

MEng Individual Project Report

HIGH PERFORMANCE PARALLEL DESIGN BASED ON SESSION PROGRAMMING

Nicholas Ng (cn06@doc.ic.ac.uk)
Department of Computing
Imperial College London

Supervisor
Nobuko Yoshida (yoshida@doc.ic.ac.uk)

Second Marker
Wayne Luk (wl@doc.ic.ac.uk)

Contents

1	Introduction	11
1.1	This project	13
1.2	Contributions	13
2	Background	15
2.1	Pi-calculus	15
2.1.1	Asynchronous π -calculus	15
2.2	Session types	17
2.2.1	Syntax of session calculus	17
2.3	Multiparty session types	19
2.4	Session programming with SJ	20
2.4.1	Branching	22
2.4.2	Iteration	23
2.4.3	Delegation	24
2.4.4	Non session-based alternatives	24
2.4.5	Related work	25
2.5	Axel	25
2.5.1	Hardware arrangement	25
2.5.2	Software	26
2.5.3	Performance	28
2.6	Parallel Algorithms	28
2.6.1	N-body simulation	28
2.7	Summary	30
3	Design and Implementation	31
3.1	Design goals	31

3.2	Session Java on Axel	32
3.2.1	Overall design	32
3.2.2	Application topology and session typing	34
3.3	SJ with FPGA	36
3.3.1	The need for cross-language features	36
3.3.2	Java Native Interface	37
3.3.3	Java Native Access	37
3.3.4	Problems encountered	38
3.4	C-translation of SJ n-body implementation	39
3.4.1	Why would this work?	40
3.4.2	A SJ primitives library for C	40
3.4.3	Shortcomings of the library	42
3.5	Summary	42
4	Correctness Proof of N-body Implementation	43
4.1	Session calculus with multichannel in/outwhile	43
4.2	Syntax	44
4.3	Operational semantics	44
4.4	Type system	46
4.4.1	Duality	47
4.4.2	Typing environment	48
4.4.3	Typing rules	48
4.5	Subject reduction	52
4.5.1	Well-formed topology	53
4.5.2	Subject congruence theorem	54
4.5.3	Subject reduction theorem	56
4.6	Progress property	60
4.7	Correctness proof for n-body simulation	61
4.7.1	N-body simulation in session calculus	61
4.8	Summary	62
5	Testing and Evaluation	63
5.1	Alternative designs	63
5.1.1	SJ and acceleration hardware allocation	63
5.1.2	Communication medium	64

<i>CONTENTS</i>	5
5.2 Pre-implementation tests: inner product	65
5.2.1 JNA direct mapping	65
5.2.2 JNA interface mapping and direct mapping	66
5.2.3 Execution in CPU and FPGA	66
5.3 Benchmarks	68
5.3.1 Benchmark methodology	69
5.3.2 Benchmark results	69
5.3.3 Comparing with Axel's implementation	73
5.3.4 Benchmark results conclusion	73
5.4 Summary	74
6 Conclusion	75
6.1 Future work	76
A Appendix	79
A.1 Java Native Interface (JNI) example	79
A.2 Java Native Access (JNA) example	80
A.3 Comparison of SJ and C-translation implementation	82
A.4 SJ + FPGA speedup over SJ implementation	85
Bibliography	78

Abstract

Session programming is a programming model based on the theory of session types, a typing system for π -calculus. Session types is developed to model structured interaction between processes and correctly typed process will have the property of communication safety. Session Java (SJ) is a full implementation of session types in Java. In this project, We aim to introduce the session programming model to Axel, a heterogeneous cluster with both FPGAs and GPUs as hardware accelerators to design communication safe parallel algorithms.

We give an implementation of a parallel algorithm, n-body simulation, on the Axel cluster, using SJ and FPGAs. We also give a translation of our SJ n-body simulation into C to get a higher performance. We find good performance improvements in both implementations, without compromising safety property of our program.

Finally, we present a formalisation of two new *multichannel* SJ primitives for parallel programming. We use the formalisation to prove the correctness of our n-body implementation and generalise the proof to a ring topology used by parallel algorithms in SJ.

Acknowledgements

I would like to thank the following people, without whom this project might not be as successful. My supervisor, Dr. Nobuko Yoshida for her enthusiasm and guidance throughout the project, my second marker, Prof. Wayne Luk for his constructive advices and feedback on the project, Andi Bejleri for the crash course on session types, Brittle Tsoi for his help on FPGA and the Axel cluster, Olivier Pernet for his insights and advices to the direction of the project, Raymond Hu for advices and support on SJ, Wilhelm Kleiminger for moep and proofreading my final report, and finally, my family for their support and love throughout my four years at university.

Chapter 1

Introduction

In 1965, Gordon Moore predicted that the number of transistors on a chip doubles about every two years [?]. 45 years later, Moore's Law remained valid and is generalised to describe the performance growth of microprocessors. Until recent years, microprocessor manufacturers have enjoyed performance increase simply by cramming more transistors on a single microprocessor. As the cost of performance rose to an unfeasible level due to power consumption, to keep up with Moore's Law, research and development on computer architecture turned towards parallelising techniques on existing hardware. Multicore processor architecture rose in popularity, today it is easy to find dual-, quad-, even hexa-cores on a single processor dice, and we are expecting to see processors with as much as 80-cores in the next few years [?].

On a much bigger scale, another type of parallel architecture is *computer clusters*. It is a form of distributed computing, where multiple standalone cluster nodes are connected and computation jobs are shared between the nodes. Each node can work on their partition of jobs in parallel. From outside the cluster, the jobs submitted to the cluster are completed as if a single computer is used. As a cluster can be built using commodity hardware, it is a cost-effective way of building supercomputers.

Parallel models of computing has shown it self as a promising direction to higher performance computer architecture.

Another trend that saw a lot of interest lately is the use of *hybrid architectures* to achieve high performance. Instead of using a centralised computation model based on the CPU, parts of computations are delegated to other specially designed hardware which can perform the computation more efficiently.

Field Programmable Gate Array (FPGA) is a type of reconfigurable integrated circuit. FPGAs can be configured to represent software instructions directly in hardware during runtime. On CPUs, all instructions need to go through the fetch-decode-execute cycle before they can be executed. Implementations in FPGAs do not require the fetch-decode phase since the instructions are already represented in the hardware circuit. As a result, FPGAs are much more efficient than CPUs in computation-heavy tasks. Also because the computations are done in hardware circuit, pipelined instructions can be executed in parallel as a physical property of electric circuits.

A modern GPU can double as a many-core general purpose processor. A GPU has hundreds of processing cores which are very capable at floating point computation, as they are usually used for graphics calculations. *Common Unified Device Architecture* (CUDA) and ATI's *Stream* are software frameworks that allow the use of the GPU cores for non-graphical computations. This is known as *General Purpose computation on Graphics Processing Units* (GPGPU). Compared to traditional CPUs where support for *Single Instruction Multiple Data* (SIMD) is limited to the *Streaming SIMD Extensions* (SSE) instruction set which can work on at most four single precision floating point number in parallel, the GPU can work on hundreds of data in parallel in the GPU cores.

The performance edge in parallel architectures do not come without its own problems. Parallel programming is a much less understood model than traditional serial programming model. Some of the problems were solved by the implicit programming model, where programmers do not need to understand parallel programming and parallelisation is done implicitly in hardware or by compiler optimisations. The advantage of this technique is a guarantee of a certain degree of parallelism and the correctness of the parallelised section since most of the optimisations are rather conservative. On the other hand, this approach might not always give the optimal result if there exists specific ways to parallelise the code as the programmer do not have control over the implicit optimisations. The alternative, explicit parallelism comes typically in form of message-passing. Often, small trivial mistakes in the program will result in parallel synchronisation issues, race conditions, or deadlocks. Combined with the interleaving of executions, parallel programming in this model cannot guarantee communication safety and is difficult to identify problems as the execution sequence can be undeterministic. For years, computer science theorists seek to understand parallel programming model and search for solutions by formalising and modelling concurrent processes. Amongst the more active researches, the Actor model and process algebras, such as Calculus of Communicating Systems CCS [?] and its successor π -calculus [?, ?], are fields that found most success. With a model of processes interactions, it is much easier to understand the properties of concurrent systems to prevent the issues common in parallel programming.

Session types [?, ?, ?] were developed as a typing system for π -calculus. Interactions between parties are conducted over private channels called sessions. A session type specifies the sequence and typing of interactions in the session. Session types captures the fact that communicating parties must have a compatible typing between them. For example, if a sender intends to send an integer, the receiver must be expecting to receive an integer as well; otherwise there would be problem. By analysing the session typing of communicating processes and making sure only compatible processes can start a session, we are able to show that *communication safety property* holds for session-typed processes - deadlocks are not possible. If a parallel system is modelled in π -calculus is shown to be type-safe by session types, we are confident to say that the system is communication safe.

Session types as a typing system alone cannot be used directly as a design tool. *Session Java* (SJ) [?] is an implementation of session types as an extension of Java. It is designed to be a non-intrusive addition to Java and integrates well with the object oriented setting. SJ brought the full theoretical session programming framework to a programming language in common use. It is a powerful tool for programmers that can ensure session compatibility within the program-

ming language, without first modelling interactions in the pure theoretical framework, to create communication safe code.

1.1 This project

This project aims to explore ways of applying session programming to heterogeneous clusters, which uses acceleration hardware such as FPGAs or GPUs.

We wish to demonstrate uses of session-type based Java to design parallel algorithms which are communication safe, efficient and easily readable on our target platform. We also aim to generalise and extend our approach to other similar platforms so designs on these platforms can also make use of session programming.

1.2 Contributions

In this report we document our findings and results on applying session programming to design parallel algorithms on heterogeneous clusters. We made the following contributions:

- Introduced an architecture for high performance parallel application design with *session Java* (SJ) using heterogeneous hardware (§3.2).
- Implemented a parallel n-body simulation in SJ accelerated by FPGA on a heterogeneous cluster (§3.3), with full sets of benchmark results to compare the performance with and without acceleration hardware (§5.3). Our implementation using SJ and FPGA yields up to 2 times speedup in the best performance.
- Implemented a C library that can be used for session programming in the C programming language (§3.4.2), and a C implementation of n-body simulation translated from SJ to demonstrate the use of the library (§3.4). The translated code has on average 5 times speedup over SJ implementation.
- Presented a formalisation and first correctness proof of a pair of new multichannel SJ primitives - **inwhile** and **outwhile** in SJ, designed to implement parallel algorithms. (§4.1) The two primitives can represent parallel topologies more naturally and were shown to be more efficient than its single-channel counterpart.
- Proved our implementation of n-body simulation deadlock free, using the formalisation of the new multichannel SJ primitives (§4.7).

Report organisation

The report contains six chapters and each chapter is organised as follows:

- Chapter 2: Background will cover the background theories behind session types and session Java. We will also introduce our target platform, a heterogeneous cluster called Axel which contains FPGAs and GPUs as processing elements.
- Chapter 3: Design and Implementation will give the design and implementation details of the main result of the project, an implementation of n-body simulation using SJ and FPGA. We will also include a version of the implementation translated from SJ to C, which is more suitable for deployment on high performance clusters.
- Chapter 4: Correctness proof of n-body implementation will detail an extension to the session type introduced in [?] to include multichannel `inwhile` and `outwhile` SJ primitives used for parallel programming in SJ. We will then show a correctness proof of our n-body implementation based on the updated session type.
- Chapter 5: Testing and Evaluation will discuss and evaluate alternative designs and compare benchmark results of different implementations of n-body simulation.
- Chapter 6: Conclusion will conclude our findings of the project and outline potential future works.

Chapter 2

Background

In this chapter we will discuss fundamental background theories of which session programming is based on. This includes π -calculus (§2.1) - the process calculi which is the modelling basis for session programming, session types (§2.2) - the typing system of π -calculus for sessions-based communication and an introduction to session Java (§2.4) which is the main programming tool we are going to use in the project.

Next, we will introduce the target platform for the project (§2.5) - Axel, a heterogeneous cluster with *Field Programmable Gate Arrays* (FPGAs) and *Graphics Processing Units* (GPUs) as acceleration hardware.

Finally, we will briefly look at the current parallel programming model of Axel and session Java to implement our choice of parallel algorithm - n-body simulation (§2.6).

2.1 Pi-calculus

π -calculus [?] is a process calculus proposed by Milner, Parrow and Walker as a successor to *Calculus of Communicating Systems* (CCS) as a model to study concurrent mobile systems. It uses message passing and is distinguished from CCS by its use of names in messages rather than values in CSS. The difference between value passing and name passing is that only name passing allows sending and receiving of *channel names* as messages, so channels can be ‘reconfigured’ at run time. This makes π -calculus more expressive and suited for mobile processes [?]. There are many variants of π -calculus for different applications, due to the notational differences in different domains. π -calculus and its variants lie a foundation for modelling communication systems, from simple asynchronous π -calculus and spi calculus for session types and cryptography [?] respectively, to more advanced calculi such as 3π for developmental and systems biology [?].

2.1.1 Asynchronous π -calculus

Asynchronous π -calculus is the simplest variant of π -calculus and is the variant which session type (§2.2) is based on.

$P, Q ::=$	processes
$\mathbf{0}$	nil process
$P \mid Q$	parallel composition of P and Q
$(\nu a)P$	generation of a with scope P
$!P$	replication of P
$\bar{a}\langle v \rangle$	output of v on channel a
$a(x).P$	input of <i>distinct</i> variables x on a , with continuation P

Fig. 2.1: Syntax of asynchronous π -calculus

Syntax

The main difference between full π -calculus and asynchronous π -calculus is asynchronicity. This means after an output action, there is no continuation and the process terminates when the message is delivered. Communication with asynchronous π -calculus is therefore deadlock free. Asynchronous communication can be used to simulate synchronous communication, and is common in distributed systems. Fig. 2.1 shows the syntax of asynchronous π -calculus.

Reduction rules

$$\begin{array}{r}
 \bar{a}\langle v \rangle \mid a(x).P \rightarrow P\{v/x\} \quad \text{[COM]} \\
 \frac{P \rightarrow P'}{P \mid Q \rightarrow P' \mid Q} \quad \text{[PAR]} \\
 \frac{P \rightarrow P'}{(\nu a)P \rightarrow (\nu a)P'} \quad \text{[RES]} \\
 \frac{P \equiv Q \rightarrow Q' \equiv P'}{P \rightarrow P'} \quad \text{[STRUCT]}
 \end{array}$$

Fig. 2.2: Base reduction rules of asynchronous π -calculus

Reduction is the method for processes in π -calculus to interact. If reduction is not possible, no action can be performed.

It should be noted that in π -calculus interactions between processes are initiated by parallel composition; with processes alone, π -calculus and the reduction rules are meaningless.

$P ::= \text{request } a(k) \text{ in } P$	session request
$ \text{accept } a(k) \text{ in } P$	session acceptance
$ k![\tilde{e}]; P$	data sending
$ k?(\tilde{x}) \text{ in } P$	data reception
$ k \triangleleft l; P$	label selection
$ k \triangleright \{l_1 : P_1 \parallel \dots \parallel l_n : P_n\}$	label branching
$ \text{throw } k[k']; P$	channel sending
$ \text{catch } k(k') \text{ in } P$	channel reception
$ \text{if } e \text{ then } P \text{ else } Q$	conditional branch
$ P \mid Q$	parallel composition
$ 0$	inaction
$ (\nu u) P$	name/channel hiding
$ \text{def } D \text{ in } P$	recursion
$ X[\tilde{e}\tilde{k}]$	process variables
$e ::= c$	constant
$ e + e' \mid e - e' \mid e \times e \mid \text{not}(e) \mid \dots$	operators
$D ::= X_1(\tilde{x}_1\tilde{k}_1) = P_1 \text{ and } \dots \text{ and } X_n(\tilde{x}_n\tilde{k}_n) = P_n$	declaration for recursion

Fig. 2.3: Session calculus syntax from [?]

2.2 Session types

A session is a predefined sequence of exchanging messages otherwise known as a protocol. Session types were developed as a typing system of the π -calculus for use by communication-based concurrent programming languages with basic communication constructs. The theory of session types is defined in terms of session calculus based on asynchronous π -calculus, originally introduced by Honda et al. [?] and their work was subsequently revised by Yoshida and Vasconcelos [?] which became the basis of session Java. Session calculus is a building block of session types, where session type defines compatible sessions in terms of session calculus.

2.2.1 Syntax of session calculus

Basic constructs Processes exchange messages by pairs of send and receive actions. Data sending and receiving are the basic constructs in session type, along with inaction process, parallel composition of processes and name restriction (ν). These basic constructs can be directly translated to asynchronous π -calculus.

$$\begin{aligned}
&\text{accept } a(k) \text{ in } P = \text{request } a(k) \text{ in } P \\
&\quad k![\tilde{x}]; P = k?(x) \text{ in } P \\
&\quad k \triangleleft l; P = k \triangleright \{l_1 : P_1 \parallel \dots \parallel l_n : P_n\} \\
&\text{throw } k[k']; P = \text{catch } k(k') \text{ in } P \\
&\quad \mathbf{0} = \mathbf{0}
\end{aligned}$$
Fig. 2.4: *Dual actions*

Label branching Label branching is a feature in session calculus for *structured external choice*. Without label branching, it is not possible for sessions to exhibit different behaviour on different conditions or they will be incompatible. Thus the use of session calculus and session types will be very limited and only useful in serial and simple communications, if choices are not possible.

Label branching is done by sending and receiving a *label*, and based on the content of the label, the session type following the selection action can be different as long as the session type of the counterpart participant remains compatible with the receiver after sending the said label.

Session delegation Because of the name-passing property of π -calculus, it is possible to pass more information and use the sessions more flexibly. This allows sessions to be passed to other processes as a parameter via a channel, subprocesses can use the received session as a session rather than a value. By offloading parts of responsibilities of the parent process to the subprocesses, we can distribute processing to lower level or smaller processes, without sacrificing the advantages of using session-types because all delegated processes also follow a subset of session-type from the top level. Most importantly, the top-level process does not need to be informed about the delegation which allows a higher level view when designing distributed systems with session types.

Duality We mentioned the importance of interaction in the previous section (§2.1.1). A correct session type for two interacting process requires the sessions in the same channel to be ‘associated with complementary behaviours’ [?, Definition 5.2], this important requirement provides the theoretical basis for communication safe processes. In the syntax given in Fig. 2.3, complementary actions are shown in Fig. 2.4; these pairs of actions, when composed together, will reduce without getting stuck. A sound type system will not cause stuck errors if the implementation is correct.

As an example of session type, Example 1 is definition of a simple sum server system that adds and returns the sum of two numbers in *SumServer* supplied by *SumClient*. This will be used in subsequent sections to demonstrate the similarities and difference between session type and its derivatives.

$G ::=$	Global types
$p \rightarrow p' : \langle U \rangle . G$	Message
$p \rightarrow p' : \{l_k : G_k\}_{k \in K}$	Branching
$\mu x . G$	Recursion
x	Type variable
$G \ i$	Application
end	Null

Fig. 2.5: Global types and type reduction from [?]

Example 1.

$$\begin{aligned}
SumServer &= \text{accept } a(k) \text{ in } k?(\tilde{x}) \text{ in } k?(\tilde{y}) \text{ in } k![x+y]; \mathbf{0} \\
SumClient &= (\nu k) \text{ request } a(k) \text{ in } k![42]; k![77]; k?(result) \text{ in } \mathbf{0} \\
SummingSystem &= (\nu a) SumClient \mid SumServer
\end{aligned}$$

2.3 Multiparty session types

Session types introduced in §2.2 describes communication between two parties. When a communication involves more than two participants, the communication can be modelled by multiple binary sessions between any two of the participants. All communications between any two participants can be guaranteed compatible and error free by the safety property of binary session type. However, binary sessions cannot prevent interleaving of sessions in a communication with multiple binary sessions. Interleaving sessions might allow incorrect communication logic or cause problems such as deadlocks because the execution sequence is not determined at design time.

This will be a problem when using binary session types in a system where different participants of communication are implemented by different parties, and each party is given a protocol specification which they code according to. The final product will be correct in the local view but could be incorrect in the global view because the parties do not have the information of the global session type. Therefore to design a correct multiparty protocol, no assumption of execution order should be made and the communication of different participants should be specified explicitly.

The basic constructs of a multiparty session type are almost identical to binary session type, but add a global type (Fig. 2.5) on top of the endpoint (or local) session type. The global type specifies the global progress of the communication, and *projects* to endpoint session type for each of the participants.

Global session types provide a theoretical basis to prove that a communication is correct in the global view, where the participants and order of communication are defined explicitly in the design. The projection of global type to local type will ensure such properties are preserved after the operation.

2.4 Session programming with SJ

Session types are a basis for session-based programming, but do not describe a standalone programming language. As a result, session types and the object-oriented programming language Java are combined to create *Session Java* (SJ) [?, ?]¹. SJ is an extension of Java and thus the syntax of SJ is identical to Java, with extra primitives for session programming. This is best illustrated with a simple example:

```

1  public class Server {
2      // Session declaration
3      final noalias protocol p_server {
4          sbegin.?(int).?(int).!<int>
5      }
6      public run(int port) {
7          final noalias SJServerSocket svr;
8          final noalias SJSocket sock;
9          try (svr) {
10             svr = SJServerSocketImpl.create(p_server, port);
11             try (sock) {
12                 sock = svr.accept();
13                 int x = sock.receiveInt(); int y = sock.receiveInt();
14                 sock.send(x + y);
15             } catch ( ... ) {}
16         } catch ( ... ) {}
17     }
18 }
19
20 public class Client {
21     // Session declaration
22     final noalias protocol p_client {
23         cbegin.!<int>.!<int>.?<int>
24     }
25     public run(String host, int port) {
26         final noalias SJService svc
27             = SJService.create(p_client, host, port);
28         final noalias SJSocket sock;
29         try (sock) {
30             sock = svc.request();
31             sock.send(42); sock.send(77);
32             int result = sock.receiveInt();
33         } catch ( ... ) {}
34     }
35 }

```

Listing 2.1: *SJ sum server/client*

¹Currently, Session Java is only an implementation of *binary* session type and extension for *multiparty* session type is planned.

We now look at the *SumServer/SumClient* example again to show a basic communication system with SJ. A session calculus version was given in Example 1 in the previous section.

1. A client sends two numbers to the server
2. The server replies with the sum of the two numbers received

The communication primitives of SJ are similar to conventional socket programming (`request`, `accept`, `send`, `receive`), except for the `protocol` code block. The `protocol` defines the *session typing* of the program, introduced in the previous section (§2.2). With the session type of the program defined, it is possible to [?]:

1. Ensure the implementation conforms to the specified `protocol` by static checks at compile time.
2. Check that the two communicating programs are compatible, by a duality check of the `protocol` at the start of communication.
3. Simplify checking protocol correctness by abstracting away implementation details and checking only the session type.

Table 2.1 shows the relationship between protocol and Java code, which allows the communication safety checks mentioned above.

Protocol	SJ code		Line
<code>sbegin</code>	<code>accept()</code>	Starts a server session	9
<code>cbegin</code>	<code>request()</code>	Starts a client session	23
<code>!<datatype></code>	<code>send()</code>	Sends an object with <i>datatype</i>	24
<code>? (datatype)</code>	<code>receive()</code> ²	Receives an object with <i>datatype</i>	10

Table 2.1: Session type and the corresponding Java code

Below we list three scenarios in the *SumServer* example that can demonstrate the benefits of the safety checks:

The protocol is correct but the implementation does not conform to the protocol If the implementation of *SumServer* receives three integers instead of two as stated in the protocol, the SJ compiler will throw an exception.

The protocol and implementation are both correct but the protocols are not dual of each other If *SumServer* replies with the result at the end of the execution, and *SumClient* is not receiving the final result, ie. Server: `sbegin.?(int).?(int).!<int>` and Client: `cbegin.!<int>.>.!<int>`. When the connection between the two processes is established, an incompatible session exception is thrown.

²`receiveInt()` is a shortcut to receive an `int`

A logic error exists in the protocol design If *SumServer* sends a result before receiving any values, and *SumClient* is also compatible with the server, ie. Server: `sbegin.!<int>.?(int).?(int)` and Client: `cbegin.?(int).!<int>.!<int>`. The same applies to analysing problems with distributed deadlocks, where the processing of the values is less important than the main source of problem - communication primitives. With SJ the developer can reason about the problem in the protocol level, eg. “Because the result is sent before the numbers are received and processed, therefore changing the arguments to *SumServer* do not influence the result” without necessarily understanding the operation done on the two arguments. Also because of the conformance check, we can be assured that the code implements the protocol *without* looking at the code to find the problem.

