are distinct. Two connections  $\theta_1$  and  $\theta_2$  are *compatible* if  $\theta_1 \cup \theta_2$  is a connection. We assume that all the participants in a session are connected by compatible connections.

For example, consider a cart and an arm. The physical variables can provide the position and velocities of the center of mass of the cart and of the arm. Note that if the arm is attached to the cart, then its position changes when the cart moves. Thus, the position and velocity of the cart are external inputs to the arm, and play a role in determining its own position. However, the arm can also move relative to the cart and the position of its end effector is determined both by the external inputs as well as its relative position and velocity. Furthermore, the mass and the position of the center of mass of the arm are external inputs to the cart, because these variables affect the dynamics of the cart.

**Motion Primitives.** Let  $X$  and  $W$  be two sets of real-valued variables, representing internal state and external input variables of a robotic system, respectively. A motion primitive updates the values of the variables in  $X$  over time, while respecting the values of variables in  $W$  set by the external world. This dynamic process results in a pair of state and input trajectories  $(\xi, \nu)$ , i.e., a valuation over time to variables in  $X$  and  $W$ .

Formally, a motion primitive  $m$  is a tuple  $(T, \text{Pre}, \text{Inv}, \text{Post})$  consisting of a *duration*  $T$ , a *pre-condition*  $\text{Pre} \subseteq \mathbb{R}^{|X|} \times \mathbb{R}^{|W|}$ , an *invariant*  $\text{Inv} \subseteq ([0, T] \rightarrow \mathbb{R}^{|X|}) \times ([0, T] \rightarrow \mathbb{R}^{|W|})$ , and a *post-condition*  $\text{Post} \subseteq \mathbb{R}^{|X|} \times \mathbb{R}^{|W|}$ . A *trajectory* of duration  $T$  of the motion primitive  $m$  is a pair of continuous functions  $(\xi, \nu)$  mapping the real interval  $[0, T]$  to  $\mathbb{R}^{|X|}$  and  $\mathbb{R}^{|W|}$ , respectively, such that  $(\xi, \nu) \in \text{Inv}$ ,  $(\xi(0), \nu(0)) \in \text{Pre}$ , and  $(\xi(T), \nu(T)) \in \text{Post}$ .

Correspondingly, we need to update the semantics of our motion calculus:

- The participant executing a program  $p \triangleleft P$  now also carries a store containing a valuation for the physical state of the robot:  $p, \rho \triangleleft P$ .
- The motion transitions  $\xrightarrow{\text{dt}(a)}$  get labelled with trajectories:  $\xrightarrow{\text{dt}((\xi, \nu))}$ .
- The semantics rule for choice can use values from the store:

$$\frac{[\text{T-CONDITIONAL}] \quad \rho(e) \downarrow \text{true}}{p, \rho \triangleleft \text{if } e \text{ then } P \text{ else } Q} \longrightarrow p, \rho \triangleleft P \quad \frac{[\text{F-CONDITIONAL}] \quad \rho(e) \downarrow \text{false}}{p, \rho \triangleleft \text{if } e \text{ then } P \text{ else } Q} \longrightarrow p, \rho \triangleleft Q$$

where  $\rho(e)$  replaces the variables from  $Var$  in  $e$  with their value according to  $\rho$ .

## 28:18 Motion Session Types for Robotic Interactions

- The semantics of a motion checks the trajectories against the motion primitive specification and the store:

$$\begin{array}{c}
 \text{[MOTION]} \\
 \frac{a = (T, \text{Pre}, \text{Inv}, \text{Post}) \quad \text{range}(\xi) = [0, T] \quad \rho = \xi(0) \quad \rho' = \xi(T) \\
 (\xi(0), \nu(0)) \in \text{Pre} \quad (\xi(T), \nu(T)) \in \text{Post} \quad \forall t \in [0, T]. (\xi(t), \nu(t)) \in \text{Inv}}
 {p, \rho \triangleleft \text{dt}(a).P \xrightarrow{\text{dt}(\langle \xi, \nu \rangle)} p, \rho' \triangleleft P}
 \end{array}$$

The rule checks that the trajectory is valid w.r.t.  $a$ : the duration of the trajectory must match the duration of the motion primitive, the start and end of the trajectory match the state of  $\rho$  and  $\rho'$  respectively. Furthermore, the pre-condition, post-condition, and invariant must be respected.

