

Formalising Java RMI with Explicit Code Mobility

Alexander Ahern and Nobuko Yoshida

Department of Computing, Imperial College London

This paper presents a Java-like core language with primitives for object-oriented distribution and explicit code mobility. We apply our formulation to prove the correctness of several optimisations for distributed programs. Our language captures crucial but often hidden aspects of distributed object-oriented programming, including object serialisation, dynamic class downloading and remote method invocation. It is defined in terms of an operational semantics that concisely models the behaviour of distributed programs using machinery from calculi of mobile processes. Type safety is established using invariant properties for distributed runtime configurations. We argue that primitives for explicit code mobility offer a programmer fine-grained control of type-safe code distribution, which is crucial for improving the performance and safety of distributed object-oriented applications.

Categories and Subject Descriptors: D.3.1 [**Programming Language**]: Formal Definitions and Theory; D.3.3 [**Programming Language**]: Language Constructs and Features; F.3.2 [**Theory of Computation**]: Semantics of Programming Languages

General Terms: Languages, Theory

Additional Key Words and Phrases: Distribution, Java, RMI, Types, Optimisation, Runtime, Code mobility

1. INTRODUCTION

Language features for distributed computing form an important part of modern object-oriented programming. It is now common for different portions of an application to be geographically separated, relying on communication via a network. Distributing an application in this way confers many advantages to a programmer such as resource sharing, load balancing, and fault tolerance. However this comes at the expense of increased complexity for that programmer, who must now deal with concerns—such as network failure—that did not occur in centralised programs.

Remote procedure call mechanisms attempt to simplify such engineering practice by providing a seamless integration of network resource access and local procedure calls, offering the developer a programming abstraction familiar to them. Java Remote Method Invocation [Microsystems Inc. 2005] (RMI) is a widely adopted remote procedure call implementation for the Java platform, building on the customisable class loading system of the underlying language to further hide distribution from the programmer. When objects are passed as parameters to remote methods, if the provider of that method does not have the corresponding class file, it may attempt to obtain it from the sender. Such code mobility is important as it reduces the need for strong coupling between communicating parties, while preserving the type safety of the system as a whole.

The implicit code mobility in RMI allows almost transparent use of remote ob-

Author's address: Department of Computing, Imperial College London, 180 Queen's Gate, London SW7 2AZ, UK.

, Vol. V, No. N, Month 20YY, Pages 1–0??.

jects and services. However when rigorously analysing the dynamics of distributed programs, or when providing programmers with source-level control over code distribution [Christ 2000], it becomes essential to model their behaviour explicitly. This is because elements such as distribution, network delay and partition crucially affect the behaviour and performance of programs and systems. As an example, communication-oriented RMI optimisations, often called *batched futures* [Bogle and Liskov 1994] or *aggregation* [Yeung and Kelly 2003; Yeung 2004], use code distribution as their central element. To analyse these optimisations formally, or to make the most of them in applications, explicit primitives for code mobility are essential.

This paper proposes a Java-like distributed object-oriented core language with communication primitives (RMI) and distributed runtime. The formalism exposes hidden runtime concerns such as code mobility, class downloading, object serialisation and communication. The operational semantics concisely models this behaviour using machinery from calculi of mobile processes [Milner et al. 1992; Sangiorgi 1992; Honda and Tokoro 1991]. One highlight is the use of a *linear* type discipline [Kobayashi et al. 1996; Honda 1996; Yoshida et al. 2001] to ensure correctness of remote method calls. Another is the application of several invariant properties. These are conditions that hold during execution of distributed programs, and they allow type safety to be established.

Our language supports *explicit code mobility* by providing primitives that allow programs to communicate fragments of code—closely related to closures in functional languages—for later execution. This subsumes the standard serialisation mechanism by sending not only object graphs but also executable code. Code passing offers a programmer fine-grained control of type-safe code distribution, improving the safety and performance of their distributed applications. For example, a program fragment accessing a resource remotely could be frozen into a closure. This code could then be passed to the remote site, co-locating it with that resource. This effectively turns remote accesses into local accesses, reducing latency and increasing available bandwidth [Christ 2000; Bogle and Liskov 1994; Yeung and Kelly 2003; Yeung 2004].

As an application of our formalism, we show that the RMI *aggregation* optimisations proposed in [Yeung and Kelly 2003; Yeung 2004] are type- and semantics-preserving. The generality of the primitive we introduce plays an essential role in this analysis: one optimisation relies on the use of second-order code passing, i.e. passing code that in turn passes code itself. Similar optimisations naturally arise whenever latency and bandwidth are a limiting factor in the performance of distributed programs, suggesting a wide applicability of this primitive in similar endeavours.

We summarise our major technical contributions below.