In session types and π -calculus, a new channel is ‘generated’ or ‘restricted’ by using the operator ν *channelname*. In session programming, this corresponds to the action of creating a new Socket. The socket will contain transmission between the participants of communication once it is created. The send and receive object methods of the socket, and in session calculus you can only input and output on a given channel.

2.4.1 Branching

When programming conditional statements, often different choices will branch to different behaviours of the program. To model that, label branching in session type is used:

Protocol	SJ code	
<pre>!{ LABEL₀:session₀, LABEL₁:session₁, ... }</pre>	<pre>sock.outbranch(LABEL₀) { code₀ } sock.outbranch(LABEL₁) { code₁ } ...</pre>	<pre>Send LABEL₀ Send LABEL₁</pre>
<pre>?{ LABEL₀:session₀, LABEL₁:session₁, ... }</pre>	<pre>sock.inbranch { case LABEL₀: { code₀ } case LABEL₁: { code₁ } ... }</pre>	<pre>Receive LABEL₀ Receive LABEL₁</pre>

Table 2.2: *Branching in session programming*

Table 2.2 shows the receiving and sending of labels. The code blocks `?{}` and `!{}` represent receive label and send label respectively. Sending of labels is usually used in conjunction with conditional statements, for example in Listing 2.2.

```

1 final noalias protocol p_hwc {
2     cbegin.!{LOWER:?(int), UPPER:?(String)}
3 }
4 ...
5
6 if (userInput.equalsIgnoreCase("lower")) {
7     sock.outbranch(LOWER) {
```

```

8     System.out.println("LOWER branch; Server replies with #"
9     + sock.receiveInt());
10    }
11 } else {
12     sock.outbranch(UPPER) {
13         System.out.println("UPPER branch; Server replies with "
14         + (String) sock.receive());
15     }
16 }

```

Listing 2.2: Example usage of label sending

2.4.2 Iteration

Iteration in session programming translates to replication in π -calculus. In π -calculus processes can be repeated and this forms the loop-body of an iteration. Iteration is not part of the session calculi defined in [?] but will be formalised in this report. By using an explicit looping construct that is similar to normal Java programming (**outwhile/inwhile** vs. **while**), the reasoning of iteration is thus simpler to the programmer. Table 2.3 shows the syntax of **outwhile** and **inwhile**. The only difference between the two is **outwhile** controls the looping condition, and **inwhile** reacts passively. The iteration construct therefore also work as a synchronisation mechanism between the sessions. To implement iteration with the same semantic in MPI, the single line code might be expanded to ³:

```

1 // outwhile(condition)
2
3
4 while (condition) {
5     MPI.COMM_WORLD.Barrier();
6     /* outwhile code */
7     MPI.COMM_WORLD.Send(
8         condition, ...);
9 }

```

Listing 2.3: **outwhile** in MPJ Express

```

1 // inwhile()
2 MPI.COMM_WORLD.Recv(
3     condition, ... );
4 while (condition) {
5     MPI.COMM_WORLD.Barrier();
6     /* inwhile code */
7     MPI.COMM_WORLD.Recv(
8         condition, ...);
9 }

```

Listing 2.4: **inwhile** in MPJ Express

Protocol	SJ code
! [<i>session in iteration</i>]*	s1.outwhile(<i>condition</i>){ ... } s1.outwhile(s2.inwhile; <i>condition</i>){ ... }
? [<i>session in iteration</i>]*	sock.inwhile(){ ... }

Table 2.3: Iteration in session programming

³Example uses syntax of MPJ Express, a Java MPI implementation

Note that the alternate form of `outwhile` that uses `inwhile` as condition is a new SJ primitive for implementing parallel algorithm. The formalisation and proofs will be given in §4.1.

2.4.3 Delegation

Sessions can be delegated to other components, to expression delegation of session in the protocol, we simply replace the type of the message by a session, as shown in Table 2.4. *delegated session* in the table represents the session type of the *initialised SJSocket* and of `rcvdSession`. If we look closely the primitives of session delegation is identical to ordinary `send` and `receive` in SJ, except the content is a session rather than a usual data type (but lends itself to the Java Object model where every type is a subtype of `Object`). Session delegation is an important tool to distribute tasks.

Protocol	SJ code	
! <i><session></i>	<code>sock.send(delegateSJSocket)</code>	Send a session
?(<i>session</i>)	<code>SJSocket session = sock.receive()</code>	Receive a session

Table 2.4: Session delegation in session programming

2.4.4 Non session-based alternatives

The implementation of SJ is most similar to that of the MPI standard (§2.4.4) in terms of communication model (message passing) and design. There are also other distributed message passing system such as Java *Remote Method Invocation* RMI, but the design and uses are in a different domain compared to SJ.

Message-Passing Interface (MPI)

MPI is a message-passing library interface specification [?] and is commonly used in the high performance computing field for message-passing based parallelism.

Using MPJ Express, an implementation of the MPI standard in Java, it was shown in [?] that there are many similarities between the two but the main differences are:

MPI has more features SJ does not have multicast-type message send primitive, but theory for multiparty session type §2.3 have been developed for future implementation [?] in SJ.

MPI is a low level protocol which makes it prone to communication mismatch or deadlocks due to explicit message passing [?]. A communication mismatch in MPI such as a `MPI_Send` without a corresponding `MPI_Recv`, will not cause problem until some point in the execution. Scenario 2 of safety check examples above will will cause MPI but not with SJ. In a distributed database system,

it would require a rollback on all previous calculations. SJ's safety properties (§2.4) will prevent the incompatible sessions from starting.

SJ has high level session abstraction so the code is more structured and more readable than MPI, this gives the programmer an advantage to focus on more important communication/protocol details.

SJ is not an external library SJ was designed to be a full object-oriented programming language. Implementations of MPI are external libraries since it is only a communications standard. As a domain-specific language, syntax for tasks common to communications programming can be built into the syntax and will be more natural to use, despite the small difference between Java and SJ. Examples include the try-channel syntax to catch exceptions from within a specific channel (line 8 in Listing 2.1), and the different forms of `outwhile/inwhile` as a session-type specific looping technique.

Taking the example of `outwhile` and `inwhile` in Listing 2.3 and 2.4 again, the iteration feature in MPI is less readable than in SJ because the special iteration syntax is not found in MPI.

2.4.5 Related work

Implementation of session types had been developed for other languages such as Haskell [?, ?]. Other work on session-type with C-like languages [?, ?] does not take the direction of implementing the full session type system. SJ is the first practical session-type based object-oriented programming language.

2.5 Axel

Axel [?] is a heterogeneous computer cluster built at Imperial College. The cluster consists of 18 computing nodes, and each of the nodes contains a x86 CPU, a number of *Graphics Processing Units* (GPU) and most of the nodes contain a *Field Programmable Gate Array* (FPGA) device. FPGAs and GPUs are used on Axel as hardware accelerating components.

Axel is the target platform for this project. We wish to deploy SJ on the cluster and use session programming to improve parallel design. Below is an overview of hardware and software currently on the cluster.

2.5.1 Hardware arrangement

NNUS clusters and UNNS clusters There are two ways of grouping hardware accelerators (or *Processing Elements*, PE) in a heterogeneous cluster, namely *Nonuniform Node Uniform Systems* (NNUS) and *Uniform Node Nonuniform Systems* (UNNS).

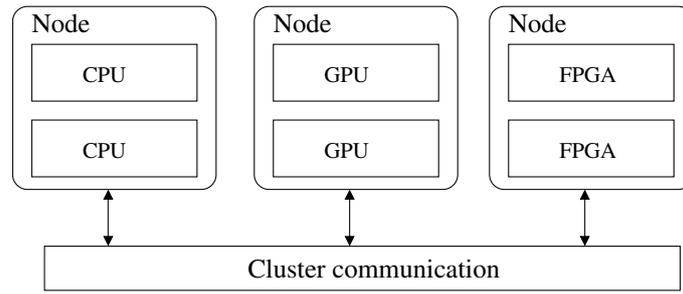


Fig. 2.6: Nodes in a generic UNNS cluster

In a UNNS cluster, each node of the cluster hosts a single type of PEs. In the example of Fig. 2.6, three nodes of the cluster hosts CPUs, GPUs and FPGAs respectively. For nodes that hosts non-CPU PEs, special hardware are needed to control the nodes because they cannot run ordinary operating systems. Examples of UNNS clusters are SRC-7 MapStation and RASC server from SGI [?].

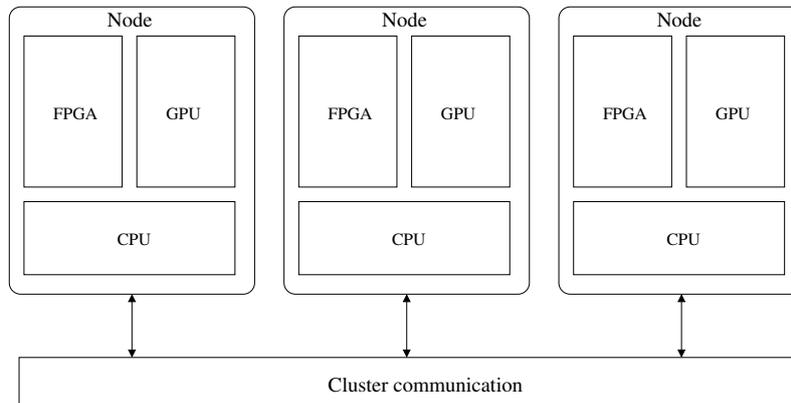


Fig. 2.7: Nodes in a generic NNUS cluster

On the other hand, in a NNUS cluster, each node of the cluster contains different PEs (thus *Nonuniform Node*). The PEs of the NNUS cluster example in Fig. 2.7 shown to be on the same node are CPU, FPGA and GPU. All nodes in the cluster have the same arrangement. This makes it easy to put together commodity hardware and build a cluster (eg. Beowulf clusters)

Axel is a **NNUS** heterogeneous cluster, meaning that each node in the cluster will contain different types of PEs. Each node on Axel can be used as an independent x86 PC equipped with hardware accelerators (FPGA board and GPUs). The details are shown in Fig. 2.8.

2.5.2 Software

All the nodes in the cluster run a standard Ubuntu Linux (amd64 architecture). The following software and frameworks are installed to program different hardware components of Axel:

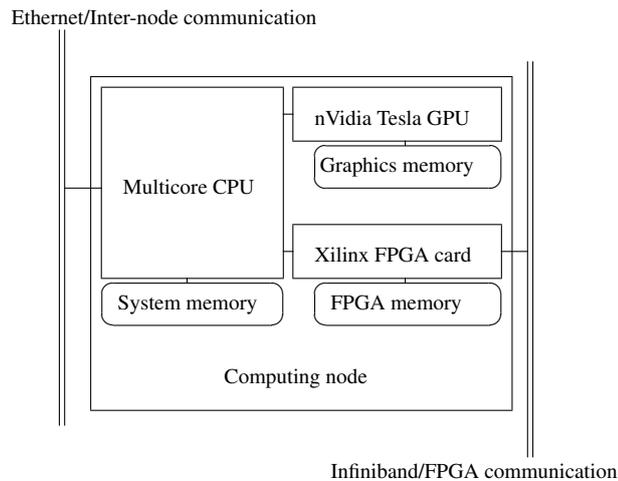


Fig. 2.8: *Axel's NNUS arrangement*

CPU The CPUs are standard multicore x86 CPUs and GCC is used with OpenMPI⁴ to produce executables to run on the CPU. The main use of the CPU in a complete Axel application is to coordinate communication between computing nodes using MPI, but it can also be used for general CPU based computation.

nVidia GPU All the GPU used are nVidia Tesla cards, designed for high performance computing rather than general graphics rendering. nVidia provides the *Common Unified Device Architecture* (CUDA) framework for GPU programming. CUDA is the standard *General-Purpose computing on Graphics Processing Unit* (GPGPU) framework for nVidia products and provides a C-like environment for the Tesla GPU platform [?].

FPGA Xilinx ISE 10.1 is used for development of hardware logic for FPGA hardware compilation to the FPGA devices, and all the runtime access to the FPGA devices are done in a very low-level memory mapped I/O and DMA, via a vendor supplied library, exposing an API in a C programming environment.

The compilation and execution of an Axel application is not done in a single executable. The application consists of a CPU-part that initialises the data and distributes to the GPU-part and the FPGA-part of the application. The workload split is described in an XML file for maximum flexibility, which is first read by the CPU-part to segment the data. It is possible to setup the application such that the CPU will do part of the calculation but it is usually used exclusively for I/O coordination and inter-component/inter-node communication. A **map-reduce** framework is used in the cluster for parallel programming where segments of calculations are 'mapped' to the computing elements (eg. FPGA, GPU cores) and the results are 'reduced' and collected by the I/O handling code running on the CPU which, in turn, 'reduces' the results to the master node which started off 'mapping' the input on each computing node.

⁴An implementation of MPI in C, see §2.4.4 for details of MPI

2.5.3 Performance

CPUs are designed for general purpose computation, and because of the underlying von Neuman architecture, the CPU works like an instruction interpreter and follows a fetch-decode-execute cycle to execute stored instructions. It is thus very flexible, but sacrifices the performance in a computation heavy program, where most of the execution time was spent on fetch and decode instead of the more important execute step.

Until recently, GPUs are used solely as graphics rendering hardware. With the increasing demand of modern gaming software and high throughput graphics calculations on display cards, the graphics manufacturers moved from using fixed numbers of dedicated vertex and pixel shaders to unified shaders. This allows a dynamic allocation of graphics hardware for different purposes and higher utilisation of the graphics hardware. GPU becoming less specific to graphics processing gave rise to GPGPU, where GPUs can be used for non-graphics calculations in a limited way. The advantage of GPGPU is the number of graphics processing **cores** available. These cores can do only a small number of tasks at one time but with multiple cores processing is done in parallel. Typical gaming GPUs come with a few hundred processing cores.

FPGAs essentially allow hardware execution of computer programs. As there is no need to interpret or decode the program code and can be executed directly in hardware, compared to the interpreted code model of CPU, there is less wastage of resources. Most important of all, FPGA can be easily reconfigured to carry other tasks unlike immutable dedicated hardware accelerators.

With the reasons stated above, GPUs and FPGAs are much more suited to the use cases of parallel high performance computing. Results of benchmarks [?] showed that utilising all the components gives a much better performance than using the components individually, where the magnitude of acceleration is in the descending order of *FPGA*, *GPU* and *CPU*.

On the GPU-only version and the multithreaded CPU version, the execution time of an n-body simulation of 81902 particles in a single time step is 1.1 times and 3.5 times slower than 10-core FPGA version respectively. In the heterogeneous version which uses both GPU and FPGA in a processing node, the workload is load balanced by assigning 2/3 of workload to FPGA and 1/3 to GPU, this gives the overall speedup of 2.1 times over the FPGA-only version.

2.6 Parallel Algorithms

2.6.1 N-body simulation

N-Body simulations are systems to simulate particle movement and interaction due to gravitational forces action on each other.

Each particle in the n-body simulation has a position, vector velocity and mass. In each time step, the positions and velocity of the vectors are recalculated using the velocity and acceleration. The base algorithm is shown in the algorithm below (from [?]).

```

1  for (i = 0; i < N; i++) {
2      for (j = 0; j < N; j++) {
```

```

3     if (j != i) {
4         rx = p[j].x - p[i].x;
5         ry = p[j].y - p[i].y;
6         rz = p[j].z - p[i].z;
7
8         dd = rx*rx + ry*ry + rz*rz + EPS;
9         d = 1 / sqrtf(dd * dd * dd);
10
11        s = p[j].m * d;
12
13        a[i].x += rx * s;
14        a[i].y += ry * s;
15        a[i].z += rz * s;
16    }
17 }
18 }

```

Listing 2.5: *An algorithm for n-body simulation*

N-body algorithm is highly parallelisable because every particle in the algorithm does not interfere with the content of other particles during the calculation in a time step. The calculation for each article can be done individually in parallel by different computing components.

Current Implementation on Axel On Axel, the implementation is a straightforward translation of the algorithm, by looping over each particles. The *i*-loop is split evenly to different nodes and the *j*-loop is partitioned for FPGA and GPU to compute. The results are then distributed using the `MPI_AllToAll` function to other nodes.

Implementation in SJ In SJ, due to the lack of a multicast-type send, the algorithm cannot be implemented directly. Instead, the n-body algorithm is implemented in 3 parts using a ring topology [?]. The `Master` process forwards initial data to a number of `Worker` processes, which is chained together and the last `Worker` process connects to the `Master` process to complete the ring. Since the session type of the first and last worker that communicates directly to the `Master` process, the last `Worker` is slightly different from other `Workers`, we need to further distinguish them from other worker components. The `Worker` processes will carry out most of the computation.

In each iteration the data is forwarded to each `Worker` through the chain, and adds the results of previous iteration to the data set to be forwarded to the next `Worker` and continues until all nodes have received set of particles from other nodes once. When the data is seen by all `Workers`, the positions of each particles are updated using the overall velocities and acceleration acting on each particles.

2.7 Summary

In this chapter we have introduced a theoretical framework for modelling structured communications in concurrent systems. Session calculus, a process calculi based on the asynchronous π -calculus and its typing system - session types, forms the basis of a session programming. Session typing ensures that only compatible processes can establish a session and guarantees communication safety.

We then described a full implementation of session-based programming language *Session Java* (SJ), combining Java with sessions.

Next we detailed the target platform for our project, Axel, a cluster with FPGAs and GPUs as acceleration hardware. We also included a comparison of performance between the different components in existing implementations to see a general picture.

Finally, we finished the backgrounds by looking at the differences in parallel programming model of the existing n-body algorithm implemented in Axel and SJ.

Chapter 3

Design and Implementation

In this chapter we will first look at an overall design of parallel applications with SJ on Axel (§3.2), then an implementation of the n-body simulation with SJ on the Axel cluster using CPU and FPGA (§3.3).

Next, we will present a translation of our n-body implementation in SJ to C (§3.4). We will detail the main contribution of the translation, a SJ communication primitives library for C (§3.4.2). The C translation brings session typed SJ programs to C, which is a much more suitable target language and programming environment for high performance computing.

3.1 Design goals

The aim of the project is to develop an approach for designing parallel high performance applications on heterogeneous clusters with session programming. The main criteria we considered were:

Efficiency Existing implementations of parallel algorithms on heterogeneous clusters are very efficient. We aim to keep the performance of our designs as close to current implementations as possible, while getting the advantages of session programming.

(Communication) Safety Session types were designed such that a communication between incompatible sessions will not begin and programs will behave according to its predefined protocol. It is difficult to verify correctness of complex parallel applications design, but session types and their safety property can guarantee the programs are free from incompatible interaction patterns.

Readability Session programming is a high level description of communication. Existing parallel design processes on clusters typically involves using low level libraries and working close to the metal. The instructions to develop for these libraries are verbose and require very explicit instructions (eg. MPI, §2.4.4) for simple tasks, which obfuscates the more important details of process communication. In contrast, the high level SJ abstracts most

of the implementation in the runtime system such as the transport medium (TCP vs. UDP vs. shared memory) and puts the focus on communication. Design errors can be identified easily with help of session typing §2.2 and automatic type-checking with SJ §2.4.

However, readability in session programming is a property that comes from structured code and the high level of abstraction and cannot be easily quantified by numerical metrics.

Before we commence our discussion on design and implementation of SJ applications on Axel, we should point out that Axel is an example of NNUS cluster (§2.5.1 contains details of two kinds of cluster arrangement). While our design is targeted to Axel specifically, in theory the design principals can be generalised and applied to NNUS clusters with similar architecture.

3.2 Session Java on Axel

In this section we will discuss the design and the overall architecture for SJ applications to run on Axel.

3.2.1 Overall design

Session Java is an extension of Java. Features of object oriented programming are available in SJ and allow us to create structured and easily reusable code.

The main processing elements on the Axel cluster are CPUs, GPUs and FPGAs. In the initial phase of development, we first implement our choice of algorithm in pure SJ. This allow us to identify any problems with the communication design before involvement of the new hardware, and familiarise with the facilities available on Axel.

Class organisation

By organising classes into suitable packages and class hierarchy, we reduced the efforts needed when implementing for FPGAs or any other acceleration hardware. This is possible because acceleration hardware do not participate in the flow-control of the algorithms. Typically complicated and computation heavy sections of an algorithm are isolated to a single function. The function will then be implemented on acceleration hardware, and no changes to the other parts of the program are required.

FPGA: do one thing and do it well It is uncommon to delegate multiple tasks during a single execution on one piece of acceleration hardware. Suppose two different tasks are implemented on the same piece of hardware, and the two tasks are executed in parallel. Tasks and programs are mapped to the physical hardware on FPGAs. When #1 of the two tasks are being executed, only a portion of the hardware is used - the rest of the hardware will be idle because they are designed to run task #2. Therefore we are not fully utilising the hardware every time the hardware is used.

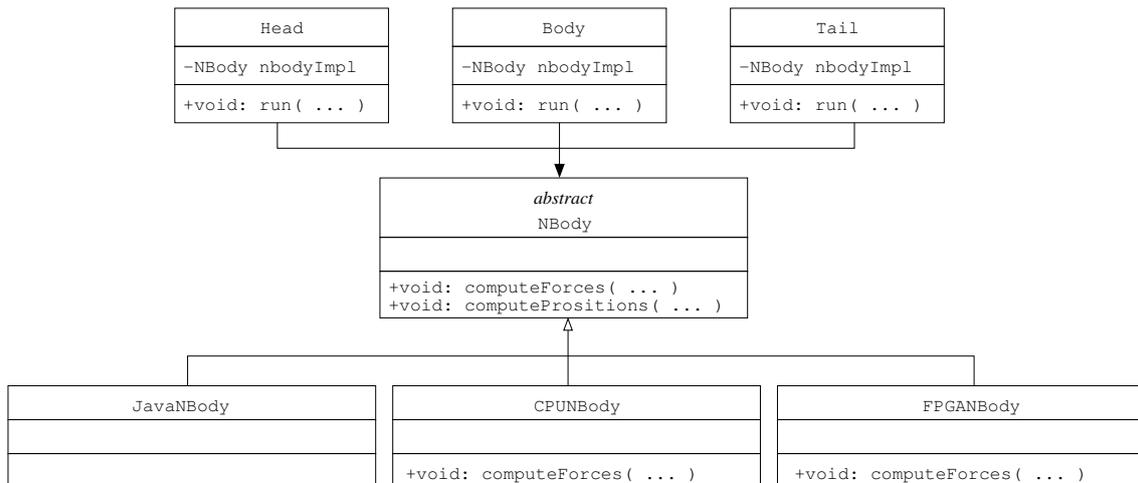


Fig. 3.1: Simplified UML diagram of our system

Despite FPGAs are known for the ability to reconfigure at runtime, it is a lengthy process and offers no practical advantage if we need to switch between tasks and reconfigure constantly.

In the class diagram shown above, `NBody` class can be replaced by the main component of other algorithms. This will allow different algorithm to use the same class structure for their implementation, such as an implementation of inner product we will detail in next chapter in §5.2.

NBody The abstract class contains all functions used by the algorithm. This class should be replaced by a similar class that contains all core functionalities when implementing other algorithms.

JavaNBody, CPUNBody, FPGANBody These classes are solid implementations of the `NBody` class. The main functions will take the input values and map them on different hardware or software implementations, then return the results as a Java array. The process is transparent to the caller.

Head, Body, Tail classes contain code to set up the topology of the application and SJ communication between nodes of the cluster.

Head is the component that runs on the first node and act as the initiator.

Body is the worker component and can be chained together with other `Body` nodes.

Tail is the last worker component that connects `Body` with `Head` on the other side of the ring. More details on the sessions of the components will be given in the next subsection. §3.2.2

All of the classes take a constructor of class `NBody` to select the implementation to use. This allows combinations such as `FPGANBody/Head`, `CPUNBody/Body` or `JavaNBody/Tail` to be constructed flexibly and easily.