- The parallel composition of motions connects the external inputs of each process according to the connections. For the notations, we use subscript to denote that an element belongs to a particular process  $p$ , e.g.,  $X_p$  for the internal variables of  $p$ . We denote the restriction of a trajectory  $\xi$  over a subset  $X$  of the dimensions by  $\xi|_X$ .

$$\begin{array}{c}
 \text{[M-PAR]} \\
 \frac{\forall i \quad \xi_i = \xi|_{X_{p_i}} \quad \nu_i = \theta_{p_i}(\xi)|_{W_{p_i}} \quad p_i, \rho_i \triangleleft P_i \xrightarrow{\text{dt}(\langle \xi_i, \nu_i \rangle)} p_i, \rho'_i \triangleleft P'_i \\
 \forall i, j, t. i \neq j \Rightarrow \text{rsrc}_{p_i}(\xi|_{X_{p_i}}(t), \theta_{p_i}(\xi)|_{W_{p_i}}(t)) \cap \text{rsrc}_{p_j}(\xi|_{X_{p_j}}(t), \theta_{p_j}(\xi)|_{W_{p_j}}(t)) = \emptyset}
 {\Pi_i p_i, \rho_i \triangleleft P_i \xrightarrow{\text{dt}(\langle \xi, \nu \rangle)} \Pi_i p_i, \rho'_i \triangleleft P'_i}
 \end{array}$$

Even at the top level, there is a  $\nu$  as there can be elements which are under the control of the environment. Then, for each process we create the appropriate trajectory  $(\xi, \nu)$  by applying the appropriate connection  $\theta$ . Also, the resources used by each participants during the motion needs to disjoint from each other. This last check ensures the absence of collision between robots. We use this check to avoid the complexity of modelling collisions.

► **Example 5.1.** Let us look at the cart from Example 3.3. The cart is moving on the ground, a 2D plane and, therefore, we model its physical state ( $X_{\text{Cart}}$ ) by its position  $\mathbf{p}_{\text{Cart}} \in \mathbb{R}^2$ , orientation  $\mathbf{r}_{\text{Cart}} \in [-\pi; \pi)$ , and speed  $\mathbf{s}_{\text{Cart}} \in \mathbb{R}$ .

A trivial motion primitive  $\text{idle}(\mathbf{p}_0, \mathbf{r}_0)$  keeps the cart at its current position  $\mathbf{p}_0$  and orientation  $\mathbf{r}_0$ ; the pre-condition is  $\mathbf{s}_{\text{Cart}} = 0$  (i.e., it is at rest), the post-condition is  $\mathbf{s}_{\text{Cart}} = 0 \wedge \mathbf{p}_{\text{Cart}} = \mathbf{p}_0 \wedge \mathbf{r}_{\text{Cart}} = \mathbf{r}_0$ , and the invariant is  $\mathbf{p}_{\text{Cart}}(t) = \mathbf{p}_0 \wedge \mathbf{r}_{\text{Cart}}(t) = \mathbf{r}_0 \wedge \mathbf{s}_{\text{Cart}} = 0$  for all  $t \in [0, T]$ .

A slightly more interesting motion primitive is  $\text{move}(\mathbf{p}_0, \mathbf{p}_t)$ , which moves the cart from position  $\mathbf{p}_0$  to  $\mathbf{p}_t$ . The pre-condition is  $\mathbf{s}_{\text{Cart}} = 0 \wedge \mathbf{p}_{\text{Cart}} = \mathbf{p}_0$ . The post-condition is  $\mathbf{s}_{\text{Cart}} = 0 \wedge \mathbf{p}_{\text{Cart}} = \mathbf{p}_t$ . The invariant can specify a bound on the velocity, e.g.,  $0 \leq \mathbf{s}_{\text{Cart}} \leq v_{\text{max}}$ , that the cart moves in straight line between  $\mathbf{p}_0$  and  $\mathbf{p}_t$ , etc.

We can also include external input. For instance, we may add an external variable  $\mathbf{w}_{\text{obj}}$  to represent the weight of any carried object, e.g., the arm attached on top. Then, the pre-condition of  $\text{move}$  may include an extra constraint  $0 \leq \mathbf{w}_{\text{obj}} \leq w_{\text{max}}$  to say that the cart can only move if the weight of the payload is smaller than a given bound.

## 5.2 Motion Calculus with Guarded Choice

Before executing some motion, a process may need to test the state of the physical world and, according to the current state, decide what to do. Therefore, we extend the calculus with the ability for a process to test predicates over its  $\text{Var}$  as part of the  $\text{if} \cdot \text{then} \cdot \text{else} \cdot$ . On the specification side, we also add predicates to the internal choice.