- (1) Introduction of a core calculus for a class based typed object-oriented programming language with primitives for concurrency and distribution, including RMI, explicit code mobility, thread synchronisation and dynamic class downloading.
- (2) A technique to systematically prove type safety for distributed networks using distributed invariants. Not only are they essential for proving type safety but also they are a useful analytical tool for developing consistent typing rules.
- (3) Justification of several inter-node RMI optimisations employing explicit code

mobility, using a semantically sound syntactic transformation of the language and runtime. The analysis also demonstrates the greater control that explicit code mobility offers to programmers.

In the remainder, Section 2 informally motivates the present work through concrete examples of RMI optimisations. Section 3 introduces the syntax of the language. Sections 4 and 5 respectively discuss the dynamic semantics (reduction) and static semantics (typing) of the language. Section 6 establishes type safety and progress properties using the invariants. Section 7 studies contextual congruence of the core language and applies the theory to justify the optimisations in Section 2. Section 8 discusses related work. Section 9 concludes the paper with further topics.

This paper is a full version of [Ahern and Yoshida 2005b], with complete definitions and detailed proofs. The emphasis is on illustrating correctness of formalising RMI and a use of the distributed invariants for proving the type safety, and developing a theory of an observational congruence for a distributed Java. The present version also gives more examples on dynamic semantics of RMI and comparisons with related work.

2. MOTIVATION: REPRESENTING AND JUSTIFYING RMI OPTIMISATION

The RMI optimisations introduced in this section are used as running examples, culminating in their justification by the behavioural theory in § 7. These are (arguably) typical inter-node optimisations of distributed object-oriented programs. Just as inter-procedure or inter-module optimisations are hard to analyse, RMI optimisation poses a new challenge to the semantic analysis of distribution. They also motivate the use of explicit code mobility for fine-grained control of distributed behaviour and to improve performance.

Original RMI program 1. In optimisations for sequential languages, we can aim to improve execution times by removing redundancy and ensuring our programs exploit features of the underlying hardware architecture. In distributed programs these are still valid concerns, but other significant optimisations exist, in particular how latency and bandwidth overheads can be reduced. One typical example of this sort, centring on Java RMI [Flanagan 2000] but which is generally applicable to various forms of remote communication, is *aggregation* [Bogle and Liskov 1994; Yeung and Kelly 2003; Yeung 2004]. We explain this idea using the simple program in Listing 1.

```

1  int m1(RemoteObject r, int a) {
2      int x = r.f(a);
3      int y = r.g(a, x);
4      int z = r.h(a, y);
5      return z;
6  }
```

Listing 1: Original RMI program 1

This program performs three remote method calls to the same remote object `r` with eight items transferred across the network (counting each parameter and return value as one). `x` is returned as the result of the call to `f` from the remote

server, but is subsequently passed back to the server during the next call. The same occurs with the variable y . These variables are unused by the client, and are merely returned to the remote object r (where they were created) as parameters to the next call. We can immediately see that there is no need for x or y to ever be passed back to the client at all. Hence these three calls can be *aggregated* into a single call, reducing by a factor of three the network latency incurred by the method $m1$ and approximately reducing by a factor of four the amount of data that must be shipped across the network.

This optimisation methodology is implemented in the Veneer virtual Java Virtual Machine (vJVM) [Yeung and Kelly 2003; Yeung 2004], where sequences of adjacent calls to the same remote object are grouped together into an execution *plan* in bytecode format. This is then uploaded to and executed by the server, with the result of the computation being returned to the client. This simple idea—remote evaluation of code [Stamos and Gifford 1990]—can speed up distributed programs significantly, especially when operating across slower networks or when significant amounts of data may be transmitted otherwise. As a concrete example, in [Yeung and Kelly 2003] the authors reported that over a moderate bandwidth and moderate latency ADSL connection, call aggregation yields a speedup over a factor of four for certain examples [Flanagan 2000].

Optimised program 1. Call aggregation implicitly uses code passing: we first collect all the code that can be executed at a remote site and then send it, in one bundle, for execution there. This aspect is hidden as the transfer of bytecode in the implementation of [Yeung and Kelly 2003; Yeung 2004], but requires explicit modelling if one wishes to discuss its properties or justify that it preserves the original program semantics. For this purpose we introduce two primitives, `freeze` and `defrost`. In Listing 2, we illustrate these primitives using the optimised version of the code of Listing 1.


```

1  // Client
2  int mOpt1(RemoteObject r, int a) {
3      thunk<int> t = freeze {
4          int x = r.f(a);
5          int y = r.g(a, x);
6          int z = r.h(a, y);
7          z
8      };
9      return r.run(t);
10 }
11 // Server
12 int run(thunk<int> x) {
13     return defrost(x);
14 }

```

Listing 2: Optimised program 1

Here the client uses the `freeze` expression of the language to create a frozen representation of three calls with a closure of free variable a , sending the resulting “thunk” to the server. `thunk<int>` says the frozen code contains an expression of type `int`. We now make only one call across the network to send the frozen



We annotate call arrows with the method invocation and return arrows with the name of the variable the client will use to store the return value of the method.