This design uses the well-known design patterns of *strategy* and *template method* [?]. The two patterns help separating communication from the algorithm body, and make it easier to reuse

```

1 public abstract class NBody {
2     public abstract void computeForces( ... );
3     ...
4 }
5
6 public class JavaNBody extends NBody {
7     public void computeForces( ... ) { ... }
8     ...
9 }
10
11 public class Head {
12     NBody nbody; // set by constructor or injected
13     public void algorithmBody( ... ) {
14         ...
15         nbody.computeForces( ... );
16         ...
17     }
18 }

```

Listing 3.1: Strategy/template method pattern

the same algorithm outline and implement it in different hardware. Listing 3.1 outlines how the classes are used.

3.2.2 Application topology and session typing

The implementation of a n-body simulation follows a similar ring topology as described in §2.6.1, the only difference is the addition of initial hardware set-up and tear-down phases at the end of execution.

Session interaction

The structure we are going to describe is not limited to n-body simulations and can be adapted to any algorithm that uses a ring topology. Fig. 3.2 shows the interaction between the node and Table 3.1 gives the session declaration in each of the nodes.

We have covered the meanings of components of SJ protocols in §2.4, and we will revisit the declaration shown here in the next chapter (§4.7.1).

Partitioning of data

The input data is uniformly partitioned into n parts, where n is the number of cluster nodes in the simulation at execution time. Each node is responsible for outputting the particle positions of its allocated set of particles, and will keep track of their velocities and acceleration components at each steps of calculation.

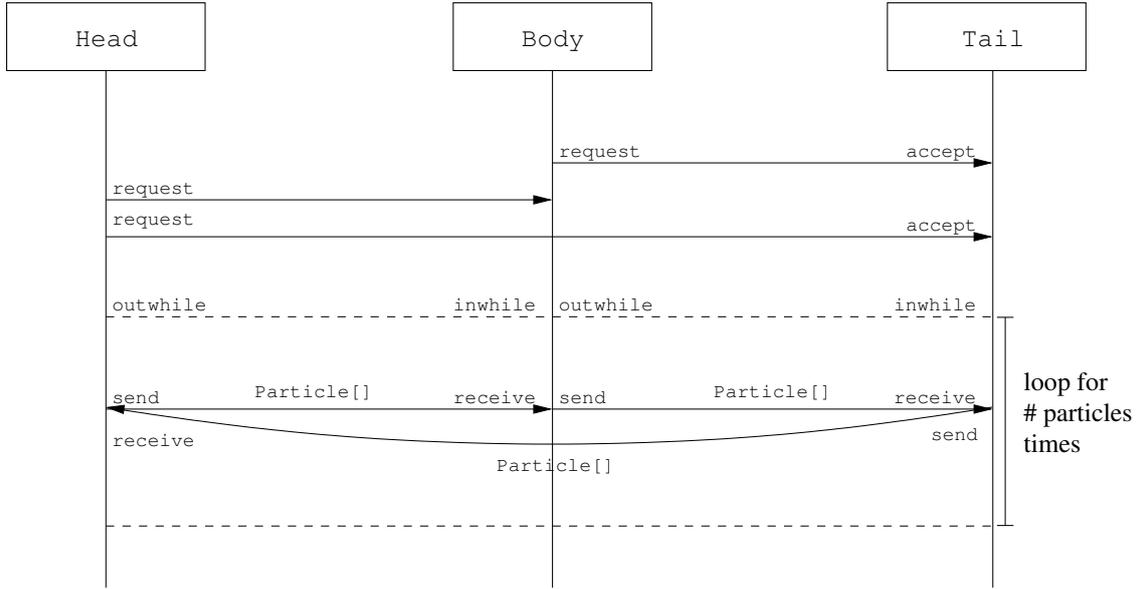


Fig. 3.2: Interaction between Head, Tail and a single Body node

Node	Session between	SJ session declaration (protocol)
Head	Head Tail	cbegin.! [?(Particle[])]*
Head	Head Body	cbegin.?(int)! [!<Particle[]>]*
Body	Body $i-1$ Body i	sbegin.!<int>.?[?(Particle[])]*
Body	Body i Body $i+1$	cbegin.?(int)! [![Particle[]]]*
Tail	Body Tail	sbegin.!<int>.?[?(Particle[])]*
Tail	Head Tail	sbegin.?[!<Particle[]>]*

Table 3.1: SJ session declaration for ring topology

The numbers in Fig. 3.3 represent the node number which the particles are from. In the *initial* round, velocities and acceleration components of each particle are calculated against each other in the same node.

Next, each node forwards the initial set (or the received set after the initial round) of particle positions to the adjacent node. Since velocities and acceleration components can be accumulated, when a set of particles is received, each node can immediately update the velocities and acceleration of their **own** set of particles without keeping a copy of the received particles.

After 3 rounds, all nodes will have seen all the particles and can perform calculations to update the positions of the particles they were allocated. With n nodes participating in the simulation, the calculate-and-forward step is repeated for $n - 1$ steps instead of 3 in our example above in order for all nodes to see all particles at once.

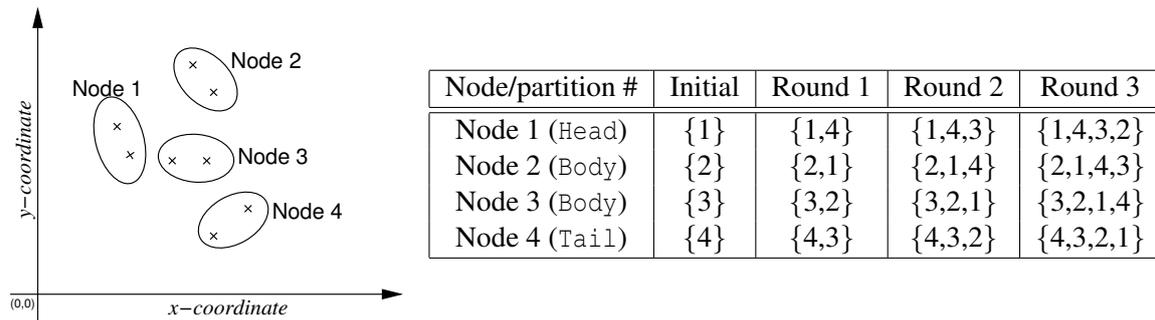


Fig. 3.3: Partitioning of data and calculation for 4 nodes

3.3 SJ with FPGA

After looking at the general architecture of SJ parallel applications on Axel, in this section we will detail an implementation of the n-body simulation using FPGA.

3.3.1 The need for cross-language features

As we have described in the introduction of the Axel cluster §2.5.2, the development environment for all of the heterogeneous components of Axel is C. However SJ is based on Java, and cannot use executables or shared libraries in C (or other *native* code) directly.

Java is an interpreted programming language that runs in the *Java Virtual Machine* (JVM). By design, the underlying architecture of the hardware is abstracted by the JVM completely, making Java a very portable language but it is not possible to access the memory or execute native instructions directly.

We also have the following considerations in programming language choice:

- Java is not as fast as native compiled languages such as C/C++ in most scenarios [?].
- SJ is a communication based language, the most important feature are its *communication capabilities* and *safety properties*.
- FPGAs and GPUs can be accessed by libraries supplied by vendor, but only C APIs and C-based development environments are available.

Combining the best points of SJ and acceleration hardware, we should delegate all communication and I/O coordination to SJ and all computation-related tasks to acceleration hardware for its performance.

In this design, SJ/Java will have to inter-operate with native libraries to access and control the acceleration hardware.

Alternative 1 It is certainly possible to take an extreme and approach this problem by re-implementing SJ and create a new session-types based language in C/C++. Despite the performance advantage, this will be a lot more involved than extending Java to SJ. C/C++ does not come with the rich set of readily usable network and datastructure libraries found in Java and upon which SJ depends heavily on. For the purpose of high performance computing, this would be the best option. In the next section we will present a C-translation of SJ built around the concept of a session-type based C/C++ programming language.

Alternative 2 The other end of the extreme is to translate all hardware drivers to Java so SJ can initiate computations from within the JVM. In the case of GPU, we will be translating the complete CUDA framework to Java. Some current unofficial implementations exists [?, ?]; and for FPGA, none of the vendors provide Java APIs for control and access. Moreover, ways of accessing system memory are very limited in Java. The design of Java is to *prevent* this mode of operation to decouple applications running in the JVM from the underlying hardware. This method is not easily generalisable.

Alternative 3 With the reasons above combined, the remaining option is to mix Java and C using a suitable bridging library. This way SJ can be used in its natural form, and hardware drivers and shared library code written in C can be used as it is designed. The bridging library should handle type conversions and data access between the two sides.

3.3.2 Java Native Interface

With *Java Native Interface* (JNI), applications written in languages other than Java can be accessed by Java using an interface understandable by the JVM. The user will write native code and export selected native subroutines via JNI. This is the standard native programming environment supported by the Java specification, and provides very fine grain control for data shared between the two sides. An example is given in the appendix §A.1.

3.3.3 Java Native Access

Java Native Access (JNA) [?] is an API built on top of JNI and takes on Java-C bridging in a very different approach. It does not require any boilerplate code in the native language, and therefore can use existing native libraries *without modifications* at the cost of performance.

JNA uses `libffi` to analyse the structure of the shared library at runtime. FFI stands for *Foreign Function Interface*, the purpose of FFI is to convert calling conventions and coordinate between programming languages [?]. JNA also came with a basic type-mapping infrastructure to allow data exchanged between the two sides. In addition to primitive type mappings, the library also maps `Structure` class to C-struct, and Java arrays (non-contiguous in memory) to C arrays (contiguous). An example is given in the appendix §A.2.

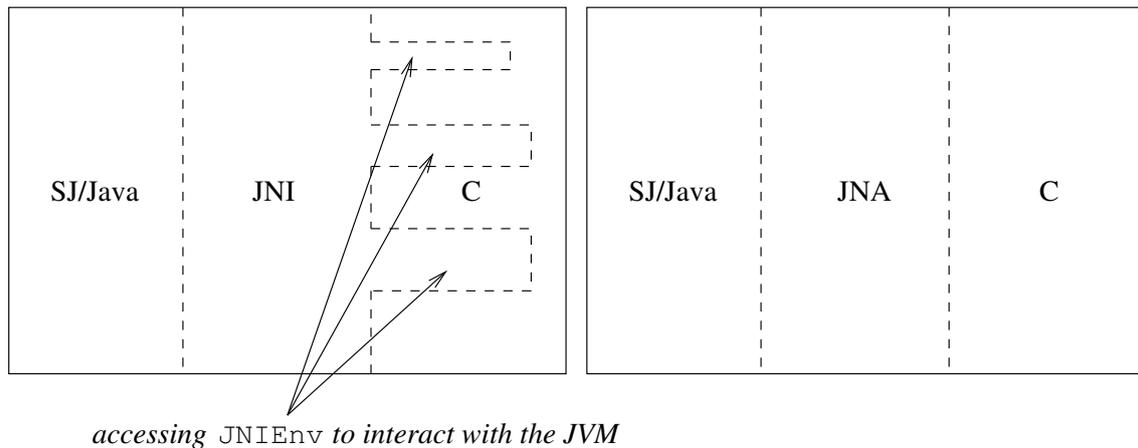


Fig. 3.4: Platform integration using JNI and JNA. JNA has a much cleaner interface than JNI

The above short introduction we showed that JNA is more flexible than JNI, where JNI forces a very tight integration between JVM and native platform.

For our implementation, we have chosen JNA over JNI because the bridging code has very little to do with session programming. Using JNA allows much more rapid prototyping to check correctness of communication. In JNA, we could simply load a different hardware driver if the underlying hardware is changed, as long as the interface of the compiled code remains the same.

In the next section §3.4 we will see a version of the implementation translated from SJ to C, which takes advantage of the shared library being an isolated component from Java.

3.3.4 Problems encountered

During the implementation, we encountered some problems and limitations

Java send buffer Java TCP sockets, which is what SJ uses when running on the cluster, is non-blocking. The sockets are made non-blocking by queueing the values to send in a send buffer. However, when the buffer is full, the semantics of Java TCP sockets became blocking. While there is a way to change the send buffer size, the send buffer size has a hard upper limit of 131071 bytes. Setting the send buffer size above this value will not change the actual buffer size. This caused some problems when trying to benchmark the performance of any SJ implementations using a very high number of particles.

This problem was addressed recently (and indirectly), by a new TCP socket implementation that uses custom TCP send and receive queues for events-based session programming [?].

JNA data conversion In C and most native languages, arrays are represented by contiguous blocks of main memory. In Java, array elements can be distributed all over the Java mem-

```

1  for (i = 0; i < particlesPerNode; i++) {
2      for (j = 0; j < particlesPerNode; j++) {
3
4          ri = receivedParticles[j].x - particles[i].x;
5          rj = receivedParticles[j].y - particles[i].y;
6          m = receivedParticles[j].m
7
8          if (ri != 0) {
9              ai += (ri < 0 ? -1 : 1) * G * m / (ri * ri);
10         }
11
12         if (rj != 0) {
13             aj += (rj < 0 ? -1 : 1) * G * m / (rj * rj);
14         }
15     }
16
17     particleVelocities[i].ai += ai;
18     particleVelocities[i].aj += aj;
19 }

```

Listing 3.2: *algorithm for SJ implementation's n-body simulation*

ory space for efficient management of free space. JNA allows force creation of contiguous blocks of memory in Java, only then the piece of memory can be passed to a native function.

Our implementation of n-body simulation receives an array of `Particles` in each iteration and passes the array to a native function to process. Since the array is received *as a Java object*, the memory is not contiguous. In order to pass the array to the native functions, the array needs to be copied to another contiguous-memory array in every iteration. This overhead could not be eliminated.

3.4 C-translation of SJ n-body implementation

This section will present a translation of the described in the previous section (§3.3). The motivation behind this translation is **performance**. When using a bridging library between two languages with completely disjoint language runtime, there is *always* overhead associated with the conversion between data formats. Also, we have discussed in previous sections (§3.3) that acceleration hardware libraries are provided in a C programming environment, it would be much more convenient if we could use session programming in a C environment. However, we also noted that the standard C do not have a comprehensive networking and datastructure library as in Java, and difficulties of building a completely new programming language based on C similar to how SJ is built on top of Java.

Our proposed solution is a C library that provide networking primitives available in SJ to C. The primitives made available to the programmer should have the same or very similar semantics

as their counterpart in SJ. For example, SJ's `send` primitive is non-blocking. The library should similarly implement a non-blocking `send` in the library. This library would be a step forward towards a session-types based C programming language we discussed in previous section (§3.3.1).

It should be reminded that the library is not designed to be used directly. The purpose of the library is to provide a building block of SJ programs translated to C, usage of the library does not automatically imply communication correctness if not translated from SJ.

3.4.1 Why would this work?

C/C++ shares a very similar language syntax. For basic language constructs and flow-control, conversion between the two languages are trivial. When the source SJ program uses a SJ primitive, the target C program will use one of the primitives provided in our library. This forms a backbone of our translated C program, as translation will be line by line, the sequence of invoking the SJ primitives will be identical in the translated version.

3.4.2 A SJ primitives library for C

Table 3.2 compares the equivalent primitives in the two languages:

C	SJ
<code>server_socket(port)</code>	<code>SJServerSocketImpl.create(p, port)</code>
<code>client_socket(host_addr, port)</code>	<code>SJService.create(p, host, port)</code>
<code>accept_connection(node)</code>	<code>node.accept()</code>
Included in <code>client_socket()</code>	<code>node.request()</code>
<code>send_int(node, value)</code> ¹	<code>node.sendInt(value)</code>
<code>inwhile(nodes[], nodes_count)</code>	<code><node1,node2,..>.inwhile{ }</code>
<code>outwhile(cond, node[], nodes_count)</code>	<code><node1,node2,..>.outwhile(cond) { }</code>

Table 3.2: C and SJ session primitives

In the section where we introduced JNA (§3.3.3), we have briefly discussed that *shared library*² to access FPGA can be reused without modifications. In most use cases, the shared library will provide a single function that takes input data and forward them to the FPGA. For n-body simulation, the main function in the shared library is `compute_forces(...)` shown in Fig. 3.5.

The translation do not need to worry about any new interface to access the acceleration hardware.

The syntax for `inwhile` and `outwhile` above shows slightly different syntax between the two languages. C-version of both constructs are typically used in conjunction with a while loop. The

¹`send` is implemented for all primitive types

²shared library refers to the native code component in the SJ/FPGA implementation, compiled as a shared object (.so)

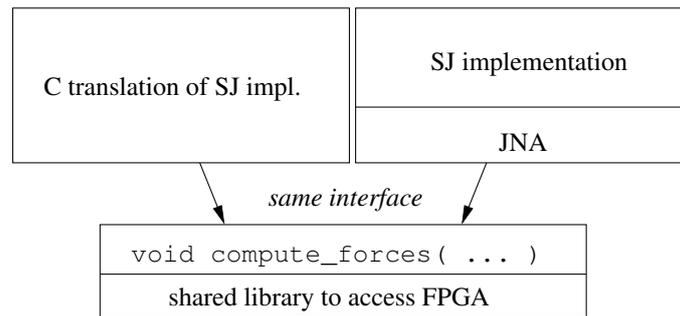


Fig. 3.5: Shared library used from both SJ/JNA and C

reason for the difference is because as a library external to the programming language, we are unable to modify the syntax of the language without going into parser of the compiler. We can, however, use C preprocessor macros to use a more familiar syntax. Listing 3.3 shows the usages of `inwhile` and `outwhile` in C.

```

1  #define OUTWHILE( COND, MOEP, NR_OF_MOEP ) \
2      while( outwhile( (COND), (MOEP), (NR_OF_MOEP) ) )
3  #define INWHILE( MOEP, NR_OF_MOEP ) \
4      while( inwhile( (MOEP), (NR_OF_MOEP) ) )
5  ...
6
7  outwhile_sfds[0] = next_fd;
8  outwhile_sfds[1] = tail_fd;
9  loop_index = 0;
10 while ( inwhile( loop_index < iter_count, outwhile_sfds, 2 ) ) { ... }
11 // OUTWHILE( loop_index < iter_count, outwhile_sfds, 2 ) { ... }
12
13 inwhile_sfds[0] = prev_fd;
14 inwhile_sfds[1] = head_fd;
15 while ( inwhile( inwhile_sfds, 2 ) ) { ... }
16 // INWHILE( inwhile_sfds, 2 ) { ... }

```

Listing 3.3: C `inwhile` and `outwhile`

In addition to the control flow and communication primitives, the library also has support for error handling in form of UNIX signal handlers. The C language do not have try-catch or similar system for exceptions. In SJ, runtime errors in sessions such as network error will `throw` an exception and all communicating processes will receive a `SJFIN` exception and terminate immediately. The library keeps track of all active connections in each process. If any of the connections encountered a problem, the process will force close all the active connections from its side. The other ends of the connections will then receive a `SIGPIPE` signal and promptly close all active connections. The signal is then propagate to all other connected processes until all processes are terminated.

```
1 #include <signal.h>
2 #include "sighandlers.h"
3 ...
4 int main(int argc, char *argv[])
5 {
6     signal(SIGPIPE, &sigpipe_handler);
7     signal(SIGSEGV, &sigsegv_handler);
8     ...
9 }
```

Listing 3.4: *Error handling in translated C*

To use the signal handlers provided, users only need to set the signal handlers to those provided by the library (header `sighandlers.h`) as shown in Listing 3.4.

Another consideration when designing the library is to make the translated C version have the same structure as the SJ version. For example, socket options when creating TCP sockets are completely hidden from the user. The options exposed to the user are the same as in Java version. Listing A.7 and A.8 shows the outline of the two versions without variable declarations and other unimportant details. The SJ implementation can almost map to the C-translation line by line.

3.4.3 Shortcomings of the library

The library is not a complete implementation of the full SJ. Session delegation and higher-order session manipulation is not possible with our library. The main reason is we lack a representation of *sessions*, which we omitted when we design our library for *translating* a session-based language to a non session-based language, rather than to add sessions to C.

3.5 Summary

In this chapter, we have discussed the details of our n-body implementation in SJ on the NNUS heterogeneous cluster Axel.

We looked at both the application architecture such as class layout and the session interaction pattern between the nodes in a ring-topology. We also looked at rationale to use a cross-language library *Java Native Access* (JNA), with a short introduction to the usage.

Finally we showed a manual translation of the SJ program to C, and a library that provides SJ primitives to C. The library contains a collection of light weight SJ primitives that mimics the SJ implementation, and can be a building block for an automatic SJ-to-C translator.

Chapter 4

Correctness Proof of N-body Implementation

In this chapter we will look at the formalisation of multichannel `inwhile` and `outwhile` primitives in SJ. The new primitives are designed for programming parallel algorithms in SJ, and had not been formalised in session calculus previously.

With the new formalisation, we will prove that our implementation of n-body simulation in SJ is *communication safe* from a global view and generalise the proof to algorithms with similar design and topology as our implementation.

We will first present an updated session calculus (§4.1) to include the new SJ primitives, where we will look at the syntax (§4.2), followed by the operational semantics (§4.3) and the type system (§4.4). Finally, we will look at *subject reduction* (§4.5) which will be used to prove communication safety property (§4.6)

4.1 Session calculus with multichannel in/outwhile

We will now present an extension of the session calculus to include the multichannel `inwhile` and `outwhile` SJ primitives used in parallel algorithms design with SJ. The original `inwhile` and `outwhile` primitives described in [?, ?] only operates in a single session channel. Synchronisation of `outwhile` loop condition between multiple session channels are not possible without re-opening the last session in each iteration of `outwhile` loop.

The multichannel constructs are the key components of parallel design with SJ, and parallel topologies can be expressed more naturally. Fig. 4.1 shows how the constructs improve ring topology design, first introduced in [?], and is used in our n-body implementation.

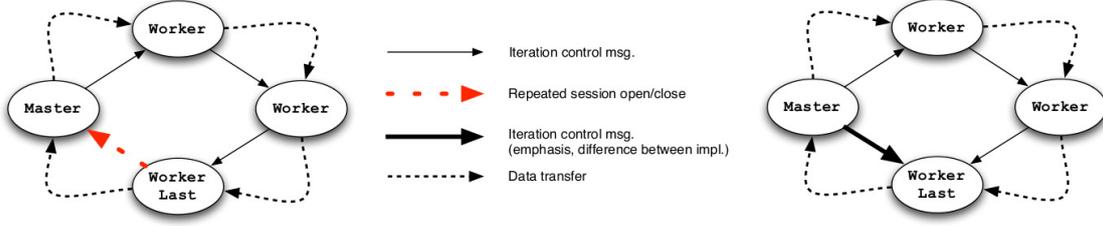


Fig. 4.1: Comparison of ring topology in unichannel and multichannel `inwhile` and `outwhile`

4.2 Syntax

The syntax of the updated session calculus with multichannel `inwhile` and `outwhile` is shown in Fig. 4.2.

The process definition is modified to include an `Err` process which represents a *while condition mismatch* in an `inwhile/outwhile` composition. *while condition mismatch* is further explained in the operational semantics §4.3.

Single channel `inwhile` and `outwhile` is sometimes written in the calculus as $k.inwhile\{Q\}$ and $k.outwhile(e)\{P\}$. This syntax is a shorthand for $\langle k \rangle.inwhile\{Q\}$ and $\langle k \rangle.outwhile(e)\{P\}$.

4.3 Operational semantics

The operational semantics are based on the reduction relation \rightarrow , and the reduction rules are given in Fig. 4.5. The session calculus is π -calculus extended with session primitives $[?]$, so definition of structural congruence \equiv is similar to π -calculus. Fig. 4.3 lists the structural congruence rules in our updated session calculus. An additional structural congruence rule in this calculus is $\mathbf{0};P \equiv P$, which allows continuation in sequential composition (Definition 1 below).

To keep session reasoning simple, we introduce evaluation contexts. Evaluation contexts isolate subprocesses and allow subprocesses to reduce independent of influences external to the context. Our evaluation contexts are defined as:

$$E[] := [] \mid E[];P \mid E[] \mid P \mid (\nu u) E[] \mid \text{def } D \text{ in } E[]$$

If the *head subprocess* P in $E[P]$ can be reduced using the reduction rules, then there is a dual *head subprocess* Q in $E[Q]$ that can be reduced $[?]$. This simplifies the reduction rules and allows us to avoid including explicit reduction rules such as sequential composition $(P;Q)$. The resulting reduction rules are shown in Fig. 4.5.