Let  $\mathcal{P}$  range over predicates. The global and local motion session types are modified as follows:

- The branching type for global session types becomes  $\mathbf{p} \rightarrow \mathbf{q} : \{\{\mathcal{P}_i\} \ell_i(S_i).G_i\}_{i \in I}$ .
- The branching type for local session types becomes  $\oplus\{\{\mathcal{P}_i\} \mathbf{q}! \ell_i(S_i).T_i\}_{i \in I}$ .

To make sure the modified types can be projected and then used for typing they need to respect the following constraints. Assume that  $Var_{\mathbf{p}}$  are the variables associated with the robot executing the role of  $\mathbf{p}$ . (1) The choices are *local*, i.e., for  $\mathbf{p} \rightarrow \mathbf{q} : \{\{\mathcal{P}_i\} \ell_i(S_i).G_i\}_{i \in I}$  we have that  $\text{fv}(\mathcal{P}_i) \subseteq Var_{\mathbf{p}}$  for all  $i$  in  $I$ . (2) The choices are *total*, i.e., for  $\mathbf{p} \rightarrow \mathbf{q} : \{\{\mathcal{P}_i\} \ell_i(S_i).G_i\}_{i \in I}$  we have that  $\bigvee_{i \in I} \mathcal{P}_i$  is valid. The local types have similar constraints.

The subtyping and typing relation are updated as follows:

$$\begin{array}{c} \text{[SUB-OUT]} \\ \hline \forall i \in I : \quad S_i \leq S'_i \quad T_i \leq T'_i \quad \mathcal{P}_i \Rightarrow \mathcal{P}'_i \\ \hline \oplus\{\{\mathcal{P}_i\} \mathbf{p}! \ell_i(S_i).T_i\}_{i \in I} \leq \oplus\{\{\mathcal{P}'_i\} \mathbf{p}! \ell_i(S'_i).T'_i\}_{i \in I \cup J} \end{array}$$

The change in this rule is the addition of checking the implication  $\mathcal{P}_i \Rightarrow \mathcal{P}'_i$  to make sure that if the pre-condition of a motion primitive relies on  $\mathcal{P}'_i$ , it still holds with  $\mathcal{P}_i$ . Notice that  $\oplus\{\{\mathcal{P}_i\} \mathbf{p}! \ell_i(S_i).T_i\}_{i \in I}$  which can have more restricted predicates needs to be a valid local type and the guards still need to be total.

$$\begin{array}{c} \text{[T-CHOICE]} \\ \hline \Gamma \vdash e : \text{bool} \quad \exists k \in I \quad e \Rightarrow \mathcal{P}_k \quad \Gamma \vdash P_1 : T_k \quad \Gamma \vdash P_2 : \oplus\{[e \vee \mathcal{P}_i]T_i\}_{i \in I \setminus \{k\}} \\ \hline \Gamma \vdash \text{if } e \text{ then } P_1 \text{ else } P_2 : \oplus\{[\mathcal{P}_i]T_i\}_{i \in I} \end{array}$$

Type checking the rules propagates the expression from **if then else** and matches it into a branch of the type. To deal with the **else** branch we modify the predicate in the remaining branches of the type. For the last **else** branch of a, possibly nested, **if then else** we need the following extra rule:

$$\begin{array}{c} \text{[T-CHOICE-FINAL]} \\ \hline \Gamma \vdash P : T \\ \hline \Gamma \vdash P : \oplus\{[\text{true}]T\} \end{array}$$

► **Example 5.2.** Usually, for the propagation of tested expressions through the branches we modify the type. Let us make an example of how this works. Consider we have the following process **if**  $e_1$  **then**  $P_1$  **else**  $P_2$  which has the type  $\oplus\{[e_1]T_1, [\neg e_1]T_2\}$ . Assuming that  $P_i : T_i$  for  $i \in \{1, 2\}$  we can build the following derivation:

$$\frac{e_1 \Rightarrow e_1 \quad \Gamma \vdash P_1 : T_1 \quad \frac{\Gamma \vdash P_2 : T_2}{\Gamma \vdash P_2 : \oplus\{[e_1 \vee \neg e_1]T_2\}}}{\Gamma \vdash \text{if } e_1 \text{ then } P_1 \text{ else } P_2 : \oplus\{[e_1]T_1, [\neg e_1]T_2\}}$$

With a bit of boolean algebra, we can show that  $e_1 \vee \neg e_1 \Leftrightarrow \text{true}$ .