Fig. 1: Example optimisation (1)

expression, by `r.run(t)`. When the server receives the thunked code, it evaluates it and returns the result typed by `int` to the client, again across the network. These mimic primitives found in well-known functional languages, for example the quotation and evaluation of code in Scheme, or the higher-order functions found in ML and Haskell.

In Figure 1 we show a diagram of the situation. As can be seen, the original sequence of calls (the paler arrows) requires 6 trips across the network. By aggregating the calls at the server, where they effectively become local, we see that only two trips are required (the thicker arrows).

Original RMI program 2. A more advanced form of communication-oriented optimisation, which reduces latency and uses bandwidth intelligently, is the idea of *server forwarding* [Yeung and Kelly 2003; Yeung 2004]. It takes advantage of the fact that servers typically reside on fast connections, whilst the client-server connection can often be orders of magnitude slower. Consider the program in Listing 3.

```

1  int m2(RemoteObject r1,
2      RemoteObject r2, int a) {
3      int x1 = r1.f1(a);
4      int y1 = r1.g1(a, x1);
5      int z1 = r1.h1(a, y1);
6      int x2 = r2.f2(z1);
7      int y2 = r2.g2(z1, x2);
8      int z2 = r2.h2(z1, y2);
9      return z2;
10 }
```

Listing 3: Original RMI program 2

The results of the first three calls are used as arguments to methods on another remote object `r2` in a second server. It would be better for the first server to communicate directly with the second.




Fig. 2: Example optimisation (2)

Optimised program 2. Server forwarding again uses code passing as an execution mechanism. Listing 4 lists the optimised code of the original program in Listing 3. We use closure passing for representing this optimised code, in which thunked code is *nested*, i.e. we are using higher-order code mobility. Figure 2 gives a diagram of the situation.

```

1  int mOpt2(RemoteObject r1, RemoteObject r2, int a) {
2      thunk<int> t1 = freeze {
3          int x1 = r1.f1(a);
4          int y1 = r1.g1(a, x1);
5          int z1 = r1.h1(a, y1);
6          thunk<int> t2 = freeze {
7              int x2 = r2.f2(z1);
8              int y2 = r2.g2(z1, x2);
9              int z2 = r2.h2(z1, y2);
10             z2;
11         };
12         r2.run(t2);
13     };
14     return r1.run(t1);
15 }

```

Listing 4: Optimised program 2

Original RMI program 3. The semantics of RMI is different from normal, local method invocation. Passing a parameter to a remote method (or accepting a return value) can involve many operations hidden from the end-user; these runtime features make automatic semantic-preserving optimisation of RMI much harder, in particular, when calls contain *objects as arguments*. To observe this, let us change

the type of `a` from `int` to class `MyObj` as in the code in Listing 5:

Here we have two cases:

- (1) `MyObj` is *remote* i.e. when `MyObj` implements the `java.rmi.Remote` interface and is therefore remotely callable. In this situation, `a` is effectively passed by *reference*.
- (2) `MyObj` is *local* i.e. when `MyObj` is not remotely callable (it does not implement the `Remote` interface), `a` is automatically *serialised* and passed to the server where it is automatically *deserialised*. In this situation, `a` is effectively passed by *value*.

```

1  int m3(RemoteObject r, MyObj a) {
2      int x = r.f(a);
3      int y = r.g(a, x);
4      int z = r.h(a, y);
5      return z;
6  }

```

Listing 5: Original RMI program 3

Sending a serialised value to a remote consumer can be thought of as passing an object by *value*. Informally, the serialisation process explores the graph under an object in local memory, copying all objects directly or indirectly referred to. When passing such local objects as parameters to remote methods, the Java RMI system automatically performs this copying.

Consider the method `m3` above: if the call `r.f` performs an operation that side-effects the parameter `a`, then in the original program this side-effect is lost. The version of `a` supplied to the next method `r.g` is still just a copy of the initial `a` held in the client's memory, which has not changed. If we naïvely apply code passing optimisations to the problem, we might rewrite method `m3` to look a lot like `mOpt1`. Unfortunately now the next call `r.g` no longer has a copy of the original `a` to work on: it instead receives the version modified by `r.f`, potentially altering the meaning of the program and rendering the optimisation incorrect.

By applying explicit serialisation we can simulate the original program behaviour. By insisting each method call in the server operates on a fresh copy of the original `a`, we regain correctness as is shown below.