We have defined reduction rules for `inwhile` and `outwhile` such that they can reduce on their own. In particular, a single `outwhile` can generate an infinite number of $k \dagger [b]$ without constraints. Suppose the loop condition is `true` in the first run of the `outwhile`. Since `outwhile` can reduce without constraints, `outwhile` can reduce again, and the loop condition in this iteration is `false`.

$P ::= 0$	inaction
$ T$	prefixed process
$ P ; Q$	sequential composition
$ P \mid Q$	parallel composition
$ (\nu u) P$	name/channel hiding
$ \text{Err}$	error
$T ::= \text{request } a(k) \text{ in } P$	session request
$ \text{accept } a(k) \text{ in } P$	session acceptance
$ k![\tilde{e}]$	data sending
$ k?(\tilde{x}) \text{ in } P$	data reception
$ k \triangleleft l$	label selection
$ k \triangleright \{l_1 : P_1 \parallel \dots \parallel l_n : P_n\}$	label branching
$ \text{throw } k[k']$	channel sending
$ \text{catch } k(k') \text{ in } P$	channel reception
$ \text{if } e \text{ then } P \text{ else } Q$	conditional branch
$ X[\tilde{e}\tilde{k}]$	process variables
$ \text{def } D \text{ in } P$	recursion
$ \langle k_1 \dots k_n \rangle . \text{inwhile} \{ Q \} \quad n \geq 1$	multichannel inwhile
$ \langle k_1 \dots k_n \rangle . \text{outwhile}(e) \{ P \} \quad n \geq 1$	multichannel outwhile
$ k \dagger [b] \quad (b \in \text{true}, \text{false})$	(runtime syntax)
$e ::= c$	constant
$ \langle k_1 \dots k_n \rangle . \text{inwhile} \quad n \geq 1$	inwhile expression
$ e + e' \mid e - e' \mid e \times e \mid \text{not}(e) \mid \dots$	operators
$D ::= X_1(\tilde{x}_1 \tilde{k}_1) = P_1 \text{ and } \dots \text{ and } X_n(\tilde{x}_n \tilde{k}_n) = P_n$	declaration for recursion

Fig. 4.2: Session calculus syntax with multichannel inwhile and outwhile

This gives us $k.\text{outwhile}(e)\{P\} \mid k \dagger [\text{true}] \mid k \dagger [\text{false}] \mid k.\text{inwhile}\{Q\}$ which causes us problems because we do not know which $k \dagger [b]$ to compose with inwhile.

This differs from our implementation where the while loop synchronises nodes and delivers loop conditions in order. To correctly reflect the actual behaviour of multichannel inwhile and outwhile constructs in the calculus, we have an extra constraint that **inwhile rules have precedence over outwhile**. This way, loop conditions holders, $k \dagger [b]$ will prevent the sessions from continuing without first consuming $k \dagger [b]$ with a matching inwhile.

$$\begin{aligned}
& P \equiv Q \text{ if } P \equiv_{\alpha} Q \\
& P \mid \mathbf{0} \equiv P \quad P \mid Q \equiv Q \mid P \quad (P \mid Q) \mid R \equiv P \mid (Q \mid R) \\
& (\nu u) P \mid Q \equiv (\nu u) (P \mid Q) \quad \text{if } u \notin \text{fu}(\cdot)Q \\
& (\nu u) \mathbf{0} \equiv \mathbf{0} \\
& \text{def } D \text{ in } \mathbf{0} \equiv \mathbf{0} \\
& (\nu u) \text{def } D \text{ in } P \equiv \text{def } D \text{ in } (\nu u) P \quad \text{if } u \notin \text{fu}(\cdot)D \\
& (\text{def } D \text{ in } P) \mid Q \equiv \text{def } D \text{ in } (P \mid Q) \quad \text{if } \text{dpv}(\cdot)D \cap \text{fpv}(\cdot)Q = \emptyset \\
& \text{def } D \text{ in } (\text{def } D' \text{ in } P) \equiv \text{def } D \text{ and } D' \text{ in } P \quad \text{if } \text{dpv}(\cdot)D \cap \text{dpv}(\cdot)D' = \emptyset. \\
& \mathbf{0}; P \equiv P
\end{aligned}$$

Fig. 4.3: Structural Congruence

$$E[] := [] \mid E[]; P \mid E[] \mid P \mid (\nu u) E[] \mid \text{def } D \text{ in } E[]$$

Fig. 4.4: Evaluation context

4.4 Type system

The type system in this section is designed to guarantee communication safety and progress property with the new syntax and operational semantics. The full type syntax is given in Fig. 4.6.

Sorts contain the standard types and the pair of dual sessions $\langle \alpha, \bar{\alpha} \rangle$.

Partial session types are session types that does not include the end type. Partial session types are distinguished from completed session types so that they can be sequentially composed.

Completed session types are types that end with end or are equal to \perp .

In above syntax, $![\alpha]$ and $?[\alpha]$ are session delegation and session receive respectively. This makes use of the name-passing property from π -calculus that allows sending and receiving of channels (or *sessions* in the session calculus). The same typing syntax is used for ordinary type sending and receiving ($![\tilde{S}]$, $?[\tilde{S}]$). Iteration types ($?[\tau]^*$ and $![\tau]^*$) are introduced for *inwhile* and *outwhile* respectively. With iteration types, the partial type definition τ can be repeated for a number of times until the *outwhile* condition is no longer true.

In the syntax given, $\&\{l_1 : \tau_1, \dots, l_n : \tau_n\}.\text{end} \equiv \&\{l_1 :: \tau_1.\text{end}, \dots, l_n : \tau_n.\text{end}\}$. This equivalence ensures all partial types $\tau_1 \dots \tau_n$ of label selection choices ends and are compatible with each other in the completed session type (and vice versa).

ε is an empty type, and it is defined so that $\varepsilon; \tau \equiv \tau$ and $\tau; \varepsilon \equiv \tau$. The two equivalences allows us to continue reducing when one of the two processes $P; Q$ reduces to empty.

$$\begin{array}{l}
E_1[\text{accept } a(k) \text{ in } P_1] \mid E_2[\text{request } a(k) \text{ in } P_2] \rightarrow (E_1[P_1] \mid E_2[P_2]) \quad (k \text{ is fresh}) \quad [\text{LINK}] \\
E_1[k![\tilde{e}]] \mid E_2[k?(\tilde{x}) \text{ in } P_2] \rightarrow E_1[\mathbf{0}] \mid E_2[P_2[\tilde{c}/\tilde{x}]] \quad (\tilde{e} \downarrow \tilde{c}) \quad [\text{COM}] \\
E_1[k \triangleleft l_i; P] \mid E_2[k \triangleright \{l_1 : P_1 \parallel \dots \parallel l_n : P_n\}] \rightarrow E_1[P] \mid E_2[P_i] \quad (1 \leq i \leq n) \quad [\text{LABEL}] \\
E_1[\text{throw } k[k']] \mid E_2[\text{catch } k(k') \text{ in } P_2] \rightarrow E_1[\mathbf{0}] \mid E_2[P_2] \quad [\text{PASS}] \\
E[\text{if } e \text{ then } P \text{ else } Q] \rightarrow E[P] \quad (e \downarrow \text{true}) \quad [\text{IF1}] \\
E[\text{if } e \text{ then } P \text{ else } Q] \rightarrow E[Q] \quad (e \downarrow \text{false}) \quad [\text{IF2}] \\
\text{def } D \text{ in } (E[X[\tilde{e}\tilde{k}]]) \rightarrow \text{def } D \text{ in } (E[P[\tilde{c}/\tilde{x}]]) \quad (\tilde{e} \downarrow \tilde{c}, X(\tilde{x}\tilde{k}) = P \in D) \quad [\text{DEF}] \\
E[\langle k_1 \dots k_n \rangle.\text{inwhile}\{ P \}] \mid k_1 \dagger [b_1] \mid \dots \mid k_n \dagger [b_n] \rightarrow E[P; \langle k_1 \dots k_n \rangle.\text{inwhile}\{ P \}] \\
\quad (\forall i \in 1..n, b_i = \text{true}) \quad [\text{INWHI1}] \\
E[\langle k_1 \dots k_n \rangle.\text{inwhile}\{ P \}] \mid k_1 \dagger [b_1] \mid \dots \mid k_n \dagger [b_n] \rightarrow E[\mathbf{0}] \\
\quad (\forall i \in 1..n, b_i = \text{false}) \quad [\text{INWHI2}] \\
E[\langle k_1 \dots k_n \rangle.\text{inwhile}\{ P \}] \mid k_1 \dagger [b_1] \mid \dots \mid k_n \dagger [b_n] \rightarrow E[\text{Err}] \\
\quad (\exists i, j \ b_i = \text{true} \wedge b_j = \text{false} \wedge 1 \leq i, j \leq n) \quad [\text{INWHI3}] \\
E[\langle k_1 \dots k_n \rangle.\text{outwhile}(e)\{ P \}] \rightarrow E[P; \langle k_1 \dots k_n \rangle.\text{outwhile}(e')\{ P \}] \mid k_1 \dagger [b_1] \mid \dots \mid k_n \dagger [b_n] \\
\quad (\forall i \in 1..n, b_i = \text{true}) \quad (E[e] \rightarrow E[\text{true}]) \quad [\text{OUTWHI1}] \\
E[\langle k_1 \dots k_n \rangle.\text{outwhile}(e)\{ P \}] \rightarrow E[\mathbf{0}] \mid k_1 \dagger [b_1] \mid \dots \mid k_n \dagger [b_n] \\
\quad (\forall i \in 1..n, b_i = \text{false}) \quad (E[e] \rightarrow E[\text{false}]) \quad [\text{OUTWHI2}] \\
P \equiv P' \text{ and } P' \rightarrow Q' \text{ and } Q' \equiv Q \Rightarrow P \rightarrow Q \quad [\text{STR}] \\
P \rightarrow P' \Rightarrow E[P] \rightarrow E[P'] \quad [\text{EVAL}]
\end{array}$$

Fig. 4.5: Reduction rules

4.4.1 Duality

To ensure communication compatibility, all session types have a dual-type in a well-typed program.

A simple example is $![\text{bool}].\text{end}$ and $?[\text{bool}].\text{end}$. The two session types are dual so that sending of a `bool` matches with receiving of a `bool`. If the typing of the receiver is changed to $?[\text{bool}]; ?[\text{bool}].\text{end}$ then there is a communication mismatch after the first receive. Session types can ensure such incompatibilities between two communicating parties does not happen. Fig. 4.7 is complete list of dual-types in our type system.

Sort	$S ::= \text{nat} \mid \text{bool} \mid \langle \alpha, \bar{\alpha} \rangle$
Partial session type	$\tau ::= \varepsilon \mid \tau; \tau$ $\mid ?[\tilde{S}] \mid ?[\alpha] \mid \&\{l_1 : \tau_1, \dots, l_n : \tau_n\} \mid ![\tau]^* \mid \mathbf{x}$ $\mid ![\tilde{S}] \mid ![\alpha] \mid \oplus\{l_1 : \tau_1, \dots, l_n : \tau_n\} \mid ?[\tau]^* \mid \mu\mathbf{x}.\tau$
Completed session type	$\alpha ::= \tau.\text{end} \mid \perp$
Runtime session type	$\beta ::= \alpha \mid \alpha^\dagger \mid \dagger$

Fig. 4.6: Type syntax

$\bar{\varepsilon} = \varepsilon$	$\bar{\tau}; \bar{\tau} = \bar{\tau}; \bar{\tau}$	$\overline{\alpha^\dagger} = \overline{\alpha}^\dagger$
$\overline{![\tilde{S}]} = ?[\tilde{S}]$	$\overline{\oplus\{l_1 : \tau_1, \dots, l_n : \tau_n\}} = \&\{l_i : \bar{\tau}_i, \dots, l_n : \bar{\tau}_n\}$	$\overline{![\tau]} = ?[\tau]$
$\overline{?[\tilde{S}]} = ![\tilde{S}]$	$\overline{\&\{l_1 : \tau_1, \dots, l_n : \tau_n\}} = \oplus\{l_i : \bar{\tau}_i, \dots, l_n : \bar{\tau}_n\}$	$\overline{?[\tau]} = ![\tau]$
$\overline{![\tau]^*} = ?[\tau]^*$	$\bar{\mathbf{x}} = \mathbf{x}$	$\overline{\tau.\text{end}} = \bar{\tau}.\text{end}$
$\overline{?[\tau]^*} = ![\tau]^*$	$\overline{\mu\mathbf{x}.\tau} = \mu\mathbf{x}.\bar{\tau}$	$\overline{\perp} = \perp$

Fig. 4.7: Dual types

4.4.2 Typing environment

The typing environment is defined in Fig. 4.8.

$$\begin{aligned} \Gamma &::= \emptyset \mid \Gamma \cdot x : S \mid \Gamma \cdot X : \tilde{S}\tilde{\alpha} \\ \Delta &::= \emptyset \mid \Delta \cdot k : \alpha \mid \Delta \cdot k : \dagger \end{aligned}$$

Fig. 4.8: Typing environments

Γ is the *standard environment* that maps variables to sort types.

Δ is the *runtime environment* that contains session to session type mappings and the typing for $k \dagger [b]$, which holds `inwhile` and `outwhile` loop conditions.

4.4.3 Typing rules

Most of the typing rules remained the same as in [?] The major changes between the two version are

$\overline{\Gamma \vdash 1 \triangleright \text{nat}}$	$\overline{\Gamma \vdash \text{true}, \text{false} \triangleright \text{bool}}$	$\frac{\Gamma \vdash e_i \triangleright \text{nat}}{\Gamma \vdash e_1 + e_2 \triangleright \text{nat}}$	[NAT],[BOOL],[SUM]
$\frac{\Gamma \vdash P \triangleright \Delta \cdot k : \varepsilon.\text{end}}{\Gamma \vdash P \triangleright \perp}$	$\overline{\Gamma \cdot a : S \vdash a \triangleright S}$	$\frac{\Gamma; \Delta \vdash e \triangleright S}{\Gamma; \Delta, \Delta' \vdash e \triangleright S}$	[BOT],[NAMEI],[EVAL]
$\frac{\Delta = \{k_1 : \varepsilon.\text{end}, \dots, k_n : \varepsilon.\text{end}, k'_1 : \perp, \dots, k'_m : \perp\}}{\Gamma \vdash \mathbf{0} \triangleright \Delta}$			[INACT]
$\frac{\Gamma \vdash a \triangleright \langle \alpha, \bar{\alpha} \rangle, \Gamma \vdash P \triangleright \Delta \cdot k : \bar{\alpha}}{\Gamma \vdash \text{request } a(k) \text{ in } P \triangleright \Delta}$	$\frac{\Gamma \vdash a \triangleright \langle \alpha, \bar{\alpha} \rangle, \Gamma \vdash P \triangleright \Delta \cdot k : \alpha}{\Gamma \vdash \text{accept } a(k) \text{ in } P \triangleright \Delta}$		[REQ],[ACC]
$\frac{\Gamma; \emptyset \vdash \tilde{e} \triangleright \tilde{S}}{\Gamma \vdash k![e] \triangleright \Delta \cdot k : ![\tilde{S}].\text{end}}$	$\frac{\Gamma \cdot \tilde{x} : \tilde{S} \vdash P \triangleright \Delta \cdot k : \alpha}{\Gamma \vdash k?(\tilde{x}) \text{ in } P \triangleright \Delta \cdot k : ?[\tilde{S}]; \alpha}$		[SEND],[RCV]
$\frac{\Gamma \vdash P_1 \triangleright \Delta \cdot k : \tau_1.\text{end} \quad \dots \quad \Gamma \vdash P_n \triangleright \Delta \cdot k : \tau_n.\text{end}}{\Gamma \vdash k \triangleright \{l_1 : P_1 \parallel \dots \parallel l_n : P_n\} \triangleright \Delta \cdot k : \&\{l_1 : \tau_1, \dots, l_n : \tau_n\}.\text{end}}$			[BR]
$\frac{\Gamma \vdash P \triangleright \Delta \cdot k : \tau_j.\text{end}}{\Gamma \vdash k \triangleleft l \triangleright \Delta \cdot k : \oplus \{l_1 : \tau_1, \dots, l_n : \tau_n\}.\text{end}} \quad (1 \leq j \leq n)$			[SEL]
$\overline{\Gamma \vdash \text{throw } k[k'] \triangleright \Delta \cdot k : ![\alpha].\text{end} \cdot k' : \alpha}$			[THR]
$\frac{\Gamma \vdash P \triangleright \Delta \cdot k : \beta \cdot k' : \alpha}{\Gamma \vdash \text{catch } k(k') \text{ in } P \triangleright \Delta \cdot k : ?[\alpha]; \beta}$			[CAT]
$\frac{\Gamma; \emptyset \vdash e \triangleright \text{bool} \quad \Gamma; \Delta \vdash P \triangleright \Delta \quad \Gamma; \Delta \vdash Q \triangleright \Delta}{\Gamma \vdash \text{if } e \text{ then } P \text{ else } Q \triangleright \Delta}$			[IF]
$\frac{\Gamma; \Delta \vdash e \triangleright \text{bool} \quad \Gamma \vdash P \triangleright \Delta \cdot k_1 : \tau_1.\text{end} \dots \dots k_n : \tau_n.\text{end}}{\Gamma \vdash \langle k_1 \dots k_n \rangle.\text{outwhile}(e) \{ P \} \triangleright \Delta \cdot k_1 : ![\tau_1]*.\text{end} \dots \dots k_n : ![\tau_n]*.\text{end}}$			[OUTWHI]
$\frac{\Gamma; \Delta \vdash Q \triangleright \Delta \cdot k_1 : \tau_1.\text{end} \dots \dots k_n : \tau_n.\text{end}}{\Gamma \vdash \langle k_1 \dots k_n \rangle.\text{inwhile} \{ Q \} \triangleright \Delta \cdot k_1 : ?[\tau_1]*.\text{end} \dots \dots k_n : ?[\tau_n]*.\text{end}}$			[INWHI]
$\frac{\Gamma \cdot a : S \vdash P \triangleright \Delta}{\Gamma \vdash (\text{va}) P \triangleright \Delta}$	$\frac{\Gamma \vdash P \triangleright \Delta \cdot k : \perp}{\Gamma \vdash (\text{vk}) P \triangleright \Delta}$		[NRES],[CRES]
$\frac{\Gamma; \emptyset \vdash \tilde{e} \triangleright \tilde{S}}{\Gamma \cdot X : \tilde{S} \tilde{\alpha} \vdash X[\tilde{e} \tilde{k}] \triangleright \Delta \cdot \tilde{k} : \tilde{\alpha}}$			[VAR]
$\frac{\Gamma \cdot X : \tilde{S} \tilde{\alpha} \cdot \tilde{x} : \tilde{S} \vdash P \triangleright \tilde{k} : \tilde{\alpha} \quad \Gamma \cdot X : \tilde{S} \tilde{\tau} \vdash Q \triangleright \Delta}{\Gamma \vdash \text{def } X(\tilde{x} \tilde{k}) = P \text{ in } Q \triangleright \Delta}$			[DEF]
$\frac{\Gamma \vdash P \triangleright \Delta \quad \Gamma \vdash Q \triangleright \Delta'}{\Gamma \vdash P; Q : \Delta; \Delta'}$	$\frac{\Gamma \vdash P \triangleright \Delta \quad \Gamma \vdash Q \triangleright \Delta'}{\Gamma \vdash P \mid Q \triangleright \Delta \circ \Delta'}$		[SEQ],[CONC]

Fig. 4.9: Typing rules

1. Abandoning the use of Θ for mapping from variables to *basis* (ie. $X : \tilde{S}\tilde{\alpha}$), and use Γ instead for the mapping (see Typing environment definition above §4.8).
2. Introducing the new typing rules [OUTWHI] and [INWHI], which corresponds to our new constructs.
3. Rules do not have a continuation after $;$ (sequential composition). This is because we introduced evaluation contexts earlier (§4.3), and $;$ will not appear in the *head subprocess* (eg. $E[P;Q]$ will be written $E[P];E[Q]$). We also have a new typing rule for sequential composition [SEQ] for this reason.

The rules [NAT], [BOOL], [SUM], [NAMEI], [EVAL] are basic language constructs (numbers, booleans, inductive definition of numbers, variables and evaluation of expressions).

[INACT] represents inaction, and has a ‘end’ typing.

[REQ], [ACC]; [SEND], [RCV] are pairs that represent establishment of session and value/name exchange respectively.

[BR], [SEL] are label branching and selection. Each of the branches have a subtype τ_i , and when the branches finishes, the whole typing $\&\{l_1 : \tau_1, \dots, l_n : \tau_n\}$ and $\oplus\{l_1 : \tau_1, \dots, l_n : \tau_n\}$ ends.

[THR], [CAT] are called session delegation, which comes from π -calculus where channels can be passed as names and use as channels. In the P following a *catch*, the process has a typing received from the *throwing* side.

[IF], [NRES], [CRES], [VAR], [DEF] are conditionals, name restriction, channel restriction, variable process and recursive process definition respectively. Note that in [IF], Δ to prove e is set to \emptyset to prevent trouble with its channel when *inwhile* is used as an expression.

[OUTWHI], [INWHI] are the main focus of this work. It represents multichannel *inwhile* and *outwhile*. (It would be easier to understand with $n = 1$, which makes it a simple *inwhile/outwhile* loop)

Finally, [SEQ], [CONC] are sequential composition and parallel composition respectively. They will be introduced in detail next as Definition 1 and 2.

Definition 1. *Sequential composition of session type are defined as [?]:*

$$\tau; \alpha = \begin{cases} \tau.\alpha & \text{if } \tau \text{ is a partial session type and } \alpha \text{ is a completed session type} \\ \perp & \text{otherwise} \end{cases}$$

$$\Delta; \Delta' = \Delta \setminus \text{dom}(\Delta') \cup \Delta' \setminus \text{dom}(\Delta) \cup \{k : \Delta(k) \setminus \text{end}; \Delta'(k) \mid k \in \text{dom}(\Delta) \cap \text{dom}(\Delta')\}$$

The first rule concatenates a partial session type τ with a completed session type α to form a new (completed) session type. The second rule can be decomposed to three parts:

1. $\Delta \setminus \text{dom}(\Delta')$ extracts session types with sessions unique in Δ
2. $\Delta' \setminus \text{dom}(\Delta)$ extracts session types with sessions unique in Δ'

3. $\{k: \Delta(k) \setminus \text{end}; \Delta'(k) \mid k \in \text{dom}(\Delta) \cap \text{dom}(\Delta')\}$ modifies session types with a common session k in Δ and Δ' by removing end type from $\Delta(k)$ and concatenates the modified $\Delta(k)$ (which is now a partial session type) with $\Delta'(k)$ as described in the first rule.

Example 1. Suppose $\Delta = \{k_1: \varepsilon.\text{end}, k_2: ![nat].\text{end}\}$ and $\Delta' = \{k_2: ?[bool].\text{end}, k_3: ![bool].\text{end}\}$. Since k_1 is unique in Δ and k_3 is unique in Δ' , we have

$$\Delta \setminus \text{dom}(\Delta') = \{k_1: \varepsilon.\text{end}\} \text{ and } \Delta' \setminus \text{dom}(\Delta) = \{k_3: ![bool].\text{end}\}$$

A new session type is constructed by removing end in $\Delta(k_2)$, so the composed set of mappings is

$$\Delta; \Delta' = \{k_1: \varepsilon.\text{end}, k_2: ![nat]; ?[bool].\text{end}, k_3: ![bool].\text{end}\}$$

Definition 2. Parallel composition of session and runtime type is defined as:

$$\Delta \circ \Delta' = \Delta \setminus \text{dom}(\Delta') \cup \Delta' \setminus \text{dom}(\Delta) \cup \{k: \beta \circ \beta' \mid \Delta(k) = \beta \text{ and } \Delta'(k) = \beta'\}$$

where $\beta \circ \beta' : \begin{cases} \alpha \circ \dagger = \alpha^\dagger \\ \alpha \circ \bar{\alpha} = \perp \\ \alpha \circ \bar{\alpha}^\dagger = \perp^\dagger \end{cases}$

The parallel composition relation \circ is commutative as the order of composition do not impact the end result.