### 5.3 Existence of Joint Trajectories and Verification

The goal of the compatibility check is to make sure that abstract motion primitives specified in a global type can execute concurrently. This requires two checks. First, for motion primitives of different processes executed in parallel, we need to make sure that there exists a trajectory

satisfying all the constraints of the motion primitives. Second, for motion primitives executed sequentially by the same process, we need to make sure that the post-condition of the first implies the pre-condition of the second motion primitive, taking into account the guards of choices in the middle.

To check that motion primitives executing in parallel have a joint trajectory, we use an assume-guarantee style of reasoning. When two processes are attached, one process relies on the invariants of the other’s output (which can be an external input) to satisfy its own invariant and vice versa. We refer to standard methods [35, 4] for the details.

For the allowed trajectories, we need to also check the absence of collision. This means that once we have the constraints defining a joint trajectory  $\xi$  to check that for any two distinct processes  $p$  and  $q$  the property  $rsrc_p(\xi_p) \cap rsrc_q(\xi_q) = \emptyset$ .

► **Example 5.3.** In Example 3.4, the cart and the carrier are moving toward each other. They need to be close enough for the arm to grab the object but far enough to avoid colliding. We model the resources of the cart by a cylinder around the cart’s position:  $rsrc_{\text{Cart}} = \{(x, y, z) \mid |(x, y) - \mathbf{p}_{\text{Cart}}| \leq r \wedge 0 \leq z \leq h\}$  where  $r$  is the “radius” of the cart and  $h$  its height. The carrier’s resources are similar but with the appropriate radius and height  $r'$ ,  $h'$ . The cart and carrier does not collide if we can prove that  $\forall t. |\xi_{\text{Cart}}|_{\mathbf{p}_{\text{Cart}}}(t) - \xi_{\text{Carrier}}|_{\mathbf{p}_{\text{Carrier}}}(t)| > r + r'$ .

## 6 Evaluation

### 6.1 Implementation

We have implemented the system we describe on top of PGCD [4]<sup>1</sup>, a system for programming and verification of robotic systems. PGCD is build on top of the Robotic Operating System (ROS) [40], a software ecosystem for robots. The core of ROS is a publish-subscribe messaging system. PGCD uses ROS’s messaging to implement its synchronous message-passing layer. On the verification side, PGCD uses a mix of model-checking (using SPIN [22]) to deal with the message-passing structure, and symbolic reasoning (using SYMPY [33]) and constraint solving (using DREAL [16]) to reason about motion primitives.

We replace the global model-checking algorithm of PGCD with motion session calculus specifications but reuse PGCD’s infrastructure to reason about the trajectories of motion primitives. Currently, our implementation uses a syntax for specifications closer to the state-machine form of session types [11] but without the parallel composition operator. This representation allows for more general guarded choice. `if · then · else ·` implicitly forces disjoint guards for the two branches. Our implementation allows overlapping guards. Algorithmically, since the types are represented in a form close to an automaton, the projection and merge operations are implemented using automata theoretic operation: morphism, minimisation, and checking determinism of the result.

The typing, including subtyping, is implemented by computing an alternating simulation [2] between programs and their respective local type. Intuitively, an alternative refinement relation check that a process implements its specification (subset of the behaviours) without restricting the other processes. For synchronous message passing programs, the subtyping relation for session type matches alternating refinement. We use this view on subtyping as the theory of alternating simulation [2] gives us an algorithm to compute this relation and, therefore, check subtyping.

<sup>1</sup> PGCD repository is <https://github.com/MPI-SWS/pgcd>. The code for this work is located in the `pgcd/nodes/verification/choreography` folder.

## 6.2 Experiments

For the evaluation, we take two existing PGCD programs and write global types in motion session calculus that describe the co-ordination in the program.

First, we describe our experimental setup, both for the hardware and for the software. Then, we explain the experiments. Finally, we report on the size of the specifications, and time to check the programs satisfy the specification.

### Setup

We use three robots: a robotic arm and two carts, shown in Figure 4. The robots are built with a mix of off-the-self parts and 3D printed parts.

**Arm.** The arm is a modified BCN3D MOVEO,<sup>2</sup> where the upper arm section is shortened to make it lighter and easier to mount on the cart. The arm with its control electronics is mounted on top of the cart.

**Cart.** The cart is shown on Figure 4a. The control electronics and motors are situated below the wooden board. The cart is an omnidirectional driving platform. It uses omniwheels to get three degrees of freedom (two in translation, one in rotation) and can move between any two positions on a flat ground. The advantage of using such wheels is that all the three degrees of freedom are controllable and movement does not require complex planning. Due to the large power consumption of the arm mounted on top, this cart is powered by a tether.

**Carrier.** We call the second cart the carrier (Figure 4b) as we use it to carry the block that is grabbed by the arm. As the first cart, it is also omnidirectional (mecanum wheels).