Optimised program 3. We show the case when `MyObj` is a local class. If there are no call-backs from the server to the client (discussed next), then the original RMI program has the same meaning as passing the code in Listing 6. First the client creates three copies of serialised object `a` by applying the explicit serialisation operator `serialize`. We write `serialize` as shorthand for the idiom in Java that involves writing objects to an instance of `ObjectOutputStream`. The server immediately deserialises the arguments, creating three independent object graphs, thus avoiding problems with methods that alter their parameters (we write `deserialize` in place of reading from an `ObjectInputStream`). In the code in Listing 6, the declaration `ser<MyObj> b1` says that `b1` is a serialised representation of an object of class `MyObj`.

```

1  int mOpt3(RemoteObject r, MyObj a) {
2      ser<MyObj> b1 = serialize(a);
3      ser<MyObj> b2 = serialize(a);
4      ser<MyObj> b3 = serialize(a);
5      thunk<int> t = freeze {
6          int x = r.f(deserialize(b1));
7          int y = r.g(deserialize(b2), x);
8          int z = r.h(deserialize(b3), y);
9          z
10     };
11     return r.run(t);
12 }

```

Listing 6: Optimised program 3

Two further problems. We have seen that code passing primitives can help us to cleanly represent communication-based optimisation of RMI programs. Analysis of the code above immediately suggests two further problems that must be addressed.

- (1) *Sharing between objects and call-backs:* the above copying method should not be applied naïvely, since marshaling should preserve sharing between objects. It also may not be applicable if a call by the client to the server results in the server calling the client.
- (2) *Overhead of class downloading:* if the server location does not contain the byte-code for `MyObj`, RMI automatically invokes a class downloading process to obtain the class from the network. In addition, verifying that the received class is safe to use may require the downloading of many others (such as all superclasses of `MyObj` and classes mentioned in method bodies and so on), which may incur many trips across the network, increasing the risk of failures and adding latency.

To illustrate the first problem, consider the following simple code with `r` remote and `x` and `y` local:

```

1  x.f = y; r.h(x, y);

```

The content of `y` is shared with `x` in the original code, but if we apply the copying method then the server creates independent copies of `x` and `y`, breaking the original sharing structure.

For the second point of (1), imagine that the body of remote method `f` invoked at line 2 of the original program involves some communication back to the local site. Then it is possible for the value of `a` to be modified at the client side and so the optimised program is no longer correct: because in our optimised program, `a` is serialised in line 4 before `r.f` is performed, any effect that a call-back would have on `a` is lost, when it *should* be visible to the call `r.g`.

The second problem, class downloading, is more subtle from the communication-based optimisation viewpoint. Although the aim of this optimisation is to reduce the number of trips across the network, if there is a deep inheritance hierarchy above

MyObj, sending code may not yield the performance benefit that the programmer expects. This is because many requests over the network may be required to obtain all the required classes.

As an example, if MyObj has a chain of n -superclasses such that $\text{MyObj} <: \text{MyObj2} <: \dots <: \text{MyObjn}$, and none of these are present at the server, there are at least n class downloads even with “verification off” in the framework of type safe dynamic linking [Liang and Bracha 1998; Qian et al. 2000]. With “verification on”, this could be even more.

These hidden features of RMI make reasoning about the behaviour of a program, and establishing a clear optimisation, hard.

Challenges. Having provided the source-level presentation of several features necessary to discuss RMI optimisations, we may ask the following questions:

- Q1. How can we precisely model this dynamic runtime behaviour, including code passing, serialisation and class downloading?
- Q2. How can we verify the correctness of the optimised code, in the sense that the original code and the optimised code have the same contextual behaviour?
- Q3. Having studied the optimisations above, can we improve our code mobility primitives to make them generally useful to application programmers?

A satisfactory solution to Q1. is a prerequisite for Q2. due to the interleaving of communication events which affect the observational behaviour of distributed programs. Various elements inherent in distributed computing make the semantic correctness of optimisations more subtle than in the sequential setting. The behaviour, hence the final answer, may differ depending on sharing of objects, timing and class downloading strategies, as well as network failure. In our paper, Q1. will be answered by giving a clean formal semantics for distributed object-oriented features usually hidden from a programmer. We shall distill key runtime features, including class downloading and serialisation, so that important design choices (for example various class downloading and code mobility mechanisms) can be easily reflected in the semantics. Q2. will be answered by semantic justification of the above optimisations using the theory of mobile processes [Milner et al. 1992; Sangiorgi 1992; Honda and Tokoro 1991]. For Q3., we summarise our proposal below.