The rule can be decomposed into three parts:

1. $\Delta \setminus \text{dom}(\Delta')$ which extracts session and runtime types with sessions unique in Δ
2. $\Delta' \setminus \text{dom}(\Delta)$ which extracts session and runtime types with sessions unique in Δ'
3. $\{k: \beta \circ \beta' \mid \Delta(k) = \beta \text{ and } \Delta'(k) = \beta'\}$ has three cases
 - If one of the β is a \dagger , combine the session type with the \dagger to form an intermediate runtime type α^\dagger .
 - If β s are duals, combine the types to \perp . This covers cases for parallel compositions that does not involve runtime types.
 - If one of the β is an intermediate runtime type α^\dagger , and the other β is the dual of α , combine the types to \perp but mark the result as an intermediate runtime type \perp^\dagger since the \dagger has not been consumed.

Example 2. Suppose $\Delta = \{k_1: ![bool].\text{end}, k_2: ![nat].\text{end}, k_3: ?[nat].\text{end}\}$ and $\Delta' = \{k_1: ?[bool].\text{end}, k_2: ![bool].\text{end}, k_4: ![bool].\text{end}\}$. Since k_3 is unique in Δ and k_4 is unique in Δ' , the two sessions are included in $\Delta \circ \Delta'$ without modification. With $\Delta(k_1) = \overline{\Delta'(k_1)}$ and $\Delta(k_2) \neq \overline{\Delta'(k_2)}$, k_1 maps to bottom and k_2 is omitted. Therefore

$$\Delta \circ \Delta' = \{k_1: \perp, k_3: ?[nat].\text{end}, k_4: ![bool].\text{end}\}$$

4.5 Subject reduction

Next we are going to present *subject reduction* theorem. Subject reduction will enable us to reduce global composition of processes under a *well-formed ring topology* (defined in Definition 1), such as our implementation of n-body simulation. The main proof can be found in page 57.

Before we go into details of subject reduction theorem, we will begin with auxiliary results for later proofs to build on. The proofs presented here are based on [?] with modifications and additions to fit our updated type system with multichannel `inwhile` and `outwhile`.

The Weakening Lemma represents *adding* of mappings to the typing environment. Formally:

Lemma 1 (Weakening Lemma). *Let $\Gamma \vdash P \triangleright \Delta$.*

1. *If $X \notin \text{dom}(\Gamma)$, then $\Gamma \cdot X : \tilde{S}\tilde{\alpha} \vdash P \triangleright \Delta$.*
2. *If $a \notin \text{dom}(\Gamma)$, then $\Gamma \cdot a : S \vdash P \triangleright \Delta$.*
3. *If $k \notin \text{dom}(\Delta)$ and $\alpha = \perp$ or $\alpha = \varepsilon.\text{end}$, then $\Gamma \vdash P \triangleright \Delta \cdot k : \alpha$.*

Proof. For the first two sequent, simple induction of the derivation tree can show that X and a do not interfere with the typing. For 3, we note that in [INACT] and [VAR], Δ contains only $\varepsilon.\text{end}$ and \perp . □

The Strengthening Lemma represents *removal* of mappings from the typing environment, given that they do not change the typing of a process. Formally:

Lemma 2 (Strengthening Lemma). *Let $\Gamma \vdash P \triangleright \Delta$.*

1. *If $X \notin \text{fpv}(P)$, then $\Gamma \setminus X \vdash P \triangleright \Delta$.*
2. *If $a \notin \text{fn}(P)$, then $\Gamma \setminus a \vdash P \triangleright \Delta$.*
3. *If $k \notin \text{fc}(P)$, then $\Gamma \vdash P \triangleright \Delta \setminus k$.*

Proof. Start from $\Delta = \emptyset$, then by induction over all session constructs, showing all three sequent hold. □

The Channel Lemma states that if a channel is free in a process then it will have a typing in Δ , otherwise the typing can only be one of the end types that cannot react with other channels. Formally:

Lemma 3 (Channel Lemma). *1. If $\Gamma \vdash P \triangleright \Delta \cdot k : \alpha$ and $k \notin \text{fc}(P)$, then $\alpha = \perp, \varepsilon.\text{end}$.*

2. *If $\Gamma \vdash P \triangleright \Delta$ and $k \in \text{fc}(P)$, then $k \in \text{dom}(\Delta)$.*

Proof. A simple induction on the derivation tree for each sequent. □

We omit the standard renaming properties of variables and channels, but present the Substitution Lemma for names. Note that we do *not* require a substitution lemma for channels or process variables, for they are not communicated.

Lemma 4 (Substitution Lemma). *If $\Gamma \cdot x : S \vdash P \triangleright \Delta$ and $\Gamma \vdash c : S$, then $\Gamma \vdash P[c/x] \triangleright \Delta$*

Proof. By induction on the derivation tree. □

We write $\Delta \prec \Delta'$ if we obtain Δ' from Δ by replacing $k_1 : \varepsilon.\text{end}, \dots, k_n : \varepsilon.\text{end}$ ($n \geq 0$) in Δ by $k_1 : \perp, \dots, k_n : \perp$. If $\Delta \prec \Delta'$, we can obtain Δ' from Δ by applying the [BOT]-rule zero or more times.

4.5.1 Well-formed topology

We now introduce the notion of well-formed *ring* topology. These are the conditions which a correctly designed parallel algorithm based on a ring topology must satisfy.

Definition 1. *A process is under a well-formed ring topology if:*

$$\begin{aligned}
 P_1 &= \langle k_{1,2}, k_{1,n} \rangle . \text{outwhile}(e) \{ Q_1[k_{1,2}, k_{1,n}] \} \\
 P_{i \in \{2..n-1\}} &= k_{i,i+1} . \text{outwhile}(\langle k_{i-1,i} \rangle . \text{inwhile}) \{ Q_i[k_{i,i+1}, k_{i-1,i}] \} \quad 2 \leq i \leq n-1 \\
 P_n &= \langle k_{1,n}, k_{n-1,n} \rangle . \text{inwhile} \{ Q_n[k_{1,n}, k_{n-1,n}] \} \\
 \text{and } \Gamma \vdash Q_1 \triangleright \{ k_{1,2} : T_{1,2}, k_{1,n} : T_{1,n} \} \\
 \Gamma \vdash Q_i \triangleright \{ k_{i,i+1} : T_{i,i+1}, k_{i-1,i} : T'_{i-1,i} \} \\
 \Gamma \vdash Q_n \triangleright \{ k_{1,n} : T_{1,n}', k_{n-1,n} : T_{n-1,n}' \} \\
 \Gamma \vdash Q_1 \mid Q_2 \mid \dots \mid Q_n \triangleright \{ \tilde{k} : \tilde{\perp} \} \\
 \text{with } \overline{T_{i,j}} = T'_{i,j}
 \end{aligned}$$

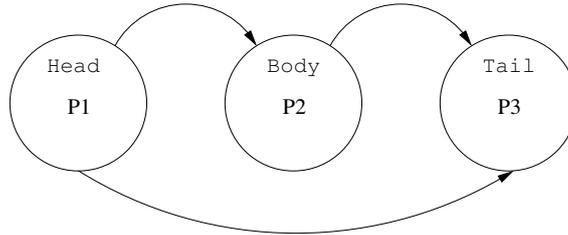


Fig. 4.10: Ring topology for 3 processes, arrow shows direction of *outwhile*

We also define a well-formed *intermediate* ring topology, which are the conditions that should hold when the reduction involves the runtime type \dagger as intermediate steps.

Definition 2. A process is under a well-formed intermediate ring topology if:

$$\begin{aligned}
P_1 &= \langle k_{1,2}, k_{1,n} \rangle . outwhile(e) \{ Q_1[k_{1,2}, k_{1,n}] \} \\
P_{i \in \{2..n-1\}} &= k_{i,i+1} . outwhile(\langle k_{i-1,i} \rangle . inwhile) \{ Q_i[k_{i,i+1}, k_{i-1,i}] \} \mid k_{i-1,i} \dagger [b] \quad b \in \{true, false\} \\
P_n &= \langle k_{1,n}, k_{n-1,n} \rangle . inwhile \{ Q_n[k_{1,n}, k_{n-1,n}] \} \mid k_{1,n} \dagger [b] \mid k_{n-1,n} [b] \quad \forall b = true \text{ or } \forall b = false \\
\text{and } \Gamma \vdash Q_1 \triangleright \{ k_{1,2} : T_{1,2}, k_{1,n} : T_{1,n} \} \\
\Gamma \vdash Q_i \triangleright \{ k_{i,i+1} : T_{i,i+1}, k_{i-1,i} : T_{i-1,i}' \} \\
\Gamma \vdash Q_n \triangleright \{ k_{1,n} : T_{1,n}'', k_{n-1,n} : T_{n-1,n}'' \} \\
\text{and } \Gamma \vdash Q_1 \mid Q_2 \mid \dots \mid Q_n \triangleright \{ \tilde{k} : \tilde{\perp}' \} \\
\text{with } \overline{T_{i,j}} = T_{i,j}'
\end{aligned}$$

4.5.2 Subject congruence theorem

Theorem 1. Subject congruence is defined by

$$\Gamma \vdash P \triangleright \Delta \text{ and } P \equiv P' \text{ implies } \Gamma \vdash P' \triangleright \Delta$$

Proof. **Case** $P \mid \mathbf{0} \equiv P$. We show that if $\Gamma \vdash P \mid \mathbf{0} \triangleright \Delta$, then $\Gamma \vdash P \triangleright \Delta$. Suppose

$$\Gamma \vdash P \triangleright \Delta_1 \quad \text{and} \quad \Gamma \vdash \mathbf{0} \triangleright \Delta_2.$$

with $\Delta_1 \circ \Delta_2 = \Delta$. Note that Δ_2 only contains $\varepsilon.\text{end}$ or \perp , hence we can set: $\Delta_1 = \Delta'_1 \circ \{ \tilde{k} : \varepsilon.\tilde{\text{end}} \}$ and $\Delta_2 = \Delta'_2 \cdot \{ \tilde{k} : \varepsilon.\tilde{\text{end}} \}$ with $\Delta'_1 \circ \Delta'_2 = \Delta'_1 \cdot \Delta'_2$ and $\Delta = \Delta'_1 \cdot \Delta'_2 \cdot \{ \tilde{k} : \tilde{\perp} \}$. Then by the [BOT]-rule, we have:

$$\Gamma \vdash P \triangleright \Delta'_1 \cdot \{ \tilde{k} : \tilde{\perp} \}$$

Notice that, given the form of Δ above, we know that $\text{dom}(\Delta'_2) \cap \text{dom}(\Delta'_1) \cdot \{ \tilde{k} : \perp \} = \emptyset$. Hence by applying Weakening, we have:

$$\Gamma \vdash P \triangleright \Delta'_1 \cdot \Delta'_2 \cdot \{ \tilde{k} : \tilde{\perp} \}$$

as required.

For the other direction, we set $\Delta = \mathbf{0}$ in [INACT].

Case $P \mid Q \equiv Q \mid P$. \circ relation is commutative by the definition of \circ (Definition 2)

Case $(P \mid Q) \mid R \equiv P \mid (Q \mid R)$. To show $(P \mid Q) \mid R \equiv P \mid (Q \mid R)$, where

$$\Gamma \vdash P \triangleright \Delta_1 \quad \Gamma \vdash Q \triangleright \Delta_2 \quad \Gamma \vdash R \triangleright \Delta_3$$

We assume $(\Delta_1 \circ \Delta_2) \circ \Delta_3$ is defined

Suppose $k : \beta_1 \in \Delta_1$ and $k : \beta_2 \in \Delta_2$, then we have

$$\left\{ \begin{array}{ll} \beta_1 = \alpha & \beta_2 = \dagger \\ \beta_1 = \alpha & \beta_2 = \bar{\alpha} \\ \beta_1 = \alpha & \beta_2 = \bar{\alpha}^\dagger \\ \beta_1 = \dagger & \beta_2 = \perp \end{array} \right.$$

Now suppose $k: \beta_3 \in \Delta_3$,
if $\beta_1 = \alpha \quad \beta_2 = \dagger$, then $\beta_3 = \bar{\alpha}$

$$\begin{aligned} (\beta_1 \circ \beta_2) \circ \beta_3 &= (\{k: \alpha\} \circ \{k: \dagger\}) \circ \{k: \bar{\alpha}\} = \{k: \perp^\dagger\} \\ \equiv \beta_1 \circ (\beta_2 \circ \beta_3) &= \{k: \alpha\} \circ (\{k: \dagger\} \circ \{k: \bar{\alpha}\}) = \{k: \perp^\dagger\} \end{aligned}$$

if $\beta_1 = \alpha \quad \beta_2 = \bar{\alpha}$, then $\beta_3 = \dagger$

$$\begin{aligned} (\beta_1 \circ \beta_2) \circ \beta_3 &= (\{k: \alpha\} \circ \{k: \bar{\alpha}\}) \circ \{k: \dagger\} = \{k: \perp^\dagger\} \\ \equiv \beta_1 \circ (\beta_2 \circ \beta_3) &= \{k: \alpha\} \circ (\{k: \bar{\alpha}\} \circ \{k: \dagger\}) = \{k: \perp^\dagger\} \end{aligned}$$

in all other cases, $k \notin \text{dom}(\Delta_3)$ and therefore no parallel composition is possible.

Case $(\nu u) P \mid Q \equiv (\nu u) (P \mid Q)$ if $u \notin \text{fv}(Q)$. The case when u is a name is standard. Suppose u is channel k and assume $\Gamma \vdash (\nu k) (P \mid Q) \triangleright \Delta$. We have

$$\frac{\Gamma \vdash P \triangleright \Delta'_1 \quad \Gamma \vdash Q \triangleright \Delta'_2}{\Gamma \vdash P \mid Q \triangleright \Delta' \cdot k: \perp}$$

with $\Delta' \cdot k: \perp = \Delta'_1 \circ \Delta'_2$ and $\Delta' \prec \Delta$ by [BOT]. First notice that k can be in either Δ'_i or in both. The interesting case is when it occurs in both; from Lemma 3(1) and the fact that $k \notin \text{fc}(Q)$ we know that $\Delta'_1 = \Delta_1 \cdot k: \varepsilon.\text{end}$ and $\Delta'_2 = \Delta_2 \cdot k: \varepsilon.\text{end}$. Then, by applying the [BOT]-rule to k in P , we have $\Gamma \vdash P \triangleright \Delta_1 \cdot k: \perp$, and by applying [CRES] we obtain $\Gamma \vdash (\nu k) P \triangleright \Delta_1$. On the other hand, by Strengthening, we have $\Gamma \vdash Q \triangleright \Delta_2$. Then, the application of [CONC] yields $\Gamma \vdash (\nu k) P \mid Q \triangleright \Delta'$. Then by applying the [BOT]-rule, we obtain $\Gamma \vdash (\nu k) P \mid Q \triangleright \Delta$, as required. The other direction is easy.

Case $(\nu u) \mathbf{0} \equiv \mathbf{0}$. Standard by Weakening and Strengthening.

Case $\text{def } D \text{ in } \mathbf{0} \equiv \mathbf{0}$. Similar to the first case using Weakening and Strengthening.

Case $(\nu u) \text{def } D \text{ in } P \equiv \text{def } D \text{ in } (\nu u) P$ if $u \notin \text{fv}(D)$. Similar to the scope opening case using Weakening and Strengthening.

Case $(\text{def } D \text{ in } P) \mid Q \equiv \text{def } D \text{ in } (P \mid Q)$ if $\text{dpv}(D) \cap \text{fpv}(Q) = \mathbf{0}$. Similar with the scope opening case using Weakening and Strengthening.

Case $\mathbf{0}; P \equiv P$. We show that if $\Gamma \vdash \mathbf{0}; P \triangleright \Delta$, then $\Gamma \vdash P \triangleright \Delta$. Suppose

$$\Gamma \vdash \mathbf{0} \triangleright \Delta_1 \quad \text{and} \quad \Gamma \vdash P \triangleright \Delta_2.$$

with $\Delta_1; \Delta_2 = \Delta$. Δ_2 only contains $\varepsilon.\text{end}$ or \perp , by definition of sequential composition (Definition 1), $\Delta(k) = \Delta_1(k). \Delta_2(k) = \varepsilon. \Delta_2(k) = \Delta_2(k)$ as required. \square

4.5.3 Subject reduction theorem

Theorem 2. *The following subject reduction rules hold for a well-formed ring topology.*

$$\Gamma \vdash P \triangleright \Delta \text{ and } P \rightarrow P' \text{ implies } \Gamma \vdash P' \triangleright \Delta' \quad \text{such that}$$

$$\begin{aligned} \Delta(k) = \alpha &\Rightarrow \begin{cases} \Delta'(k) = \alpha \\ \Delta'(k) = \alpha^\dagger \end{cases} \\ \Delta(k) = \alpha^\dagger &\Rightarrow \begin{cases} \Delta'(k) = \alpha \\ \Delta'(k) = \alpha^\dagger \end{cases} \end{aligned}$$

Under a well-formed intermediate ring topology

$$\Gamma \vdash P \triangleright \Delta \text{ and } P \rightarrow^* P' \text{ implies } \Gamma \vdash P' \triangleright \Delta' \quad \text{such that}$$

$$\begin{aligned} \Delta(k) = \alpha &\Rightarrow \begin{cases} \Delta'(k) = \alpha \\ \Delta'(k) = \alpha^\dagger \end{cases} \\ \Delta(k) = \alpha^\dagger &\Rightarrow \begin{cases} \Delta'(k) = \alpha \\ \Delta'(k) = \alpha^\dagger \end{cases} \end{aligned}$$

Proof. We assume that

$$\Gamma \vdash e \triangleright S \quad \text{and} \quad e \downarrow c \quad \text{implies} \quad \Gamma \vdash c \triangleright S \quad (4.1)$$

and prove the result by induction on the last rule applied.

Case [LINK] (accept $a(k)$ in P_1) | (request $a(k)$ in P_2) \rightarrow (vk) (P_1 | P_2). Suppose $\Gamma \vdash$ (accept $a(k)$ in P_1) | (request $a(k)$ in P_2) $\triangleright \Delta$. Then the assumption is derived from:

$$\frac{\Gamma \vdash a \triangleright \langle \alpha, \bar{\alpha} \rangle \quad \Gamma \vdash P_1 \triangleright \Delta'_1 \cdot k : \alpha}{\Gamma \vdash \text{accept } a(k) \text{ in } P_1 \triangleright \Delta'_1} \quad \text{and} \quad \frac{\Gamma \vdash a \triangleright \langle \alpha, \bar{\alpha} \rangle \quad \Gamma \vdash P_2 \triangleright \Delta'_2 \cdot k : \bar{\alpha}}{\Gamma \vdash \text{request } a(k) \text{ in } P_2 \triangleright \Delta'_2}$$

and [BOT] with $\Delta'_i \prec \Delta_i$, [CONC] with $\Delta_1 \circ \Delta_2 = \Delta'$, and [BOT] with $\Delta' \prec \Delta$. Then applying [BOT] to P_1 and P_2 , we have:

$$\frac{\Gamma \vdash P_1 \triangleright \Delta'_1 \cdot k : \alpha}{\Gamma \vdash P_1 \triangleright \Delta_1 \cdot k : \alpha} \quad \text{and} \quad \frac{\Gamma \vdash P_2 \triangleright \Delta'_2 \cdot k : \bar{\alpha}}{\Gamma \vdash P_2 \triangleright \Delta_2 \cdot k : \bar{\alpha}}$$

Then we apply [CONC] to P_1 and P_2 to obtain:

$$\frac{\Gamma \vdash P_1 \triangleright \Delta_1 \cdot k : \alpha \quad \Gamma \vdash P_2 \triangleright \Delta_2 \cdot k : \bar{\alpha}}{\Gamma \vdash P_1 | P_2 \triangleright \Delta' \cdot k : \perp}$$

Now applying [CRES] and [BOT], we are done.

Case [COM] ($k![\tilde{e}]; P_1$) | ($k?(\tilde{x})$ in P_2) \rightarrow P_1 | $P_2[\tilde{c}/\tilde{x}]$ with $\tilde{e} \downarrow \tilde{c}$. The assumption is derived from:

$$\frac{\Gamma \vdash \tilde{e} \triangleright \tilde{S} \quad \Gamma \vdash P_1 \triangleright \Delta'_1 \cdot k : \alpha}{\Gamma \vdash k![\tilde{e}]; P_1 \triangleright \Delta'_1 \cdot k : ![\tilde{S}]} \quad \text{and} \quad \frac{\Gamma \cdot \tilde{x} : \tilde{S} \vdash P_2 \triangleright \Delta'_2 \cdot k : \bar{\alpha}}{\Gamma \vdash k?(\tilde{x}) \text{ in } P_2 \triangleright \Delta'_2 \cdot k : ?[\tilde{S}]; \bar{\alpha}}$$

and [BOT] with $\Delta'_i \prec \Delta_i$, [CONC] with $\Delta_1 \circ \Delta_2 \cdot k : \perp = \Delta'$, and [BOT] with $\Delta' \prec \Delta$. Then by (4.1), we know $\Gamma \vdash \tilde{c} \triangleright \tilde{S}$. By applying Substitution Lemma, we have:

$$\Gamma \vdash P_2[\tilde{c}/\tilde{x}] \triangleright \Delta'_2 \cdot k : \bar{\alpha}$$

Now the application of [BOT] and [CONC] to P_1 and $P_2[\tilde{c}/\tilde{x}]$, then by [BOT], we complete this case.

Case [STR]. By Subject-Congruence.