All the three robots use stepper motors to move precisely. The robots do not have feedback on their position and keep track of their state using *dead reckoning*, i.e., they know their initial state and then they update their virtual state by counting the number of steps the motors turns. If we control slippage and do not exceed the maximum torque of the motors, there is little accumulation of error as long as the initial state is known accurately. In our experiments, we use markings on the ground to fix the initial state as can be seen in Figure 5. Furthermore, using stepper motors allows us to know the time it takes to execute a given motion primitive by fixing the rate of steps.

Each robot has a RaspberryPi 3 model B to run the program. The ROS master node, providing core messaging services, runs on a separate laptop to which all the robots connect. The RaspberryPi runs Raspbian OS (based on Debian Jessie) and the laptop runs Ubuntu 16.04. The ROS version is Kinetic Kame.

### Experiments

We describe two experiments:


**Handover.** This experiment corresponds to our earlier example. The two carts meet before the arm takes an object placed on top of the carrier and, then, they go back to their initial position (see Figure 5a).

**Underpass.** First, the carrier cart brings an object to the arm which is then taken by the arm. Then, the carrier cart goes around the arm passing under an obstacle which is high enough for just the carrier alone. Finally, the arm puts the object back on the carrier on the other side of the obstacle. This can be seen in Figure 5b.


---

<sup>2</sup> <https://github.com/BCN3D/BCN3D-Moveo>






(a) The cart and arm robots attached together.




(b) The carrier robot.

■ **Figure 4** Robots used in our experiments.



(a) Handover.



(b) Underpass.

■ **Figure 5** Composite images of the experiments. (a) For handover, a cart containing an object moves close to the cart with the attached arm. The arm picks up the object. (b) For underpass, the carrier containing an object moves near the underpass. The arm picks up the object. The carrier moves under the underpass and moves close to the arm. The arm places the object on the carrier.

Composite images (combination of multiple frame of the video) are shown in Figure 5. The carts implement motion explicitly using the motion primitives (move straight, strafe, rotate). For instance, when going around the cart in the second experiment, the carrier executes rotate, move straight, rotate, strafe. In the model of the resources, we exclude the gripper from the footprint and we do not model the objects gripped (gripping is a collision). For the environment, we model obstacles as regions of  $\mathbb{R}^3$  and also test for collision against these regions.

Table 3 shows the size of the programs in the language of PGCD (sum for all the robots) and the size of the global specifications. As part of the program we include a description of the environment which specifies the initial states of the robots and the obstacles used for additional collision checks. Finally, we show the number of verification conditions (#VCs) generated during the subtyping and the checks for joint trajectories. The total running time includes all the steps, i.e., checking the global specification, projection, typing, the existence of joint trajectories, and the absence of collision. The running time is dominated by the check on trajectories and collisions. The motion primitives (implementation and specification) are taken from PGCD without any change and represent around 1K lines of codes for all three robots.



■ **Table 3** Programs, Specification, and Checks.

Experiment	Program (LoC)	Specification (LoC)	#VCs	Time (sec.)
Handover	22	12	141	38
Underpass	29	22	302	56

Compared to the verification results presented with PGCD [4, Section 5], we have roughly a  $2\times$  speed-up. The reason is that PGCD is used model-checking instead of global/local types. The motion session calculus makes it possible to have an abstract global specification which is easier to check.

In conclusion, our evaluation demonstrates that session types allow the specification of non-trivial co-ordination tasks between multiple robots with reasonable effort, while allowing automated and compositional verification.

## 7 Related Work

There is considerable interest in the robotics community on designing modular robotic components from higher-level specifications [32, 19]. However, most of this work has focused on descriptions for the physical and electronic design of components or on generating plans from higher level specifications rather than on language abstractions and types to reason about concurrency and motion. The interaction between concurrency and dynamics, and the use of automated verification techniques were considered in PGCD [4]. Our work takes PGCD as a starting point and formalises a compositional verification methodology through session types.

At the specification level, hybrid process algebras and other models of hybrid systems [1, 41, 7, 38] can model concurrent hybrid systems. However, these papers do not provide a direct path to implementation. Hybrid extensions to synchronous reactive languages [6, 5] describe programs which interact through events and control physical variables. Most existing verification methodologies for these programs rely on global model checking rather than on types. Our choice of session types is inspired by efficient type checking but also as the basis for describing interface specifications for components.