Optimised program 4. Class downloading is a fundamental mechanism in distributed object-oriented programming. Yet so far we have treated it as a behind-the-scenes feature and left it as an implementation detail. However, by augmenting our primitives with a mechanism to control class downloading, a programmer is able to write down different strategies explicitly. This explicit control allows us to mitigate some of the problems class downloading induces that were explained in the previous section. For example, to represent one basic strategy of class downloading, we attach the tag **eager** to **freeze** in the original code 3 in Listing 5.

```

1  int mOpt4(RemoteObject r, MyObj a) {
2      ... // as in mOpt3
3      thunk<int> t = freeze[eager] {
4          ... // as in mOpt3
5      }

```

Listing 7: Optimised program 4

The tag `eager` in `freeze[eager]` controls the amount of class information sent along with the body of the thunk by the user. With `eager`, the code is automatically frozen together with all classes that *may* be used. In the above case all classes appearing in `MyObj` and all their superclasses are shipped together with the code (see § 4.6 for the definition). Another option is for the user to select `lazy` which essentially leaves class downloading to the existing RMI system. Further the user might write a list of specific classes $\vec{C}$ to be shipped. For example, the following program is able to notice when it is in a high latency situation and act accordingly.

```

1 // Client
2 thunk<int> t;
3 if (pingTime() > 1000) { // milliseconds
4   t = freeze[eager] {...};
5 } else {
6   t = freeze[lazy] {...};
7 }
```

Listing 8: Optimised program 5

If we imagine that the latency is very high, then it may be the case that the time to iteratively download all the superclasses exceeds the actual execution time of the frozen code being sent to the server. Because of this, the program is able to switch to the `eager` mode of class downloading, allowing improved performance. Moreover, from a point of view of failure there are fewer trips across the network with the eager policy, reducing the risk of a transient problem, such as a temporary network partition, disrupting the class downloading process.

The formal semantics for both implicit and explicit code mobility is given from the next section as part of the core language.

3. LANGUAGE

3.1 User syntax

The syntax of the core language, which we call DJ, is an extension of FJ [Igarashi et al. 2001] and MJ [Bierman et al. 2003], augmented with basic primitives for distribution and code-mobility, along with concurrent programming features that should be familiar to Java programmers. The syntax comes in two forms, and is given in Figure 3. The first form is called *user syntax*, and corresponds to terms that can be written by a programmer as source code. The second form is called *runtime syntax*. It extends the user syntax with constructs that only appear during program execution, and these are distinguished in the figure by placing them in shaded regions. We briefly discuss each syntactic category below.

Types. T and U range over types for expressions and statements, which are explained in § 5. C, D, F range over class names. $\vec{f}$ denotes a vector of fields, and $\vec{T}\vec{f}$ is short-hand for a sequence of typed field declarations: $T_1f_1; \dots; T_nf_n$. We assume sequences contain no duplicate names, and apply similar abbreviations to other sequences with ϵ representing the empty sequence. $T \rightarrow U$ denotes an *arrow*

Syntax occurring only at runtime appears in shaded regions.

(Types)	$T ::= \text{boolean} \mid \text{unit} \mid C \mid T \rightarrow U$
(Returnable)	$U ::= \text{void} \mid T$
(Classes)	$L ::= \text{class } C \text{ extends } D \{ \vec{T}\vec{f}; K \vec{M} \}$
(Constructors)	$K ::= C(\vec{T}\vec{f}) \{ \text{super}(\vec{f}); \text{this}.\vec{f} = \vec{f} \}$
(Methods)	$M ::= U \ m(C \ x) \{ e \}$
(Expressions)	$e ::= v \mid x \mid \text{this} \mid \text{if } (e) \{ e \} \text{ else } \{ e \} \mid \text{while } (e) \{ e \} \mid e.f \mid e; e$ $\mid x = e \mid e.f = e \mid \text{new } C(\vec{e}) \mid e.m(e) \mid T \ x = e$ $\mid \text{return } e \mid \text{return} \mid \text{freeze}[t](T \ x) \{ e \} \mid \text{defrost}(e; e) \mid \text{fork}(e)$ $\mid \text{sync } (e) \{ e \} \mid e.\text{wait} \mid e.\text{notify} \mid e.\text{notifyAll} \mid \text{new } C^l(\vec{e})$ $\mid \text{download } \vec{C} \text{ from } l \text{ in } e \mid \text{resolve } \vec{C} \text{ from } l \text{ in } e \mid \text{await } c$ $\mid \text{sandbox } \{ e \} \mid \text{insync } o \{ e \} \mid \text{ready } o \ n \mid \text{waiting}(c) \ n \mid \text{Error}$
(Tags)	$t ::= \text{eager} \mid \text{lazy} \mid \vec{C}$
(Values)	$v ::= \text{true} \mid \text{false} \mid \text{null} \mid () \mid o \mid \lambda(T \ x).(\nu \vec{u})(l, e, \sigma, \text{CT}) \mid \epsilon$
(Class Sig.)	$\text{CSig} ::= \emptyset \mid \text{CSig} \cdot [C \mapsto [\text{remote}] \ C \ \vec{T}\vec{f} \{ m_i : C_i \rightarrow U_i \}]$
(Identifiers)	$u ::= x \mid o \mid c$
(Threads)	$P ::= \mathbf{0} \mid P_1 \mid P_2 \mid (\nu u)P \mid \text{forked } e \mid \text{go } e \text{ with } c \mid e \text{ with } c$ $\mid \text{go } e \text{ to } c \mid \text{return}(c) \ e \mid \text{Error}$
(Configurations)	$F ::= (\nu \vec{u})(P, \sigma, \text{CT})$
(Networks)	$N ::= \mathbf{0} \mid l[F] \mid N_1 \mid N_2 \mid (\nu u)N$
(Stores)	$\sigma ::= \emptyset \mid \sigma \cdot [x \mapsto v] \mid \sigma \cdot [o \mapsto (C, \vec{f} : \vec{v}, n, \{ \vec{c} \})]$
(Class tables)	$\text{CT} ::= \emptyset \mid \text{CT} \cdot [C \mapsto L]$