Case inwhile/outwhile for 3 processes $(\nu k_{12}, k_{23}, k_{13}) (P_1 \mid P_2 \mid P_3)$. Assume well-formed ring topology (Definition 1)

Case $E[e] \rightarrow E[\text{true}]$

By [OUTWHI1],

$$\begin{aligned} & (\nu k_{12}, k_{23}, k_{13}) (\langle k_{13}, k_{12} \rangle.\text{outwhile}(e)\{ Q_1[k_{13}, k_{12}] \} \mid \\ & \quad k_{23}.\text{outwhile}(k_{12}.\text{inwhile})\{ Q_2[k_{12}, k_{23}] \} \mid \\ & \quad \langle k_{13}, k_{23} \rangle.\text{inwhile}\{ Q_n[k_{13}, k_{23}] \}) \\ \rightarrow & (\nu k_{12}, k_{23}, k_{13}) (k_{13} \dagger [\text{true}] \mid k_{12} \dagger [\text{true}] \mid \\ & \quad Q_1[k_{13}, k_{12}]; \langle k_{13}, k_{12} \rangle.\text{outwhile}(e')\{ Q_1[k_{13}, k_{12}] \} \mid \\ & \quad k_{23}.\text{outwhile}(k_{12}.\text{inwhile})\{ Q_2[k_{12}, k_{23}] \} \mid \\ & \quad \langle k_{13}, k_{23} \rangle.\text{inwhile}\{ Q_3[k_{13}, k_{23}] \}) \end{aligned}$$

$$\begin{aligned} \Gamma \vdash (k_{13} \dagger [\text{true}] \mid k_{12} \dagger [\text{true}] \mid Q_1; P_1 \mid P_2 \mid P_3) \triangleright \{ & k_{12}: T_{12}; ![T_{12}]^* \circ ?[T'_{12}]^{*\dagger}, \\ & k_{13}: T_{13}; ![T_{13}]^* \circ ?[T'_{13}]^{*\dagger}, \\ & k_{23}: ![T_{23}]^* \circ ?[T'_{23}]^* \} \end{aligned}$$

By [INWHI1],

$$\begin{aligned} & (\nu k_{12}, k_{23}, k_{13}) (k_{13} \dagger [\text{true}] \mid k_{12} \dagger [\text{true}] \mid \\ & \quad Q_1[k_{13}, k_{12}]; \langle k_{13}, k_{12} \rangle.\text{outwhile}(e')\{ Q_1[k_{13}, k_{12}] \} \mid \\ & \quad k_{23}.\text{outwhile}(k_{12}.\text{inwhile})\{ Q_2[k_{12}, k_{23}] \} \mid \\ & \quad \langle k_{13}, k_{23} \rangle.\text{inwhile}\{ Q_3[k_{13}, k_{23}] \}) \\ \rightarrow & (\nu k_{12}, k_{23}, k_{13}) (k_{13} \dagger [\text{true}] \mid \\ & \quad Q_1[k_{13}, k_{12}]; \langle k_{13}, k_{12} \rangle.\text{outwhile}(e')\{ Q_1[k_{13}, k_{12}] \} \mid \\ & \quad k_{23}.\text{outwhile}(\text{true})\{ Q_2[k_{12}, k_{23}] \} \mid \\ & \quad \langle k_{13}, k_{23} \rangle.\text{inwhile}\{ Q_3[k_{13}, k_{23}] \}) \end{aligned}$$

$$\Gamma \vdash (k_{13} \dagger [\text{true}] \mid Q_1; P_1 \mid P_2 \mid P_3) \triangleright \{ k_{12}: T_{12}; ![T_{12}]^* \circ \epsilon.\text{end}^\dagger, k_{13}: T_{13}; ![T_{13}]^* \circ ?[T'_{13}]^{*\dagger}, k_{23}: ![T_{23}]^* \circ ?[T'_{23}]^* \}$$

By [OUTWHI1],

$$\begin{aligned}
& (\forall k_{12}, k_{23}, k_{13}) (k_{13} \dagger [\text{true}] \mid \\
& \quad Q_1[k_{13}, k_{12}]; \langle k_{13}, k_{12} \rangle. \text{outwhile}(e') \{ Q_1[k_{13}, k_{12}] \} \mid \\
& \quad k_{23}. \text{outwhile}(k_{12}. \text{inwhile}) \{ Q_2[k_{12}, k_{23}] \} \mid \\
& \quad \langle k_{13}, k_{23} \rangle. \text{inwhile} \{ Q_3[k_{13}, k_{23}] \}) \\
\rightarrow & (\forall k_{12}, k_{23}, k_{13}) (k_{13} \dagger [\text{true}] \mid k_{23} \dagger [\text{true}] \mid \\
& \quad Q_1[k_{13}, k_{12}]; \langle k_{13}, k_{12} \rangle. \text{outwhile}(e') \{ Q_1[k_{13}, k_{12}] \} \mid \\
& \quad Q_2[k_{12}, k_{23}]; k_{23}. \text{outwhile}(\text{true}) \{ Q_2[k_{12}, k_{23}] \} \mid \\
& \quad \langle k_{13}, k_{23} \rangle. \text{inwhile} \{ Q_3[k_{13}, k_{23}] \}) \\
\Gamma \vdash & (k_{13} \dagger [\text{true}] \mid k_{23} \dagger [\text{true}] \mid Q_1; P_1 \mid Q_2; P_2 \mid P_3) \triangleright \{ k_{12}: T_{12}; ![T_{12}]^* \circ T_{12}'^\dagger; ?[T_{12}']^*, \\
& \quad k_{13}: T_{13}; ![T_{13}]^* \circ ?[T_{13}]^{*\dagger}, \\
& \quad k_{23}: T_{23}; ![T_{23}]^* \circ ?[T_{23}']^{*\dagger} \}
\end{aligned}$$

By [INWHI1],

$$\begin{aligned}
& (\forall k_{12}, k_{23}, k_{13}) (k_{13} \dagger [\text{true}] \mid k_{23} \dagger [\text{true}] \mid \\
& \quad Q_1[k_{13}, k_{12}]; \langle k_{13}, k_{12} \rangle. \text{outwhile}(e') \{ Q_1[k_{13}, k_{12}] \} \mid \\
& \quad Q_2[k_{12}, k_{23}]; k_{23}. \text{outwhile}(\text{true}) \{ Q_2[k_{12}, k_{23}] \} \mid \\
& \quad \langle k_{13}, k_{23} \rangle. \text{inwhile} \{ Q_n[k_{13}, k_{23}] \}) \\
\rightarrow & (\forall k_{12}, k_{23}, k_{13}) (Q_1[k_{13}, k_{12}]; \langle k_{13}, k_{12} \rangle. \text{outwhile}(e') \{ Q_1[k_{13}, k_{12}] \} \mid \\
& \quad Q_2[k_{12}, k_{23}]; k_{23}. \text{outwhile}(\text{true}) \{ Q_2[k_{12}, k_{23}] \} \mid \\
& \quad Q_3[k_{13}, k_{23}]; \langle k_{13}, k_{23} \rangle. \text{inwhile} \{ Q_3[k_{13}, k_{23}] \}) \\
\Gamma \vdash & (Q_1; P_1 \mid Q_2; P_2 \mid Q_3; P_3) \triangleright \{ k_{12}: T_{12}; ![T_{12}]^* \circ T_{12}'^\dagger; ?[T_{12}']^*, \\
& \quad k_{13}: T_{13}; ![T_{13}]^* \circ T_{13}'^\dagger; ?[T_{13}']^*, \\
& \quad k_{23}: T_{23}; ![T_{23}]^* \circ T_{23}'^\dagger; ?[T_{23}']^* \} \\
\Gamma \vdash & (Q_1; P_1 \mid Q_2; P_2 \mid Q_3; P_3) \triangleright \{ k_{12}: \perp^\dagger; \perp, k_{13}: \perp^\dagger; \perp, k_{23}: \perp^\dagger; \perp \}
\end{aligned}$$

Case $E[e] \rightarrow E[\text{false}]$

By [OUTWHI2],

$$\begin{aligned}
& (\forall k_{12}, k_{23}, k_{13}) (\langle k_{13}, k_{12} \rangle. \text{outwhile}(e) \{ Q_1[k_{13}, k_{12}] \} \mid \\
& \quad k_{23}. \text{outwhile}(k_{12}. \text{inwhile}) \{ Q_2[k_{12}, k_{23}] \} \mid \\
& \quad \langle k_{13}, k_{23} \rangle. \text{inwhile} \{ Q_n[k_{13}, k_{23}] \}) \\
\rightarrow & (\forall k_{12}, k_{23}, k_{13}) (k_{13} \dagger [\text{false}] \mid k_{12} \dagger [\text{false}] \mid \mathbf{0} \mid \\
& \quad k_{23}. \text{outwhile}(k_{12}. \text{inwhile}) \{ Q_2[k_{12}, k_{23}] \} \mid \\
& \quad \langle k_{13}, k_{23} \rangle. \text{inwhile} \{ Q_3[k_{13}, k_{23}] \})
\end{aligned}$$

$$\Gamma \vdash (k_{13} \dagger [\text{true}] \mid k_{12} \dagger [\text{true}] \mid \mathbf{0} \mid P_2 \mid P_3) \triangleright \{k_{12} : \varepsilon.\text{end} \circ ?[T'_{12}]^{*\dagger}, k_{13} : \varepsilon.\text{end} \circ ?[T'_{13}]^{*\dagger}, k_{23} : ![T_{23}]^* \circ ?[T'_{23}]^*\}$$

By [INWHI2],

$$\begin{aligned} & (\nu k_{12}, k_{23}, k_{13}) (\langle k_{13}, k_{12} \rangle.\text{outwhile}(e) \{ Q_1[k_{13}, k_{12}] \} \mid \\ & \quad k_{23}.\text{outwhile}(k_{12}.\text{inwhile}) \{ Q_2[k_{12}, k_{23}] \} \mid \\ & \quad \langle k_{13}, k_{23} \rangle.\text{inwhile} \{ Q_n[k_{13}, k_{23}] \}) \\ \rightarrow & (\nu k_{12}, k_{23}, k_{13}) (k_{13} \dagger [\text{false}] \mid \mathbf{0} \mid \\ & \quad k_{23}.\text{outwhile}(\text{false}) \{ Q_2[k_{12}, k_{23}] \} \mid \\ & \quad \langle k_{13}, k_{23} \rangle.\text{inwhile} \{ Q_3[k_{13}, k_{23}] \}) \end{aligned}$$

$$\Gamma \vdash (k_{13} \dagger [\text{true}] \mid k_{12} \dagger [\text{true}] \mid \mathbf{0} \mid P_2 \mid P_3) \triangleright \{k_{12} : \varepsilon.\text{end}, k_{13} : \varepsilon.\text{end} \circ ?[T'_{13}]^{*\dagger}, k_{23} : ![T_{23}]^* \circ ?[T'_{23}]^*\}$$

By [OUTWHI2],

$$\begin{aligned} & (\nu k_{12}, k_{23}, k_{13}) (k_{13} \dagger [\text{false}] \mid \mathbf{0} \mid k_{23}.\text{outwhile}(\text{false}) \{ Q_2[k_{12}, k_{23}] \} \mid \\ & \quad \langle k_{13}, k_{23} \rangle.\text{inwhile} \{ Q_3[k_{13}, k_{23}] \}) \\ \rightarrow & (\nu k_{12}, k_{23}, k_{13}) (k_{13} \dagger [\text{false}] \mid k_{23} \dagger [\text{false}] \mid \mathbf{0} \mid \mathbf{0} \mid \\ & \quad \langle k_{13}, k_{23} \rangle.\text{inwhile} \{ Q_3[k_{13}, k_{23}] \}) \end{aligned}$$

$$\Gamma \vdash (k_{13} \dagger [\text{true}] \mid k_{12} \dagger [\text{true}] \mid \mathbf{0} \mid P_2 \mid P_3) \triangleright \{k_{12} : \varepsilon.\text{end}, k_{13} : \varepsilon.\text{end} \circ ?[T'_{13}]^{*\dagger}, k_{23} : \varepsilon.\text{end} \circ ?[T'_{23}]^*\}$$

By [INWHI2],

$$\begin{aligned} & (\nu k_{12}, k_{23}, k_{13}) (k_{13} \dagger [\text{false}] \mid k_{23} \dagger [\text{false}] \mid \mathbf{0} \mid \mathbf{0} \mid \langle k_{13}, k_{23} \rangle.\text{inwhile} \{ Q_3[k_{13}, k_{23}] \}) \\ \rightarrow & (\nu k_{12}, k_{23}, k_{13}) (\mathbf{0} \mid \mathbf{0} \mid \mathbf{0}) \end{aligned}$$

$$\Gamma \vdash (k_{13} \dagger [\text{true}] \mid k_{12} \dagger [\text{true}] \mid \mathbf{0} \mid \mathbf{0} \mid \mathbf{0}) \triangleright \{k_{12} : \varepsilon.\text{end}, k_{13} : \varepsilon.\text{end}, k_{23} : \varepsilon.\text{end}\}$$

Finally, apply [BOT].

The result can be extended to $(P_1 \mid \dots \mid P_n)$ by expanding the middle process from P_2 to P_i ($2 \leq i \leq n$).

Case [PASS] $(\text{throw } k[k']; P_1) \mid (\text{catch } k(k') \text{ in } P_2) \rightarrow P_1 \mid P_2$. The assumption is derived from:

$$\frac{\Gamma \vdash P_1 \triangleright \Delta'_1 \cdot k : \beta}{\Gamma \vdash \text{throw } k[k']; P_1 \triangleright \Delta'_1 \cdot k : ![\alpha]; \beta \cdot k' : \alpha}$$

and

$$\frac{\Gamma \vdash P_2 \triangleright \Delta'_2 \cdot k : \bar{\beta} \cdot k' : \alpha}{\Gamma \vdash \text{catch } k(k') \text{ in } P_2 \triangleright \Delta'_2 \cdot k : ?[\alpha]; \bar{\beta}}$$

and [BOT] with $\Delta'_i \prec \Delta_i$, [CONC] with $\Delta_1 \circ \Delta_2 \cdot k : \perp \cdot k' : \alpha = \Delta'$ and [BOT] with $\Delta' \prec \Delta$. Note that $k, k' \notin \text{dom}(\Delta_1, \Delta_2, \Delta'_1, \Delta'_2)$. By applying [BOT], [CONC] to P_1 and P_2 , and then by [BOT], we obtain the required result.

Case [IF1],[IF2]. Trivial.

Case [DEF] $\text{def } D \text{ in } (X[\tilde{e}\tilde{k}] \mid Q) \rightarrow \text{def } D \text{ in } (P[\tilde{c}/\tilde{x}] \mid Q)$ with $\tilde{e} \downarrow \tilde{c}$ and $X(\tilde{x}\tilde{k}) = P \in D$. Simplifying the recursive definition to the single case, we set $D = (X(\tilde{x}\tilde{k}) = P)$. Then the assumption is derived from:

$$\frac{\Gamma X : \tilde{S}\tilde{\alpha} \vdash X[\tilde{e}\tilde{k}] \triangleright \Delta'_1 \cdot \tilde{k} : \tilde{\alpha} \quad \Gamma \cdot X : \tilde{S}\tilde{\alpha}; \vdash Q \triangleright \Delta'_2}{\frac{\Gamma \cdot X : \tilde{S}\tilde{\alpha} \cdot \tilde{x} : \tilde{S} \vdash P \triangleright \tilde{k} : \tilde{\alpha} \quad \Gamma \cdot X : \tilde{S}\tilde{\alpha} \vdash X[\tilde{e}\tilde{k}] \mid Q \triangleright \Delta'' \cdot \tilde{k} : \tilde{\alpha} \quad \Delta'' \prec \Delta'}{\Gamma \vdash \text{def } X(\tilde{x}\tilde{k}) = P \text{ in } (X[\tilde{e}\tilde{k}] \mid Q) \triangleright \Delta' \cdot \tilde{k} : \tilde{\alpha}}}$$

with $\Delta_0 = \Delta' \cdot \tilde{k} : \tilde{\alpha}$, $\Delta' = \Delta'_1 \circ \Delta'_2$ and $\Delta_0 \prec \Delta$. Note that Δ'_1 contains only \perp or ε .end. Then applying Substitution Lemma to P , we have:

$$\Gamma \cdot X : \tilde{S}\tilde{\alpha} \vdash P[\tilde{c}/\tilde{x}] \triangleright \tilde{k} : \tilde{\alpha}$$

Notice that $\tilde{k} \cap \text{dom}(\Delta'_1) = \emptyset$, since $(\Delta'_1 \circ \Delta'_2) \cdot \tilde{k} : \tilde{\alpha}$ is defined. Then by Weakening, we have:

$$\Gamma \cdot X : \tilde{S}\tilde{\alpha} \vdash P[\tilde{c}/\tilde{x}] \triangleright \Delta'_1 \cdot \tilde{k} : \tilde{\alpha}$$

Now by [CONC], we have

$$\Gamma \cdot X : \tilde{S}\tilde{\alpha} \vdash P[\tilde{c}/\tilde{x}] \mid Q \triangleright \Delta'' \cdot \tilde{k} : \tilde{\alpha}$$

Finally by [BOT] ($\Delta'' \prec \Delta'$), then by [DEF], we obtain:

$$\Gamma \vdash \text{def } X(\tilde{x}\tilde{k}) = P \text{ in } (P[\tilde{c}/\tilde{x}] \mid Q) \triangleright \Delta' \cdot \tilde{k} : \tilde{\alpha}$$

Then we can apply [BOT] to obtain Δ , as desired. □

4.6 Progress property

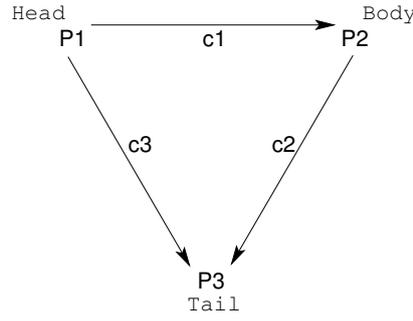
We can now model any process (composition of processes) that uses a ring topology, and show that they are deadlock free if they conform to our definition of well-formed topology (Definition 1). An exception is when there are shared names in the process P , composition with other processes might change the well-formed topology property of the process thus making our deadlock-free claim invalid.

The proof uses subject reduction theorem proven above (§4.5.3).

Theorem 1. *Suppose $\Gamma \vdash P \triangleright \Delta$ and P is under a well-formed topology without shared names Then P is deadlock free*

$$\text{ie. Suppose } P \rightarrow^* P' \text{ then } \begin{cases} \text{either } P' \equiv \mathbf{0} \\ \text{or } \exists Q (P' \rightarrow Q) \end{cases}$$

Proof. Let P be a process under well-formed topology and do not have shared names. Given the typings are correct, under subject reduction, no process will reduce to a deadlock state so P is deadlock free in all cases we have shown in above proof. □

Fig. 4.11: Shared channels c_1, c_2, c_3 between 3 n-body processes

4.7 Correctness proof for n-body simulation

In Table 3.1 from the previous section, we have shown the session declarations of the ring topology used in our n-body implementation.

4.7.1 N-body simulation in session calculus

Fig. 4.11 shows the shared channels between the processes in the n-body simulation. The processes can be represented in session calculus by

$$\begin{aligned}
 P_1 &\equiv \text{request } c_1(k_{12}) \text{ in request } c_3(k_{13}) \text{ in } k_{12}?(int) \text{ in } \langle k_{12}, k_{13} \rangle.\text{outwhile}(e)\{ Q_1 \} \\
 P_2 &\equiv \text{request } c_2(k_{23}) \text{ in accept } c_1(k_{12}) \text{ in } k_{23}?(int) \text{ in } k_{12}![int]; k_{23}.\text{outwhile}(k_{12}.\text{inwhile})\{ Q_2 \} \\
 P_3 &\equiv \text{accept } c_2(k_{23}) \text{ in accept } c_3(k_{13}) \text{ in } k_{23}![int]; \langle k_{13}, k_{23} \rangle.\text{inwhile}\{ Q_3 \}
 \end{aligned}$$

$$\begin{aligned}
 Q_1 &\equiv \langle k_{12}, k_{13} \rangle.\text{outwhile}(e)\{ k_{12}![\text{Particle}[]] \mid k_{13}?(Particle[]) \text{ in } \mathbf{0} \} \\
 Q_2 &\equiv k_{23}.\text{outwhile}(k_{12}.\text{inwhile})\{ k_{23}![\text{Particle}[]] \mid k_{12}?(Particle[]) \text{ in } \mathbf{0} \} \\
 Q_3 &\equiv \langle k_{13}, k_{23} \rangle.\text{inwhile}\{ k_{13}![\text{Particle}[]] \mid k_{23}?(Particle[]) \text{ in } \mathbf{0} \}
 \end{aligned}$$

The typing of the processes are

$$\begin{aligned}
 \Gamma \vdash P_1 \triangleright \{ k_{12}: ?[int]; ![![Particle[]].end]^*.end, k_{13}: ![?[Particle[]].end]^*.end \} \\
 \Gamma \vdash P_2 \triangleright \{ k_{23}: ?[int]; ![![Particle[]].end]^*.end, k_{12}: ![int]; ?[?[Particle[]].end]^*.end \} \\
 \Gamma \vdash P_3 \triangleright \{ k_{13}: ?[?[Particle[]].end]^*.end, k_{23}: ![int]; ?[?[Particle[]].end]^*.end \}
 \end{aligned}$$

Reduction

To prove that the our program, $((\nu k)_{12, k_{23}, k_{13}})(P_1 \mid P_2 \mid P_3)$ is deadlock free, it needs to satisfy the preconditions laid out in progress property (§4.6).

- Process is under a well-formed topology
- Process do not have shared names

We first inspect the typing of the process from the end,

$$\begin{aligned} \Gamma \vdash (k_{12}![\text{Particle}[]] \mid k_{13}?(\text{Particle}[]) \text{ in } \mathbf{0}) \triangleright \{k_{12}: ![\text{Particle}[]].\text{end}, k_{13}: ?[\text{Particle}[]].\text{end}\} \\ \Gamma \vdash (k_{23}![\text{Particle}[]] \mid k_{12}?(\text{Particle}[]) \text{ in } \mathbf{0}) \triangleright \{k_{23}: ![\text{Particle}[]].\text{end}, k_{12}: ?[\text{Particle}[]].\text{end}\} \\ \Gamma \vdash (k_{13}![\text{Particle}[]] \mid k_{23}?(\text{Particle}[]) \text{ in } \mathbf{0}) \triangleright \{k_{13}: ![\text{Particle}[]].\text{end}, k_{23}: ?[\text{Particle}[]].\text{end}\} \end{aligned}$$

When the three processes are composed, all of the sessions have a dual, satisfying the conditions for $Q_1, Q_i, \text{ and } Q_n$ in Definition 1

$$\begin{aligned} \Gamma \vdash Q_1 \triangleright \{k_{1,2}: T_{1,2}, k_{1,n}: T_{1,n}\} \\ \Gamma \vdash Q_i \triangleright \{k_{i,i+1}: T_{i,i+1}, k_{i-1,i}: T'_{i-1,i}\} \\ \Gamma \vdash Q_n \triangleright \{k_{1,n}: T_{1,n}', k_{n-1,n}: T_{n-1,n}'\} \end{aligned}$$

and $\overline{T_{i,j}} = T'_{i,j}$.

Given above, the subprocesses Q_1, Q_2, Q_3 is under a well-formed topology by matching the structures in the definition.

We can also show that the processes P_1, P_2, P_3 are under a well-formed topology, after the sessions are established with `request` and `accept` pairs and node information exchange (send and receive of a single `int`). Shared names do not exist in the process after the link phase.

Therefore, by progress property, the n-body implementation is deadlock free.

4.8 Summary

In this chapter we formalised the multichannel **inwhile** and **outwhile** construct in session calculus. An updated session calculus and session typing system is presented. We have also included a proof for subject reduction of **inwhile outwhile**, and by that shown a well-formed ring topology will have progress property and never deadlocks.

Chapter 5

Testing and Evaluation

In this chapter we will first discuss some failed attempts (§5.1) and testing results (§5.2) that influenced the current design of SJ applications on Axel.

Then we will look at the benchmark results of our n-body implementation with SJ comparing our results to non-session based message passing solution such as MPJExpress (§5.3.2). We will also look at the performance of our C translation, and compare the results with ordinary SJ and SJ with FPGA (§5.3.2).

5.1 Alternative designs

5.1.1 SJ and acceleration hardware allocation

Current design The current design of applications on cluster maps 1 SJ executable to 1 hardware accelerator. Multiple SJ executables can be run on a single node to use multiple hardware. Such as node using both FPGA and CPU shown in the example Fig 5.1. This approach is very simple and the class design can be minimalistic, with implementations for specific hardware encapsulated in a single class. eg. `FPGAHead` is a `Head` node that uses FPGA. `CPUbody` is a `Body` node that uses C on CPU. `JavaTail` is a `Tail` node that uses SJ/Java on CPU. We can execute the three nodes to have a hybrid execution with FPGA and SJ.

SJ as a coordinator We have previously considered using a single SJ application running on a node to be a coordinator between hardware accelerators and other nodes. This structure can allow dynamic load balancing if the hardware accelerators are of different performance (eg. coordinating both a GPU and a FPGA connected to the same node), especially since SJ will be in a central position that overlooks all aspects of inter-node and inter-component communications.

Typically there will be a single complicated function that we wish to accelerate. In the case of our n-body simulation, the said function is `computeForces()`. *If a SJ program controls more than one hardware accelerators*, then it would make sense to have `computeForces()` transparently

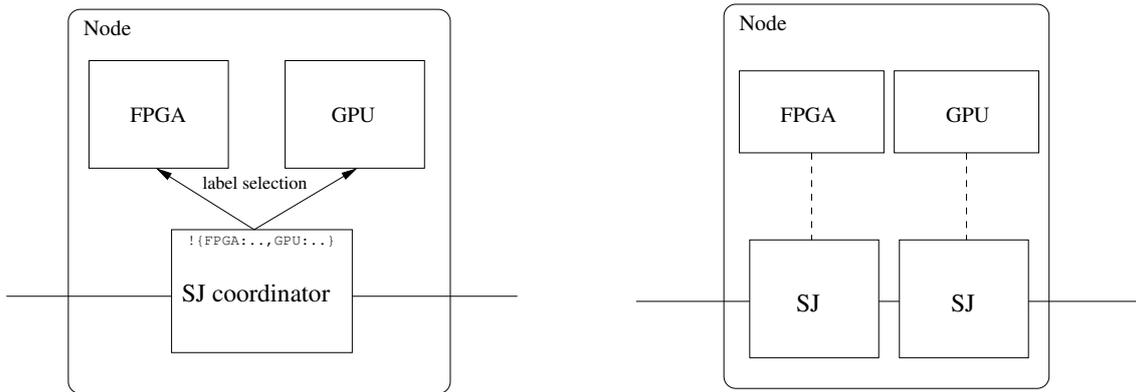


Fig. 5.1: Left: *SJ as a coordinator*, Right: *current design*

handle load balancing between the hardware accelerators based on their workload or performance. There might be a chance to exploit some advanced session programming constructs such as *label selection* if the communication between hardware accelerators and main SJ components is also session-typed.