Extensions and applications of multiparty session types have been proposed in many different settings. See, e.g. [27, 3, 17]. We discuss only most related work. The work [9] extends multiparty session types with time, to enable the verification of realtime distributed systems. This extension with time allows specifications to express properties on the causalities of interactions, on the carried data types, and on the times in which interactions occur. The projected local types correspond to Communicating Timed Automata (CTA). To ensure the progress and liveness properties for projected local types, the framework requires several additional constraints on the shape of global protocols, such as feasibility condition (at any point of the protocol the current time constraint should be satisfiable for any possible past) and a limitation to the recursion where in the loop, the clock should be always reset. The approach is implemented in Python in [34] for runtime monitoring for the distributed system. Later, the work in [8] develops more relaxed conditions in CTAs, and applies them to synthesise timed global protocols. Unlike our work, no type checking for processes is studied in [8]. The main difference from [9, 34, 8] is that our approach does not rely on CTAs and is more specific to robotics applications where the verification is divided into the two layers; (1)

a simple type check for processes with motion primitives to ensure communication deadlock-freedom with global synchronisations; and (2) additional more refined checks for trajectories and resources. This two layered approach considerably simplifies our core calculus and typing system in Section 4, allowing to verify more complex scenarios for robotics interactions.

## 8 Conclusion

We have outlined a unifying programming model and typing discipline for communication-centric systems that sense and actuate the physical world. We work in the framework of multiparty session types [25, 26], which have proved their worth in many different scenarios relating to “pure” concurrent software systems. We show how to integrate motion primitives into a core calculus and into session types. We demonstrate how multiparty session types are used to specify correct synchronisation among multiple participants: we first provide a basic progress guarantee for communications and synchronisation by motion primitives, which is useful to extend richer verification related to trajectories.

At this point, our language is a starting point and not a panacea for robotics programming. Decoupling specifications into parallel and/or sequential tasks and using distributed controllers assumes “loosely coupled dynamics.” In some examples, such as a multiple cart/arm co-ordination control, it may not be easy to assume a purely distributed control strategy based on independent motion primitives. We are thus exploring simultaneous concurrent programming and distributed controller and co-ordinator synthesis. As an example, assume that we have two cart/arm compositions which should lift one object together. In particular we can assume that lifting the object with only one arm would cause the cart/arm compositions to tilt over, which generates a strong coupling between all components during the coordinated lift of the object. Our framework allows to easily synchronise all the components. However, in any realistic scenario a robust controller would need (almost) continuous feedback between all components to fulfill the coordinated lift task. Thus, our model of loosely coupled motion primitives, one per component, may be too weak or incur too much communication and bandwidth overhead for a real implementation.

Going in this direction, we need a better way to integrate specifications of controllers (motion primitives) and their robustness. This would also enable a more realistic non-synchronous model for the communication [31] and, after checking some robustness condition on the controller, rigorously show that the synchronous idealised model is equivalent to the more realistic model, i.e., considering delay in the communication as disturbances for the motion primitives. We also plan to tackle channel passing. The challenge is that the physical world (time and space) is hard to isolate: for instance, time is an implicit synchronisation which occurs at the same time across all sessions.

Finally, robotics applications manipulate physical state and time as *resources*. An interesting open question is how resource-based reasoning techniques such as separation logics for concurrency [36, 28] can be repurposed to reason about separation of components in space and time.

---

## References

- 1 R. Alur and T.A. Henzinger. Modularity for Timed and Hybrid Systems. In *CONCUR '97: Concurrency Theory*, volume 1243 of *LNCS*, pages 74–88. Springer, 1997.
- 2 R. Alur, T.A. Henzinger, O. Kupferman, and M.Y. Vardi. Alternating refinement relations. In *CONCUR'98 Concurrency Theory*, pages 163–178. Springer, 1998.