Fig. 3: The syntax of the language DJ.

type, which is assigned to frozen expressions that expect a parameter of type T and return a value of type U . We abbreviate the type of thunked frozen expressions as $\text{thunk}\langle U \rangle \stackrel{\text{def}}{=} \text{unit} \rightarrow U$. We associate the type $\text{ser}\langle U \rangle$ with frozen *values*. If a value v has type U is frozen then the result has the type $\text{ser}\langle U \rangle$.

Expressions. The syntax is standard, including the standard synchronisation constructs of the Java language, except for two code passing primitives. The first primitive, $\text{freeze}[t](T \ x) \{ e \}$ takes the expression e and, without evaluating it, produces a flattened value representation parameterised by variable x with type T . Any parts of the local store required by the expression (such as the information held in variables free in e) are included in this new value, along with class information it may need for execution.

The tag t is a flag to control the amount of this information sent along with e by the user. If he specifies **eager**, then the code is automatically frozen together with all classes that *may* be used. If the user selects **lazy**, it is the responsibility of the receiving virtual machine to obtain missing classes. The third option is called *user-specified* information, and allows the programmer to supply a list of class names. Only these classes and their dependents (such as superclasses) are included with the frozen value.

Dual to freezing, the action **defrost**($e_0; e$) expects the evaluation of expression e to produce a piece of frozen code. This code is then executed, substituting its parameter with the value obtained by evaluating e_0 , much like invoking a method. We abbreviate freeze and defrost expressions that take no parameters as **freeze**[t]{ e } $\stackrel{\text{def}}{=} \text{freeze}[t](\text{unit } x)\{e\}$ ($x \notin \text{fv}(e)$) and **defrost**(e) $\stackrel{\text{def}}{=} \text{defrost}((); e)$ respectively. Note that $()$ denotes a constant of **unit** type.

To simplify the presentation, we only allow single parameters to methods and to frozen expressions. This does not restrict the expressiveness of programs written in DJ, as there is a semantics and type-preserving mapping from programs with multiple parameters to this subset. See Proposition 7.1 § 7 for the formal proofs.

For clarity, we introduce two *derived* constructs that are syntactic sugar for serialisation and deserialisation.

$$\text{serialize}(e) \stackrel{\text{def}}{=} \text{freeze}[\text{lazy}]\{e\} \quad \text{and} \quad \text{deserialize}(e) \stackrel{\text{def}}{=} \text{defrost}(e)$$

Class Signatures. A class signature CSig is a mapping from class names to their interface types (or signatures). We assume CSig is given globally, as a minimum agreed interface between remote parties, unlike class tables which are maintained on a per-location basis. Attached to each signature is the name of a direct super-class, as well as the declaration “**remote**” if the class is remote. For a class C , the predicate **remote**(C) holds iff “**remote**” appears in CSig(C); otherwise **local**(C) holds. Class signatures contain only expected method signatures, not their implementation. This provides a lightweight mechanism for determining the type of remote methods.

3.2 Runtime syntax

Runtime syntax extends the user syntax to represent the distributed state of multiple sites communicating with each other, including remote operations in transit.

Expressions. Location names are written $l, m, \dots$ and can be thought of as IP addresses in a network. **new** $C^l(\vec{v})$, **download** $\vec{C}$ **from** l **in** e and **resolve** $\vec{C}$ **from** l **in** e define the machinery for class downloading, which will be explained along with the operational semantics in § 4.1 and § 4.2. The key expression is **new** $C^l(\vec{e})$, indicating that the definition of class C can be obtained from a location called l should it need to be instantiated. We write C^- when the treatment of class name C is independent of whether it is labelled or not. **await** c is fundamental to the model of method invocation and can be thought of as the return point for a call. **sandbox** { e } represents the execution environment of some code e that originated from a frozen expression.

insync o { e } denotes that expression e has previously acquired the lock on object o . When an expression contains **ready** o n as a sub-term it indicates that it is ready to re-acquire the lock on object o . The expression **waiting**(c) n denotes an expression waiting for notification on channel c , at which point it may try to re-acquire a lock it was holding. n indicates the number of times that this waiting thread had entered its lock before yielding. Finally, the expression **Error** denotes the null-pointer error.