The reason this is ultimately not implemented is because the lack of advantage over ordinary method call. We currently only have a single type of hardware accelerator implemented (FPGA), and the performance compared with native implementations are far from good; or in other words, we do not have spare performance for features such as a load balancer. In the current toolchain from Axel's SDK, input values are partitioned to different hardware components based on a static XML configuration file. This has worked well so far on applications built for Axel, with [?] showing good results with a 1/3 GPU and 2/3 FPGA manual partition.

Dynamic load balancing based on SJ, however, remains a novel idea for future work.

5.1.2 Communication medium

Explicit SJ communication The original proposal was to introduce label selection to the accelerator - CPU communication, so in our main SJ program we can demonstrate using two labels for the two compute tasks, `computeForces()` invoked in every iteration and `computePositions()` invoked in every completed ring to update the positions of the particles. Because of reasons outlined in 3.2.1 that FPGA is more suitable to accelerate a single task, we instead. Seeing that there are no extra benefits on using explicit communication in accelerator - CPU link, we reverted to using an implicit (class based method call) communication between *CPU part* and *FPGA part* as in current design.

However if **SJ as a coordinator** were implemented, most likely the communication medium will be in SJ.

Shared memory between SJ and hardware accelerators This was attempted but was later abandoned. There are many advantages in using shared memory (SHM) to share data between



Fig. 5.2: Left: using SJ to share data, Right: using shared memory to share data

two different processes. Shared memory is efficient, and is the method which data is exchanged between C code and FPGA in the Axel SDK [?, ?]. Some parts of FPGA memory is mapped to main memory to pass simple function arguments (such as size of array FPGA should expect) and to kick start the computation on FPGA.

If Java have a robust mechanism to access shared memory directly, this will save us a lot of effort (and overhead) passing data as function arguments, then through a cross-language library (JNA) to communicate with the FPGA. As expected from the design principals of Java, SHM with the host operating system is not possible without the use of *Java Native Interface* (JNI), because of the closed nature of the Java Virtual Machine. This offers no obvious advantage over our JNA-based solution.

5.2 Pre-implementation tests: inner product

At the initial stage of the project, we wish to run a simple algorithm to check that all the libraries work as intended and the communication and the overall design are correct.

Two of the main features we wished to examine were methods of using JNA and the magnitude of overhead when using JNA to access native functions.

Inner product was the algorithm used for this test, where all the computing nodes are first loaded with a_i $i \in \{1..n\}$. At each step, $\sum_{i=1}^n a_i \times b_i$ where b_i $i \in \{1..n\}$ is received from a neighbour node.

5.2.1 JNA direct mapping

The JNA project states that there are two methods of using the library. **interface mapping** and **direct mapping**. The usage of the two methods are quite similar, and the developers of JNA

```

1 public class CPUInnerProduct extends InnerProduct {
2     public CPUInnerProduct () {
3         Native.register(System.getProperty("user.dir")+"/lib/libinp_cpu.so");
4     }
5
6     // Direct mapped function
7     public static native int innerproduct(int[] a, int[] b, int size);
8 }

```

Listing 5.1: JNA direct mapping example

library strongly encourages the use of direct mapping for high performance applications¹ because of the lower native function calling overhead.

To tell the difference between the two methods, we should first remember that JNA analyses native libraries at *runtime*.

For interface mapping, a Java interface needs to be supplied. For example, `LibExample` interface in Listing A.2 specifies what JNA should expect in our supplied shared library. In our example, the shared library is `libexample.so` in the calling directory. At runtime, JNA analyses the `Library` interface and instantiate any new classes for our declared type. `LibExample` does not use any external classes or datastructure, but it is common that some parameters map to a class or C struct. For example, our implementation of n-body simulation uses a `Particle` class to represent a particle.

Instead of providing a subclass of JNA's `Library` interface, the direct mapping method allows taking the function signature and declare them directly as a `native static` method, given there is a build-in mapping between chosen *primitives* or *arrays of primitives*. This is much more convenient for simple datatype than normal interface mapping; However, this method do not support all function parameter and return types. In particular, arrays of Pointer-based classes are not allowed in direct mapped methods.

5.2.2 JNA interface mapping and direct mapping

Fig. 5.3 shows the performance when direct mapping was compared to interface mapping method and native SJ. From the figure it can be seen that direct mapping has a slight performance edge over interface mapping. *CPU* in the legend refers to computation code written in C and runs in the CPU.

5.2.3 Execution in CPU and FPGA

Next, we compare the performance of running inner product in CPU and FPGA.

¹<https://jna.dev.java.net/#direct>

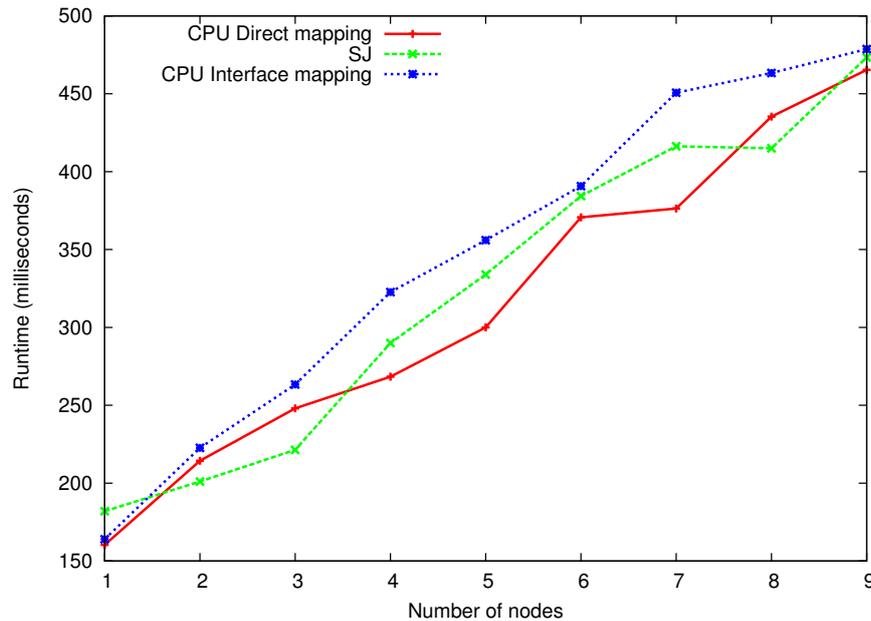


Fig. 5.3: *JNA direct mapping shows better performance over interface mapping*

In Fig. 5.4 FPGA shows a much worse performance than either SJ or CPU implementation. This can be explained by the lack of complex operations and pipelinable operations. The calculation of inner product involves n multiplications and $(n - 1)$ summation steps, so it is of $O(n)$ in each node. Our main implementation to run on the cluster, the n -body simulation, calculates the aggregate forces between a particle and other $(n - 1)$ particles. This is repeated for n particles, and requires $n \times (n - 1)$ operations in total. N -body simulation is therefore more complex than inner product at $O(n^2)$.

When a simple algorithm is implemented on the FPGA, it spends a low proportion of time in computing the results, but a high proportion of time in the transfer of data to and from the FPGA memory which is separate from the main memory. This might outweigh all benefits of using a fast hardware accelerator, since in an ordinary microprocessor, the data can be accessed directly and do not need the extra data transfer.

It is possible that with a big enough problem size, the total computation time in FPGA that includes data transfer will be less than total computation time in microprocessor. But do remember when the problem size is increased, the data transfer time will increase accordingly. If the total computation time increases in a rate equal or lower than which a CPU scales, as we saw in Fig. 5.3, then FPGA might not be a feasible solution for your task.

Nonetheless, the main reason for this test is to verify that FPGAs can be operated from SJ and calculates results correctly. It had been shown in [?] that n -body is a viable algorithm to run on FPGAs.

The conclusion of the initial testing with inner product is

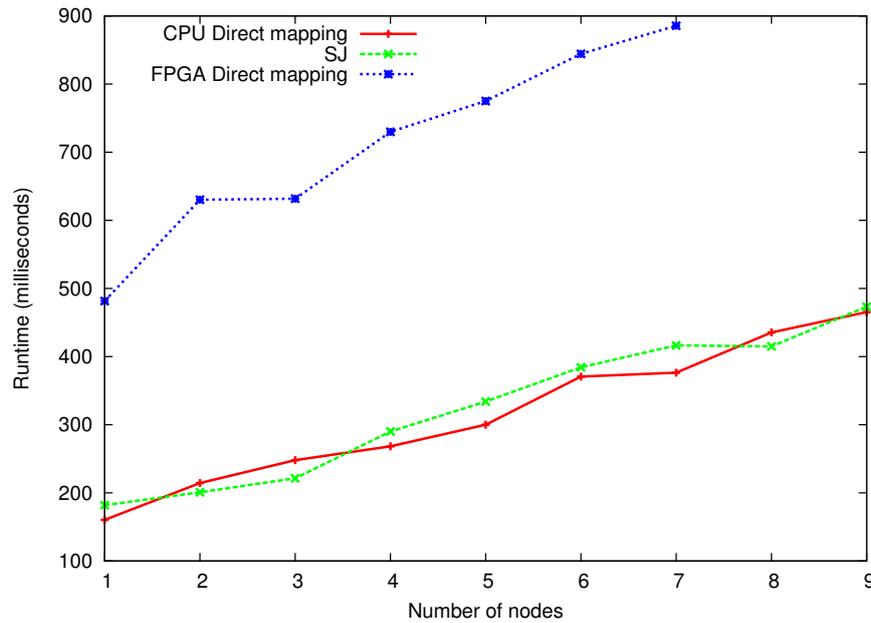


Fig. 5.4: *FPGA performance of innerproduct is much worse than JNA direct mapping and native SJ implementation*

- Direct mapping, as the authors of JNA have suggested, yields better performance.
- It is necessary that the main computation function in the parallel algorithms to be sufficiently complex to overcome the data transfer overhead

5.3 Benchmarks

This section contains all benchmark results and comparison between different implementations with Axel and SJ.

	Computation	Communication
SJ	Java	SJ
SJ + C	C with JNA	SJ
SJ + FPGA	FPGA with JNA	SJ
MPJExpress	Java	MPJExpress
Translated C	C	TCP-sockets, translated from SJ

Table 5.1: *Implementations which we will compare*

5.3.1 Benchmark methodology

The initial particles configurations came from the *Dubinski 1995 data set* available on <http://bima.astro.umd.edu/nemo/archive/>. Each node will load their partition of the particles, the starting offset of the particle indices are calculated by the node's position in the ring.

Head will be node0 and Tail will be node8, the number of particles each node loads is specified as a command-line argument.

This allows flexible assignment of number of particles to each node such that more particles can be assigned to a node if the node runs on faster hardware.

For all implementations with SJ, 5 warm-up iterations are run before timing begins. This allows the Java *Just-in-time* (JIT) compiler to optimise the code for a (marginally) better performance.

To run the implementations, the SJ application is launched in each of the nodes involved in the computation. We have put together *sessionj-tools*, a set of Perl scripts ² to automatically resolve hosts and port numbers of each component and connect the nodes in the correct order.

5.3.2 Benchmark results

SJ-based implementation

First we present the total runtime of three implementations of n-body simulation with SJ, some of the results are shown in Table 5.2 and plotted on Fig. 5.5.

SJ A pure SJ implementation that does not involve the JNA library or acceleration hardware.

SJ + C An implementation that uses JNA library to bridge SJ with C. The main computation function is implemented in C that runs on the CPU.

SJ + FPGA An implementation that uses JNA library to bridge SJ with C. The main computation function marshals the data from SJ and forwards to and from the FPGA using DMA.

MPJExpress An implementation of the MPI standard in pure Java. This is our candidate for non-session based message passing framework. In previous comparisons in [?], SJ performs competitively with MPJExpress.

The graph in Fig. 5.5 shows that with an increasing number of particles, the performance of the FPGA implementation starts to be more efficient than pure SJ. Runtime for the SJ + FPGA combination overtook SJ when the total particle number is over 33000. We have discussed the importance of the complexity of the algorithm and choosing suitable problem size in §5.2.3, this is the point where the problem size is big enough for FPGA to be feasible. The best performance of

²Full details of the miniproject can be found on <http://www.doc.ic.ac.uk/~cn06/sessionj-tools/>

Runtime (ms)		Implementation			
		SJ	SJ + CPU	SJ + FPGA	MPJExpress
Total number of particles	17600	6534	13861	9129	4450
	19800	7845	15973	10064	5393
	22000	10037	17750	11047	6896
	24200	10808	19159	13801	8708
	28600	13354	21571	14558	12497
	30800	14819	24141	15627	15979
	33000	17644	24057	15801	18747
	35200	19567	23747	16602	22801
	37400	21246	23310	17744	25900
	39600	23750	25262	18055	27054
	41800	26743	30150	18242	30266
	44000	29522	32352	19145	30950

Table 5.2: A partial table of results showing the crucial point when runtime of SJ+FPGA implementation overtakes SJ and MPJExpress

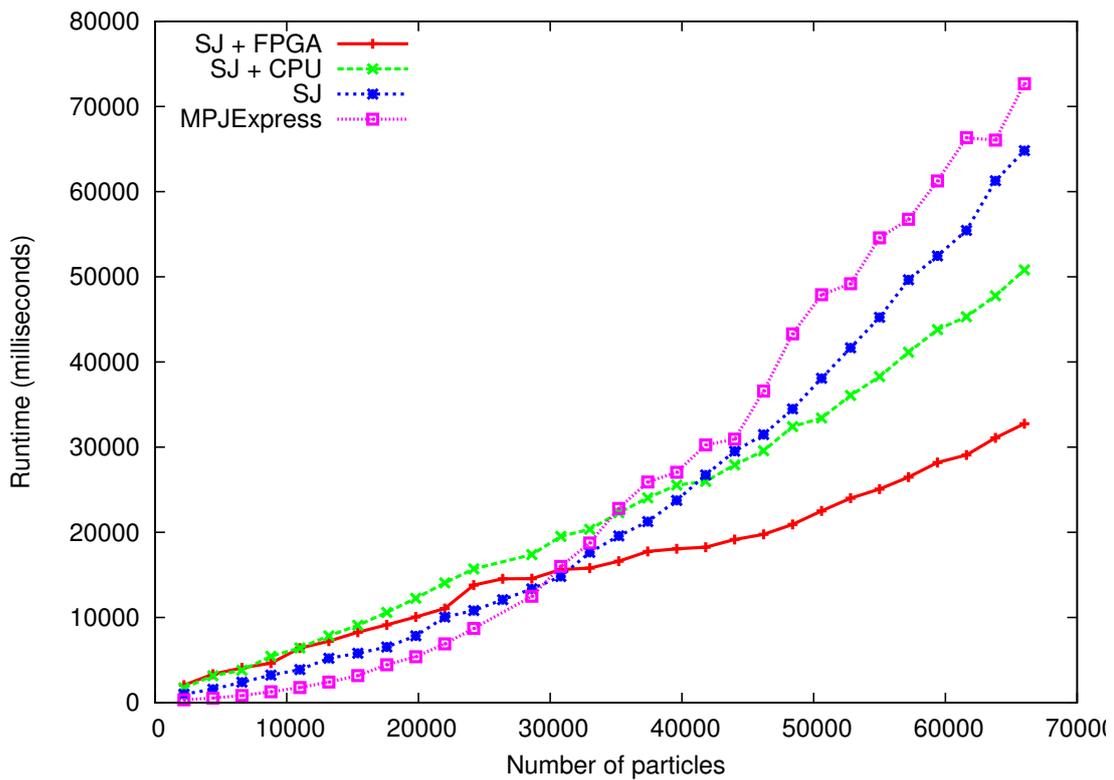


Fig. 5.5: Runtime results against number of particles in 11 nodes over 5 iterations

SJ + FPGA is when the number of particles reach maximum in our benchmark, which received almost 2 times speedup compared to SJ implementation. (Fig. A.1 in the appendix shows a complete graph of speedup against the same x-axis)

For the SJ + C version, the performance is worse than SJ. This is expected since the communication in SJ + C version is identical to that of SJ version, and the calculation uses identical hardware. The towards can only be explained by other minor factors such as JVM activities or network latency.

Comparison with C-translation

While the performance of the SJ + FPGA arrangement fared well with the SJ counterpart, the results of SJ implementations are dismal compared to a native C implementation as shown in Fig. 5.6. Speedup compared to SJ is 7 times maximum, and the average speedup is 5 times. Again, Fig. A.1 shows the speedup in more details.

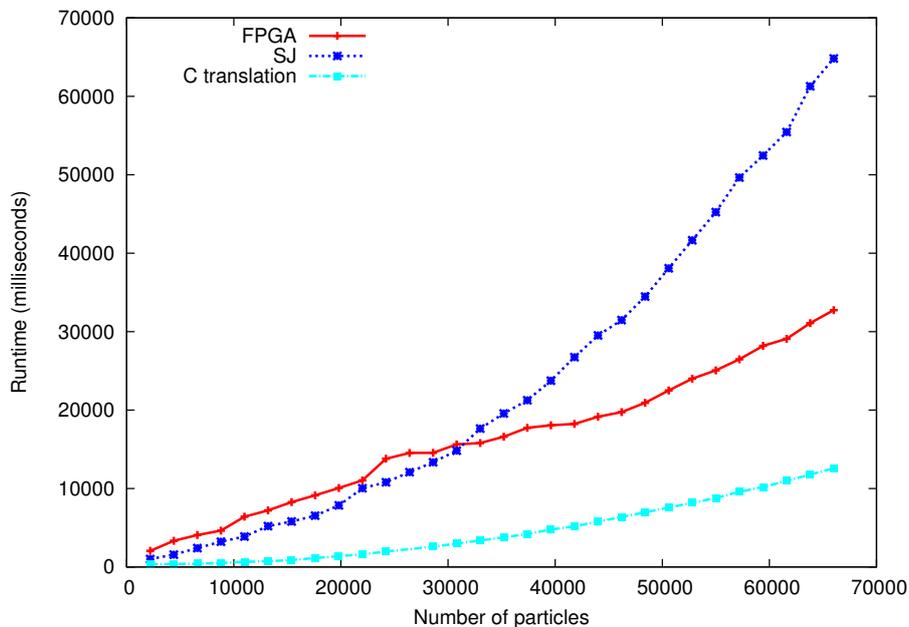


Fig. 5.6: Comparison of SJ+FPGA and its C-translation

We need to keep in mind again that the two programming languages are very different in terms of intended use and design, which we went into details in §3.3.1.

Sources of overhead Fig. 5.7 shows the flow of data during a call to the main computation function, `computeForces()` in SJ. Inputs from SJ are passed to the shared library as arguments to `compute_forces()` in C. Inside `compute_forces()`, inputs are written to and results read from the FPGA's memory. Then the results are forwarded back to the SJ `computeForces()`.

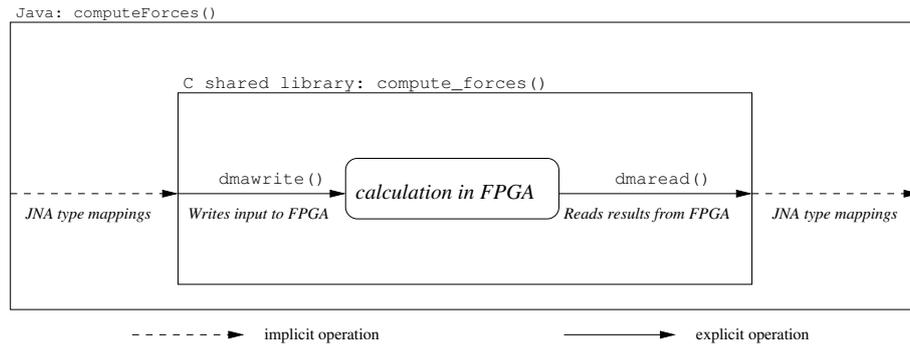


Fig. 5.7: Flow of data between SJ and FPGA in a single `computeForces()` call

In comparing the runtime of SJ+FPGA version and C-translation, the differences between the two are the time spent on conversion to and from the two languages since the communication structure of the program are identical. The conversion includes data type and format translation, as well as copying the data from JVM memory space to main memory.

Fig. 5.8 shows a comparison of time spent in the main computation (ie. `compute_forces()`), and time spent in `computeForces()` which includes all the aforementioned conversion overhead. This microbenchmark was compiled using the same configuration as in the execution of Fig. 5.6. Note that the numbers shown in the graph are average *per call* to `computeForces()`.

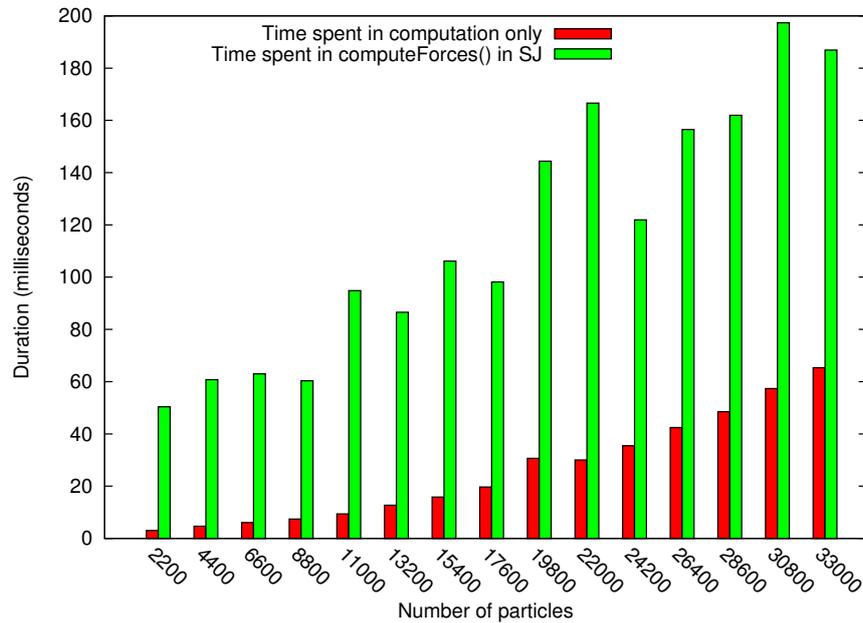


Fig. 5.8: Graph showing the actual execution time in FPGA and the execution time from SJ's perspective

From Fig. 5.8, the overhead (ie. differences of the bars) overshadows the time spent on computation. (green bars corresponds to outer box in Fig. 5.7, and red to inner box) Moreover, with the

more particles used in the simulation, the efficiency shows a slow improvement. Interpolating the results, if we can increase the input size indefinitely, the proportion of overhead would eventually be small enough to be negligible. Whether using that input size is realistic is another question - the duration of the calculation will be very long, judging from the results that computation time with 33000 particles is about a third of the total function call time. The efficiency of this function call is definitely less than 50%.

5.3.3 Comparing with Axel's implementation

In [?], the implementation of n-body simulation for 81920 particles on FPGA is quoted to have $T_{comp} = 5.62s$ and CPU $T_{comp} = 99.3s$, a total of 17.7 times improvement on computation time.

While we wish to compare the performance of our implementation directly, we have a slight difference in specific algorithm details. In this project, we have chosen to use a 2D n-body simulation, ie. particles in the universe we simulate all lies on a 2D plane. This choice is based on previous SJ parallel algorithm works [?, ?]. In the Axel implementation, 3D n-body simulation is used instead.

3D n-body simulation should be seen as an advantage for the FPGA implementation. As shown in Listing 2.5, calculation on an extra axis can be parallelised giving a better performance than CPU. If the parallel efficiency is 100%, we would be seeing 3 times speedup on 3D n-body with FPGA but only 2 times speedup on a 2D n-body simulation.

The speedup we can get from SJ version is The speedup of is C-translation from SJ implementation is on average 5 (See Fig. A.1). and on Fig. 5.6 the results of C-translation interpolates to about 20s. If we take into account the 2D/3D differences above, the runtime of Axel's FPGA implementation $5.62s \times 3/2$ is still quite a lot faster than our c-translation.

It should be reminded that this implementation runs on a different configuration and runs a simplified n-body simulation algorithm. If the *same* FPGA implementation is used, the comparison would be TCP sockets/Session sockets (c-translation) against MPI. We have shown that the Java implementation of MPI, MPJExpress compares competitively with SJ implementation. If the comparison environment is the same, it should be expected that similar results would hold for C-translations and MPI.

5.3.4 Benchmark results conclusion

In conclusion, despite the performance improvement of SJ+FPGA over vanilla SJ implementation of our n-body simulation, the biggest bottleneck of the design lies in the conversion library JNA. The translation of SJ implementation to C has eliminated the need for such *runtime* conversion library, and shows a much improved performance over versions of code that relies on Java, yet the performance still could not match implementation with MPI using the Axel toolchain.