- 3 Davide Ancona, Viviana Bono, Mario Bravetti, Joana Campos, Giuseppe Castagna, Pierre-Malo Denielou, 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. *FTPL*, 3(2-3):95–230, 2016.
- 4 Gregor B. Banusic, Rupak Majumdar, Marcus Pirron, Anne-Kathrin Schmuck, and Damien Zufferey. PGCD: robot programming and verification with geometry, concurrency, and dynamics. In Xue Liu, Paulo Tabuada, Miroslav Pajic, and Linda Bushnell, editors, *Proceedings of the 10th ACM/IEEE International Conference on Cyber-Physical Systems, ICCPS 2019, Montreal, QC, Canada, April 16-18, 2019*, pages 57–66. ACM, 2019. doi:10.1145/3302509.3311052.
- 5 Kerstin Bauer and Klaus Schneider. From synchronous programs to symbolic representations of hybrid systems. In *Proceedings of the 13th ACM International Conference on Hybrid Systems: Computation and Control, HSCC 2010, Stockholm, Sweden, April 12-15, 2010*, pages 41–50. ACM, 2010. doi:10.1145/1755952.1755960.
- 6 Albert Benveniste, Timothy Bourke, Benoît Caillaud, Jean-Louis Colaço, Cédric Pasteur, and Marc Pouzet. Building a Hybrid Systems Modeler on Synchronous Languages Principles. *Proceedings of the IEEE*, 106(9):1568–1592, 2018. doi:10.1109/JPROC.2018.2858016.
- 7 J.A. Bergstra and C.A. Middelburg. Process algebra for hybrid systems. *Theoretical Computer Science*, 335(2):215–280, 2005. Process Algebra. doi:10.1016/j.tcs.2004.04.019.
- 8 Laura Bocchi, Julien Lange, and Nobuko Yoshida. Meeting Deadlines Together. In *26th International Conference on Concurrency Theory*, volume 42 of *LIPIcs*, pages 283–296. Schloss Dagstuhl, 2015.
- 9 Laura Bocchi, Weizhen Yang, and Nobuko Yoshida. Timed Multiparty Session Types. In *25th International Conference on Concurrency Theory*, volume 8704 of *LNCS*, pages 419–434. Springer, 2014.
- 10 Mario Coppo, Mariangiola Dezani-Ciancaglini, Luca Padovani, and Nobuko Yoshida. A Gentle Introduction to Multiparty Asynchronous Session Types. In *SFM*, volume 9104 of *LNCS*, pages 146–178. Springer, 2015.
- 11 Pierre-Malo Denielou and Nobuko Yoshida. Multiparty Session Types Meet Communicating Automata. In *ESOP 2012 - European Symposium on Programming*. Springer, 2012. doi:10.1007/978-3-642-28869-2\_10.
- 12 Pierre-Malo Denielou, 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.
- 13 Pierre-Malo Denielou, 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.
- 14 Mariangiola Dezani-Ciancaglini, Silvia Ghilezan, Svetlana Jaksic, Jovanka Pantovic, and Nobuko Yoshida. Precise subtyping for synchronous multiparty sessions. In *PLACES*, volume 203 of *EPTCS*, pages 29–43, 2015. doi:10.4204/EPTCS.203.3.
- 15 Mariangiola Dezani-Ciancaglini, Silvia Ghilezan, Svetlana Jaksic, Jovanka Pantovic, and Nobuko Yoshida. Denotational and Operational Preciseness of Subtyping: A Roadmap. In *Theory and Practice of Formal Methods: Essays Dedicated to Frank de Boer on the Occasion of His 60th Birthday*, volume 9660 of *LNCS*, pages 155–172. Springer, 2016.
- 16 Sicun Gao, Soonho Kong, and Edmund M. Clarke. dReal: An SMT solver for nonlinear theories over the reals. In Maria Paola Bonacina, editor, *Automated Deduction - CADE-24 - 24th International Conference on Automated Deduction, Lake Placid, NY, USA, June 9-14, 2013. Proceedings*, volume 7898 of *Lecture Notes in Computer Science*, pages 208–214. Springer, 2013. doi:10.1007/978-3-642-38574-2\_14.
- 17 Simon Gay and Antonio Ravera, editors. *Behavioural Types: from Theory to Tools*. River Publishers, 2017.