Values. v is also extended with runtime terms. Object identifiers o denote references to instances of classes as well as the destination of RMI calls. We shall

often write “o-id” for brevity. Channels c are fundamental to the mechanism of method invocation and determine the return destination for both remote and local method calls, as illustrated in the operational semantics later. We call o and c *names*.

The most interesting extended value is a *frozen expression*, a piece of code or data that can be passed between methods as a value. Later, it can be “defrosted” at which point it is executed to compute a value. $\lambda(T\ x).(\nu\ \vec{u})(l, e, \sigma, \mathbf{CT})$ denotes an expression e frozen with class table $\mathbf{CT}$ created at l . Expression e is parameterised by variable x with type T , and σ contains data local to the expression that was stored along with it at “freezing time”. The identifiers $\vec{u}$ correspond to the domain of σ . x and $\vec{u}$ occur bound. $\mathbf{CT}$ ships class bodies that may be used during the execution of e . If it is empty and the party evaluating e lacks a required class, it should attempt to download a copy from l . If σ or $\mathbf{CT}$ is empty, then we shall omit writing them entirely for clarity. Finally, the value ϵ serves as a constant that appears at runtime as the return value of `void` methods.

Threads. $P \mid Q$ says P and Q are two threads running in parallel, while $(\nu\ u)P$ restricts identifier u local to P . $\mathbf{0}$ denotes an empty thread. This notation comes from the π -calculus [Milner et al. 1992]. It also includes **Error** which denotes the result of class downloading and communication failure. The expression **forked** e says that expression e was previously forked from another thread. The remaining constructs of P are used for representing the RMI mechanism, and are illustrated when we discuss the operational semantics in § 4.

Configurations and Networks. F represent an instance of a virtual machine. A configuration $(\nu\ \vec{u})(P, \sigma, \mathbf{CT})$ consists of threads P , a store σ containing local variables and objects, and a class table $\mathbf{CT}$. Networks are written N , and comprise zero or more located configurations executing in parallel. $\mathbf{0}$ denotes the empty network. $l[F]$ denotes a configuration F executing at location l . $N_1 \mid N_2$ and $(\nu\ \vec{u})N$ are understood as in threads.

A *store* σ consists of a mapping from variable names to values, written $[x \mapsto v]$, or from object identifiers to store objects, written $[o \mapsto (C, \vec{f} : \vec{v}, n, \{\vec{c}\})]$. This indicates an identifier o maps to an object of class C with a vector of fields with values $\vec{f} : \vec{v}$. The set of channels $\{\vec{c}\}$ contains identifiers for threads currently waiting on o , i.e. those that have executed `o.wait` and have not received notification. The number, n , indicates how many times the lock on this object has been entered by a thread.

Finally, class tables $\mathbf{CT}$, are a mapping from unlabelled class names to class definitions (metavariable L in Figure 3). Throughout the paper we write **FCT** for the *foundation class table* that contains the common classes that every location in the network should possess, roughly corresponding to the `java.*` classes.

Auxiliary functions. Several auxiliary functions are defined over the syntax of DJ. `dom` returns the domain of a store and a class table. We also write `dom(F)` etc. to denote a domain of stores which appear in F . The set of free variables $\text{fv}(N)$ and names $\text{fn}(N)$ are standard. We also use $\text{fnv}(N) \stackrel{\text{def}}{=} \text{fv}(N) \cup \text{fn}(N)$. The full definitions is given in the Appendix A.

The set of free class names for a given term is given by the function `fcl` which

is defined over expressions, threads, configurations and class table entries. The free class names of a value v is defined as $\text{fcl}(v) = \emptyset$. The free class names of an expression and threads are defined recursively as the union of the free class names of all sub-expressions, with the exception that:

$$\text{fcl}(\text{new } C(\vec{e})) = \bigcup \text{fcl}(e_i) \cup \{C\} \quad \text{and importantly:} \quad \text{fcl}(\text{new } C^l(\vec{e})) = \bigcup \text{fcl}(e_i)$$

For class tables, we retrieve the free class names appearing in the bodies of methods:

$$\text{fcl}(\text{class } C \text{ extends } D \{ \vec{T}\vec{f}; K \vec{M} \}) = \bigcup \text{fcl}(e_i) \text{ where } M_i = U_i \ m_i(C_i \ x_i) \{e_i\}$$

For stores, we set $\text{fcl}(\sigma) = \{C \mid [o \mapsto (C, \dots)] \in \sigma\}$.