5.4 Summary

In this chapter, we had some discussion on previously planned and alternative application structure. Next, we looked at an *inner product implementation*, which it showed that the JNA direct mapping method is a better way to build cross-language applications such as our implementation. Finally, we showed some benchmark data of SJ against different implementation of the same n-body algorithm, most notably, SJ + FPGA combination, MPJExpress - a pure Java MPI implementation, and a manual C-translation of the SJ + FPGA implementation. Both of SJ + FPGA and C-translation showed big improvements over native SJ, with SJ + FPGA at 2 times speedup and C-translation about 5 times speedup. We also identified the overhead of the design, JNA library, by comparing the duration of the main function call with the main computation time. It showed that only less than a third of the function call was performing useful computation.

Chapter 6

Conclusion

In our design goals outlined in §3.1, we stated that our main design criteria are *efficiency*, *safety* and *readability*. Using SJ, We have successfully shown that all the design goals are met:

Efficiency In our implementation of n-body simulation with SJ, the benchmark result shows with hardware acceleration of FPGA the performance of the simulation are improved. (§5.3.2)

Safety As the n-body simulation is designed in SJ, it is free from communication errors. Furthermore, we have proved that deadlocks Will not happen throughout the execution of the application from the global view of our implementation. (§4.7)

Readability We use SJ as the main design instrument for our implementations on Axel. SJ is a high level language that only exposes a minimal set of primitives for communications and makes use of object-oriented features to structure program design. §2.4

We also looked at a version of our n-body simulation translated to C. §3.4 The implementation shows potential of C as a target language for parallel designs in SJ. A SJ communication primitives library in C was developed as a result of the translation. (§3.4)

We have formalised a multichannel `inwhile` and `outwhile` construct used for designing parallel algorithms. We also derived a definition for well-formed ring topology as part of the formalisation above, and delivered a deadlock-free property for all processes under the topology. This involves a new mechanical proof technique that avoided the use of complex formalism such as global type (multiparty session type) and shared input queue [?, ?] to model the multichannel `inwhile` and `outwhile` semantics.

We wish formalisation of the multichannel SJ constructs will take SJ further in the field of parallel programming, given the extra confidence of a deadlock free prove. The results of our formalisation is built upon a lot of previous and ongoing session types research [?, ?, ?]. Partial session types, sequential composition, well-formedness of process structures are all additions not found in [?].

6.1 Future work

- **Multiparty session types for SJ.** We have shown global communication safety with our n-body simulation as a theoretical proof. If we have multiparty session types for SJ, we would be able to show global communication safety by asserting the property from the SJ framework without a separate proof.
- **Automatic translator.** The SJ primitives library in C and the manual C translation has shown prospects of parallel design with a SJ based approach. Avoiding the huge overhead of *runtime* translation between the two language and instead providing an SJ *runtime* in C is a much preferred approach than mixing language. If further work can extend this approach and automate the translation, we can have best of session types and HPC programming.
- **Generalise approach to GPU.** Of all the three computing elements available on Axel, GPUs are the ones that we have not implemented a SJ version. As described in previous sections, our approach is *designed* to use with different hardware and implementations. This will extend SJ to a wider range of acceleration hardware.
- **Full Session C++.** X10 programming language [?] from IBM is a research language for parallel applications in the *Partitioned Global Address Space* (PGAS) family of languages. PGAS is a parallel programming model that essentially aggregates distributed memory to a global address space and exploits locality of reference in memory access for performance. There are some interest from the SJ community to compare the two languages [?, ?]. However, an interesting feature from an implementer point of view in the language is a multi language code generator in X10 compiler. X10 can generate Java code and C++ code in their compiler¹. SJ uses the same *polyglot compiler framework* as the X10 project, and it looks like it is possible to add a similar extension to the SJ compiler. If we are able to develop a representation of sessions in C++, and incorporate this information in the runtime environment, then we could possibly get a complete session-based C++ language.
- **Integrating SJ into heterogeneous cluster toolchain** As it stands, SJ cannot be used directly as a part of the Axel toolchain.

The toolchain uses MPI as the inter-node communication tool. All the code for each hardware accelerator (eg. `fpga.c`, `gpu.c`, `cpu.c`) is compiled separately by their respective building tool, then linked to a single executable by the MPI compiler `mpicc`. `mpirun` is then invoked on the executable and the partitioning information (an XML configuration file we briefly mentioned in §2.5.2) is supplied as part of the arguments.

To fit SJ into the toolchain, it is best that we use SJ as a communication tool to replace MPI. The main application should be written in SJ, but since SJ/Java cannot perform a link operation with natively compiled code, there are two possible ways to proceed:

1. Adopt the methods described in this project and use JNA to allow SJ code interoperate with the precompiled executables. Alternatively go one step further by translating the

¹<http://docs.codehaus.org/display/XTENLANG/X10+Compiler+Overview>

SJ code to C as in §3.4. Note that in this arrangement, the architecture described is exactly **SJ as a coordinator** we discussed in §5.1.1.

Dynamic load balancing could be a feature in this arrangement, eliminating the need to specify the partition split before execution.

2. Session C++ we have just proposed would be an ideal candidate with both *performance* and *communication safety*: A session types based language that can be used **natively** with the heterogeneous components as an inter-component coordinator, and a **communication-safe** message passing framework for inter-node communication.

Bibliography

- [1] C. Austin and M. Pawlan. *Advanced Programming for the Java 2 Platform with CD-ROM*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2000.
- [2] Axel project website. http://cc.doc.ic.ac.uk/projects/prj_axel/. Accessed on 2/6/2010.
- [3] A. Bejleri, R. Hu, and N. Yoshida. Session-based programming for parallel algorithms: Expressiveness and performance. In *PLACES'09*, 2009. http://www.doc.ic.ac.uk/~ab406/parallel_algorithms.html.
- [4] A. Bejleri and N. Yoshida. Synchronous multiparty session types. *Electron. Notes Theor. Comput. Sci.*, 241:3–33, 2009.
- [5] L. Cardelli and P. Gardner. Membrane computing and biologically inspired process calculi. Slides available at [http://www.lucacardelli.name/Talks/2009-12-04PiintheSky\(ImperialLecture\).pdf](http://www.lucacardelli.name/Talks/2009-12-04PiintheSky(ImperialLecture).pdf), 2009. Accessed on 13/1/2010.
- [6] P. Collingbourne and P. Kelly. Inference of session types from control flow. In *FESCA, ENTCS*. Elsevier, 2008. To appear.
- [7] CUDA homepage. http://www.nvidia.co.uk/object/cuda_what_is.html. Accessed on 13/1/2010.
- [8] M. Dezani-Ciancaglini, U. de'Liguoro, and N. Yoshida. On progress for structured communications. In G. Barthe and C. Fournet, editors, *TGC*, volume 4912 of *Lecture Notes in Computer Science*, pages 257–275. Springer, 2007.
- [9] M. Dezani-Ciancaglini, S. Drossopoulou, D. Mostrous, and N. Yoshida. Objects and session types. *Inf. Comput.*, 207(5):595–641, 2009.
- [10] M. Fähndrich, M. Aiken, C. Hawblitzel, O. Hodson, G. C. Hunt, J. R. Larus, , and S. Levi. Language Support for Fast and Reliable Message-based Communication in Singularity OS. In *EuroSys'06*, ACM SIGOPS, pages 177–190, 2006.
- [11] libffi: A Portable Foreign Function Interface Library. <http://sourceware.org/libffi/>. Accessed on 30/5/2010.

- [12] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design patterns: elements of reusable object-oriented software*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1995.
- [13] A. D. Gordon. A calculus for cryptographic protocols: The spi calculus. *Information and Computation*, 148:36–47, 1999.
- [14] K. Honda, V. T. Vasconcelos, and M. Kubo. Language primitives and type disciplines for structured communication-based programming. In *ESOP'98*, volume 1381, pages 22–138, 1998.
- [15] K. Honda, N. Yoshida, and M. Carbone. Multiparty asynchronous session types. In *In Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008*, pages 273–284. ACM Press, 2008.
- [16] R. Hu, D. Kouzapas, O. Pernet, N. Yoshida, and K. Honda. Type-safe eventful sessions in Java. In *ECOOP '10*, 2010. To appear.
- [17] R. Hu, N. Yoshida, and K. Honda. Session-based distributed programming in java. In J. Vitek, editor, *ECOOP*, volume 5142 of *Lecture Notes in Computer Science*, pages 516–541. Springer, 2008.
- [18] Jacuzzi homepage. <http://jacuzzi.sourceforge.net>. Accessed on 13/1/2010.
- [19] JCUDA homepage. <http://www.jcuda.org>. Accessed on 13/1/2010.
- [20] JNA homepage. <https://jna.dev.java.net/>. Accessed on 13/1/2010.
- [21] Y. Kryftis. Session-based programming for message-passing-based parallel algorithms. Master's thesis, Imperial College London, 2009.
- [22] T. G. Mattson, R. Van der Wijngaart, and M. Frumkin. Programming the intel 80-core network-on-a-chip terascale processor. In *SC '08: Proceedings of the 2008 ACM/IEEE conference on Supercomputing*, pages 1–11, Piscataway, NJ, USA, 2008. IEEE Press.
- [23] R. Milner. *Communication and concurrency*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1989.
- [24] R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes, i. *Inf. Comput.*, 100(1):1–40, 1992.
- [25] G. Moore. Cramming more components onto integrated circuits. *Proceedings of the IEEE*, 86(1):82–85, Jan. 1998.
- [26] MPI: A Message-Passing Interface Standard, Version 2.1. <http://www.mpi-forum.org/docs/mpi21-report.pdf>, 2008. Accessed on 13/1/2010.
- [27] M. Neubauer and P. Thiemann. An implementation of session types. In *In PADL, volume 3057 of LNCS*, pages 56–70. Springer, 2004.

- [28] M. Sackman and S. Eisenbach. Session Types in Haskell: Updating Message Passing for the 21st Century. Technical report, June 2008.
- [29] D. Sangiorgi and D. Walker. *PI-Calculus: A Theory of Mobile Processes*. Cambridge University Press, New York, NY, USA, 2001.
- [30] The computer language benchmark game. <http://shootout.alioth.debian.org>. Accessed on 13/1/2010.
- [31] SJ homepage. <http://www.doc.ic.ac.uk/~rhu/sessionj.html>. Accessed on 13/1/2010.
- [32] K. H. Tsoi and W. Luk. Axel: a heterogeneous cluster with FPGAs and GPUs. In *FPGA '10: Proceedings of the 18th annual ACM/SIGDA international symposium on Field programmable gate arrays*, pages 115–124, New York, NY, USA, 2010. ACM.
- [33] X10 homepage. <http://x10-lang.org/>. Accessed on 11/6/2010.
- [34] N. Yoshida, P.-M. Deniérou, A. Bejleri, and R. Hu. Parameterised multiparty session types. In C.-H. L. Ong, editor, *FOSSACS*, volume 6014 of *Lecture Notes in Computer Science*, pages 128–145. Springer, 2010.
- [35] N. Yoshida and V. T. Vasconcelos. Language primitives and type discipline for structured communication-based programming revisited: Two systems for higher-order session communication. *Electr. Notes Theor. Comput. Sci.*, 171(4):73–93, 2007.

Appendix A

Appendix

A.1 Java Native Interface (JNI) example

This example is a C program from the official JNI Tutorial [?, Chapter 5].

```
1  #include <jni.h>
2  ...
3
4  JNIEXPORT jbyteArray JNICALL Java_ReadFile_loadFile(JNIEnv * env, jobject jobj,
5      jstring name)
6  {
7      caddr_t m;
8      jbyteArray jb;
9      jboolean iscopy;
10     struct stat finfo;
11     const char *mfile = (*env)->GetStringUTFChars(env, name, &iscopy);
12     int fd = open(mfile, O_RDONLY);
13
14     if (fd == -1) printf("Could not open %s\n", mfile);
15     lstat(mfile, &finfo);
16     m = mmap((caddr_t) 0, finfo.st_size, PROT_READ, MAP_PRIVATE, fd, 0);
17     if (m == (caddr_t)-1) {
18         printf("Could not mmap %s\n", mfile);
19         return(0);
20     }
21     jb = (*env)->NewByteArray(env, finfo.st_size);
22     (*env)->SetByteArrayRegion(env, jb, 0, finfo.st_size, (jbyte *)m);
23     close(fd);
24     (*env)->ReleaseStringUTFChars(env, name, mfile);
25     return (jb);
26 }
```

Listing A.1: Example C function that uses JNI from [?]

JNI provides a complete mapping between native datatype and Java datatype, such as `jstring` and `jbyteArray`. The `JNIEnv *env` pointer is the core feature of JNI, which gives the native program access to the execution environment and data in the JVM. It also helps keeping track of references for the automatic garbage collection mechanism amongst other metadata, therefore `ReleaseStringUTFChars` is issued to notify the garbage collector before the function in Listing A.1 returns.

The *Java Development Kit* (JDK) comes with a tool `javah` to generate C header and stub files as in Listing A.1, but working with code in JNI is still cumbersome and generally considered difficult.

A.2 Java Native Access (JNA) example

```

1 /* gcc -shared -o libexample.so libexample.c */
2 int sum(int x, int y)
3 {
4     return x + y;
5 }

```

Listing A.2: Source of shared library *libexample* in C

```

1 package libexample;
2 import com.sun.jna.Library;
3
4 public interface LibExample extends Library {
5     int sum(int x, int y);
6 }

```

Listing A.3: Java interface for *libexample*

```

1 package libexample;
2 import com.sun.jna.Native;
3
4 public class Example {
5     public static void main(String args[]) throws Exception {
6         LibExample libexample = (LibExample) Native.loadLibrary(
7             System.getProperty("user.dir")
8             + "/libexample.so", LibExample.class);
9         System.out.println("Sum is "+libexample.sum(42, 77));
10    }
11 }

```

Listing A.4: Java code that uses the *sum* function in *libexample*

```

1 package SJExample;
2 import sessionj.runtime.*;
3 import sessionj.runtime.net.*;

```

```

4
5 public class Client {
6     final noalias protocol p_client { cbegin.!<int>.!<int>.(int) }
7
8     public void run(int port) {
9         final noalias SJService svc = SJService.create(
10            p_client, "localhost", port);
11         final noalias SJSocket sock;
12
13         try (sock) {
14             sock = svc.request();
15             sock.send(42);
16             sock.send(77);
17             int result = sock.receiveInt();
18             System.out.println("Server replies: "+result);
19         } catch (SJIncompatibleSessionException ise) {
20             System.err.println("[C] Non-dual behaviour: " + ise);
21         } catch (SJIOException sioe) {
22             System.err.println("[C] Communication error: " + sioe);
23         } finally { /* Close socket */ }
24     }
25
26     public static void main(String argv[] throws Exception {
27         int port = Integer.parseInt(argv[0]);
28         Client client = new Client();
29         client.run(port);
30     }
31 }

```

Listing A.5: SJ code similar to `Example` class in the JNA-Java example

```

1 package SJExample;
2 import sessionj.runtime.*;
3 import sessionj.runtime.net.*;
4
5 import com.sun.jna.Native;
6 import libexample.LibExample;
7
8 public class Server {
9     final noalias protocol p_server { sbegin.(int).(int).!<int> }
10
11     public void run(int port) {
12         final noalias SJServerSocket svr;
13         final noalias SJSocket sock;
14
15         /**
16          * Get an instance of LibExample
17          */
18         String abspath = System.getProperty("user.dir")+"/libexample.so";

```

```

19     LibExample libexample = (LibExample) Native.loadLibrary(
20         abspath, LibExample.class);
21
22     try (svr) {
23         svr = SJServerSocketImpl.create(p_server, port);
24         try (sock) {
25             sock = svr.accept();
26             int x = sock.receiveInt();
27             int y = sock.receiveInt();
28             int result = libexample.sum(x, y);
29             sock.send(result);
30         } catch (SJIncompatibleSessionException ise) {
31             System.err.println("[S] Non-dual behaviour: " + ise);
32         } catch (SJIOException sioe) {
33             System.err.println("[S] Communication error: " + sioe);
34         }
35
36     } catch (SJIOException sioe) {
37         System.err.println("[S] Communication error: " + sioe);
38     } finally { /* Close socket */ }
39 }
40
41 public static void main(String argv[] throws Exception {
42     int port = Integer.parseInt(argv[0]);
43     Server server = new Server();
44     server.run(port);
45 }
46 }

```

Listing A.6: SJ code similar to Example class in the JNA-Java example

This is a full code listing of a JNA example, where the Java code invokes a function in a C shared library (libexample) to add two numbers. libexample is similar to SumServer/Client (Listing 2.1).

A.3 Comparison of SJ and C-translation implementation

This section compares SJ implementation and its C-translation of the same SJ implementation. There is almost a line-by-line correspondence between the two versions.

```

1 //
2 //$ bin/body left-port body-host right-port input-size
3
4 ... // includes, macros
5
6 volatile uint32_t* fpgaReg;
7 volatile uint8_t* fpgaMem;
8

```

```

9  int main(int argc, char **argv)
10 {
11     ... // Variable declarations
12
13     signal(SIGPIPE, &sigpipe_handler);
14     signal(SIGSEGV, &sigsegv_handler);
15
16     prepare();
17
18     // Protocol: cbegin.?(int).![!<Particle[]>]*
19     next_fd = client_socket(argv[2], atoi(argv[3]));
20     // Protocol: sbegin.!(int).?[?<Particle[]>]*
21     prevnode_fd = server_socket(atoi(argv[1]));
22     prev_fd = accept_connection(prevnode_fd);
23
24     // # of nodes
25     recv_int(next_fd, &nr_of_nodes); // ?(int)
26     ++nr_of_nodes;
27     send_int(prev_fd, &nr_of_nodes); // !<int>
28
29     size = atoi(argv[4]);
30
31     particles = (particle_t *) malloc(sizeof(particle_t)*size);
32     temp_particles = (particle_t *) malloc(sizeof(particle_t)*size);
33     pvs = (particlev_t *) malloc(sizeof(particlev_t)*size);
34
35     init(particles, pvs, size);
36
37     outer_loop_index = 0;
38     OUTWHILE(inwhile(&prev_fd, 1), &next_fd, 1) {
39
40         // This round
41         memcpy(temp_particles, particles, size * sizeof(particle_t));
42
43         // Pump particles through the ring
44         OUTWHILE(inwhile(&prev_fd, 1), &next_fd, 1) {
45
46             // !<Particle[]>, Send particles into ring
47             send_particles(next_fd, temp_particles, size);
48             compute_forces(particles, temp_particles, pvs, size);
49             // ?<Particle[]>, Receive from the other end of ring
50             recv_particles(prev_fd, temp_particles);
51
52         }
53
54         compute_forces(particles, temp_particles, pvs, size);
55         compute_positions(particles, pvs, outer_loop_index, size);
56
57         ++outer_loop_index;

```

```

58     }
59
60     // These are done by SJ automatically
61     close_socket(prev_fd); close_socket(prevnode_fd); close_socket(next_fd);
62     free(particles); free(temp_particles); free(pvs);
63
64     finish(); // Finalise FPGA etc.
65
66     return EXIT_SUCCESS;
67 }

```

Listing A.7: C translation of SJ n-body Worker node

```

1 //sessionjc - cplib : ./jna.jarsrc/nbody/Body.sj - dlib//
2
3
4 ... // imports, package declarations
5
6 public class Body {
7     // #nodes
8     final noalias protocol p_prev { sbegin.<int>.[?[(Particle[])]*]* }
9     final noalias protocol p_next { ^(p_prev) }
10    NBody nbody;
11
12    public Body(NBody nbody) {
13        this.nbody = nbody;
14    }
15
16    public void run( ... ) throws ClassNotFoundException {
17        ... // Variable declarations
18
19        nbody.prepare();
20
21        try (prevNode) {
22            prevNode = SJServerSocketImpl.create(p_prev, listenPort);
23            try (prev,next) {
24
25                next = nextNode.request();
26                prev = prevNode.accept();
27
28                // # of nodes
29                nodeIndex = next.receiveInt();
30                prev.send(nodeIndex+1);
31
32                particles = new Particle[size];
33                pvs = new ParticleV[size];
34
35                nbody.init(particles, pvs, nodeIndex);
36

```

```

37         next.outwhile(prev.inwhile()) {
38
39             // This round
40             Particle[] tempParticles = new Particle[size];
41             System.arraycopy(particles, 0, tempParticles, 0, size);
42
43             // Pump particles through ring
44             next.outwhile(prev.inwhile()) {
45                 next.send(tempParticles);
46                 nbody.computeForces(particles, tempParticles, pvs);
47                 tempParticles = (Particle[]) prev.receive();
48             }
49
50             nbody.computeForces(particles, tempParticles, pvs);
51             nbody.computePositions(particles, pvs, i);
52
53             ++i;
54         }
55     } catch (SJIncompatibleSessionException ise) {
56     } catch (SJIOException sioe) {}
57 } catch (SJIOException sioe) {
58 } finally { // Close socket
59     nbody.finish(); // Finalise FPGA etc.
60 }
61 }
62 }
63 }
64 }

```

Listing A.8: SJ n-body Worker

A.4 SJ + FPGA speedup over SJ implementation

Fig A.1 shows the speedup calculated from runtime results from Fig 5.5. We could see that the speedup increases as we increase the problem size (number of particles).

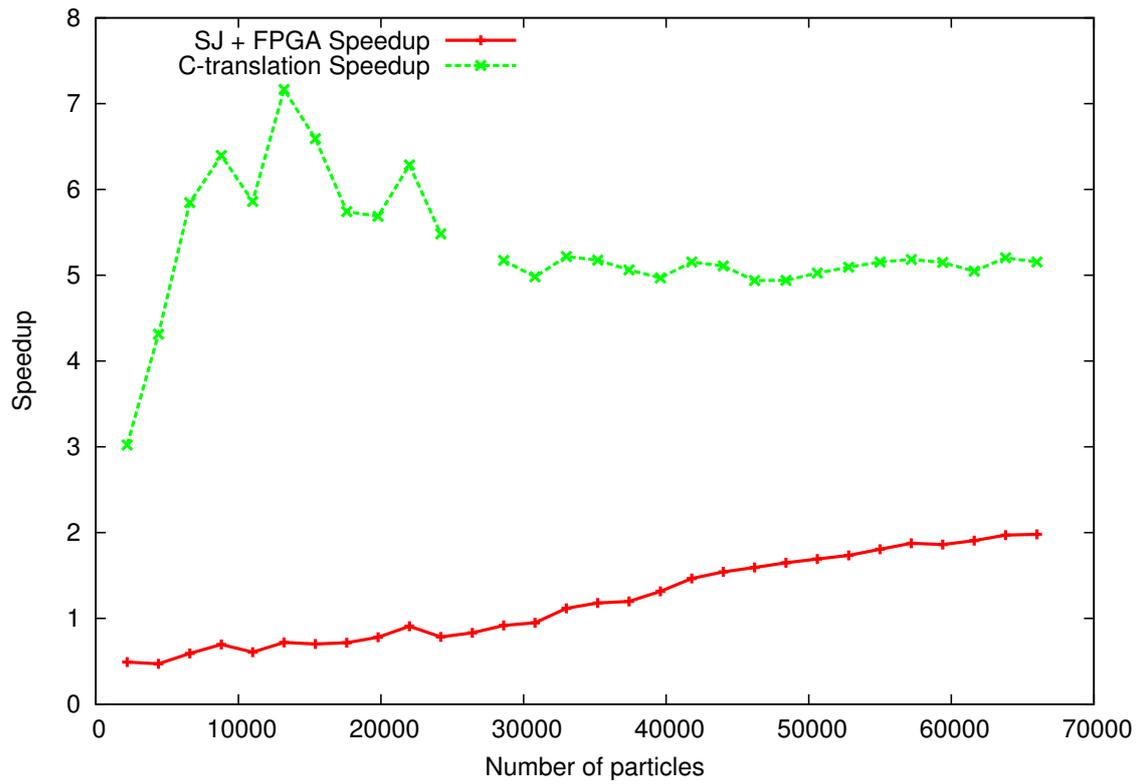


Fig. A.1: Speedup of SJ + FPGA, C translation vs. SJ in 11 nodes over 5 iterations