- 18 Silvia Ghilezan, Svetlana Jaksic, Jovanka Pantovic, Alceste Scalas, and Nobuko Yoshida. Precise subtyping for synchronous multiparty sessions. *J. Log. Algebr. Meth. Program.*, 104:127–173, 2019. doi:10.1016/j.jlamp.2018.12.002.
- 19 Sehoon Ha, Stelian Coros, Alexander Alspach, James M. Bern, Joohyung Kim, and Katsu Yamane. Computational Design of Robotic Devices From High-Level Motion Specifications. *IEEE Trans. Robotics*, 34(5):1240–1251, 2018. doi:10.1109/TR0.2018.2830419.
- 20 Thomas A. Henzinger. Sooner is Safer Than Later. *Inf. Process. Lett.*, 43(3):135–141, 1992. doi:10.1016/0020-0190(92)90005-G.
- 21 Thomas A. Henzinger. The Theory of Hybrid Automata. In *Proceedings, 11th Annual IEEE Symposium on Logic in Computer Science, New Brunswick, New Jersey, USA, July 27-30, 1996*, pages 278–292. IEEE Computer Society, 1996. doi:10.1109/LICS.1996.561342.
- 22 G.J. Holzmann. The Model Checker SPIN. *IEEE Trans. Software Eng.*, 23(5):279–295, 1997. doi:10.1109/32.588521.
- 23 Kohei Honda. Types for Dyadic Interaction. In *CONCUR'93*, pages 509–523, 1993.
- 24 Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. Language Primitives and Type Disciplines for Structured Communication-based Programming. In *ESOP*, volume 1381 of *LNCS*, pages 22–138. Springer, 1998. doi:10.1007/BFb0053567.
- 25 Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty Asynchronous Session Types. In *POPL*, pages 273–284. ACM Press, 2008.
- 26 Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty Asynchronous Session Types. *Journal of ACM*, 63:1–67, 2016.
- 27 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), 2016. doi:10.1145/2873052.
- 28 R. Jung, D. Swasey, F. Sieczkowski, K. Svendsen, A. Turon, L. Birkedal, and D. Dreyer. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *POPL 15*, pages 637–650. ACM, 2015.
- 29 Dimitrios Kouzapas and Nobuko Yoshida. Globally Governed Session Semantics. In Pedro R. D'Argenio and Hernán C. Melgratti, editors, *CONCUR*, volume 8052 of *LNCS*, pages 395–409. Springer, 2013.
- 30 Dimitrios Kouzapas and Nobuko Yoshida. Globally Governed Session Semantics. *Logical Methods in Computer Science*, 10(4), 2015.
- 31 Julien Lange, Emilio Tuosto, and Nobuko Yoshida. From communicating machines to graphical choreographies. In *42nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 221–232. ACM, 2015.
- 32 A.M. Mehta, N. Bezzo, P. Gebhard, B. An, V. Kumar, I. Lee, and D. Rus. A Design Environment for the Rapid Specification and Fabrication of Printable Robots. *Experimental Robotics*, pages 435–449, 2015.
- 33 Aaron Meurer, Christopher P. Smith, Mateusz Paprocki, Ondřej Čertík, Sergey B. Kirpichev, Matthew Rocklin, AMiT Kumar, Sergiu Ivanov, Jason K. Moore, Sartaj Singh, Thilina Rathnayake, Sean Vig, Brian E. Granger, Richard P. Muller, Francesco Bonazzi, Harsh Gupta, Shivam Vats, Fredrik Johansson, Fabian Pedregosa, Matthew J. Curry, Andy R. Terrel, Štěpán Roučka, Ashutosh Saboo, Isuru Fernando, Sumith Kulal, Robert Cimrman, and Anthony Scopatz. SymPy: symbolic computing in Python. *PeerJ Computer Science*, 3:e103, January 2017. doi:10.7717/peerj-cs.103.
- 34 Rumyana Neykova, Laura Bocchi, and Nobuko Yoshida. Timed runtime monitoring for multiparty conversations. *Formal Asp. Comput.*, 29(5):877–910, 2017.
- 35 Pierluigi Nuzzo. *Compositional Design of Cyber-Physical Systems Using Contracts*. PhD thesis, EECS Department, University of California, Berkeley, August 2015. URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2015/EECS-2015-189.html>.

- 36 P.W. O’Hearn. Resources, concurrency, and local reasoning. *Theor. Comput. Sci.*, 375(1-3):271–307, 2007. doi:10.1016/j.tcs.2006.12.035.
- 37 Benjamin C. Pierce. *Types and Programming Languages*. MIT Press, 2002.
- 38 A. Platzer. *Logical Analysis of Hybrid Systems - Proving Theorems for Complex Dynamics*. Springer, 2010.
- 39 K. V. S. Prasad. A Calculus of Broadcasting Systems. *Sci. Comput. Program.*, 25(2-3):285–327, 1995. doi:10.1016/0167-6423(95)00017-8.
- 40 Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y Ng. ROS: an open-source Robot Operating System. In *ICRA workshop on open source software*, 2009.
- 41 W.C. Rounds and H. Song. The Phi-Calculus: A Language for Distributed Control of Reconfigurable Embedded Systems. In *HSCC*, pages 435–449. Springer, 2003.
- 42 Alceste Scalas and Nobuko Yoshida. Less is More: Multiparty Session Types Revisited. *Proc. ACM Program. Lang.*, 3(POPL):30:1–30:29, January 2019. doi:10.1145/3290343.
- 43 Kaku Takeuchi, Kohei Honda, and Makoto Kubo. An Interaction-based Language and its Typing System. In *PARLE’94*, volume 817 of *LNCS*, pages 398–413, 1994. doi:10.1007/3-540-58184-7\_118.