4. OPERATIONAL SEMANTICS

This section presents the formal operational semantics of DJ, extending the standard small step call-by-value reduction of [Pierce 2002; Bierman et al. 2003]. There are two reduction relations. The first is defined over configurations executing within an individual location, written $F \longrightarrow_l F'$, where l is the name of the location containing F . The second is defined over the networks, written $N \longrightarrow N'$. $F \longrightarrow_l F'$ promotes to $l[F] \longrightarrow l[F']$. Both relations are given modulo the standard structural equivalence rules of the π -calculus [Milner et al. 1992], written $\equiv$ and given in Appendix B. We define *multi-step* reductions as: $\longrightarrow^{\text{def}} (\longrightarrow \cup \equiv)^*$ and $\longrightarrow_l^{\text{def}} (\longrightarrow_l \cup \equiv)^*$.

4.1 Local expressions

The rules for the sequential part of the language are standard [Igarashi et al. 2001; Bierman et al. 2003]. We list the reduction rules in Figure 4. When allocating a new object by **NEW**, we explicitly restrict identifiers, which represents “freshness” or “uniqueness” of the address in the store. The auxiliary function $\text{fields}(C)$ is given in Figure 5. It examines the class signature and returns the field declarations for C . We assume the existence of a distinguished class **Object** that is at the top of the inheritance hierarchy and contains no methods and no fields (written $\bullet$).

Tagged class creation takes place in **NEW_R** and **NEW_L**. The former rule is applied whenever execution attempts to instantiate an object of a tagged class whose body is not present in the local class table. Instead of immediately allocating a new object, it first attempts to download the actual body of the class from the labelled location. This is discussed in detail in § 4.2. **NEW_L** is applied when an attempt is made to instantiate a tagged class whose body is already available locally. In this case the statement reduces to a normal untagged instantiation.

To reduce the number of computation rules, we make use of the evaluation contexts in Figure 6 and the congruence rule **CONG**. Contexts contain a single hole, written $[]$ inside them. $E[e]$ represents the expression obtained by replacing the hole in context E with the ordinary expression e . The evaluation order of terms in the language is determined by the construction of these contexts.

4.2 Class downloading

Class mobility is very important in Java RMI systems, since it reduces unnecessary coupling between communicating parties. If an interface can be agreed, then any

VAR $x, \sigma, \text{CT} \longrightarrow_l \sigma(x), \sigma, \text{CT}$	COND $\text{if } (\text{true}) \{e_1\} \text{ else } \{e_2\}, \sigma, \text{CT} \longrightarrow_l e_1, \sigma, \text{CT}$ $\text{if } (\text{false}) \{e_1\} \text{ else } \{e_2\}, \sigma, \text{CT} \longrightarrow_l e_2, \sigma, \text{CT}$
WHILE $\text{while } (e_1) \{e_2\}, \sigma, \text{CT} \longrightarrow_l \text{if } (e_1) \{e_2; \text{while } (e_1) \{e_2\}\} \text{ else } \{e\}, \sigma, \text{CT}$	
FLD $\frac{\sigma(o) = (C, \vec{f} : \vec{v})}{o.f_i, \sigma, \text{CT} \longrightarrow_l v_i, \sigma, \text{CT}}$	SEQ $\frac{e_1, \sigma, \text{CT} \longrightarrow_l (\nu \vec{u})(v, \sigma', \text{CT}')}{e_1; e_2, \sigma, \text{CT} \longrightarrow_l (\nu \vec{u})(e_2, \sigma', \text{CT}')} \quad \vec{u} \notin \text{fnv}(e_2)$
ASS $x = v, \sigma, \text{CT} \longrightarrow_l v, \sigma[x \mapsto v], \text{CT}$	FLDASS $\frac{\sigma' = \sigma[o \mapsto \sigma(o)[f \mapsto v]]}{o.f = v, \sigma, \text{CT} \longrightarrow_l v, \sigma', \text{CT}} \quad o \in \text{dom}(\sigma)$
NEW $\frac{\text{fields}(C) = \vec{T}\vec{f}}{\text{new } C(\vec{v}), \sigma, \text{CT} \longrightarrow_l (\nu o)(o, \sigma \cdot [o \mapsto (C, \vec{f} : \vec{v}, 0, \emptyset), \text{CT}]}$	
NEW R $\text{new } C^m(\vec{v}), \sigma, \text{CT} \longrightarrow_l \text{download } C \text{ from } m \text{ in new } C^m(\vec{v}), \sigma, \text{CT} \quad C \notin \text{dom}(\text{CT})$	
NEW L $\text{new } C^m(\vec{v}), \sigma, \text{CT} \longrightarrow_l \text{new } C(\vec{v}), \sigma, \text{CT} \quad C \in \text{dom}(\text{CT})$	
DEC $T \ x = v; e, \sigma, \text{CT} \longrightarrow_l (\nu x)(e, \sigma \cdot [x \mapsto v], \text{CT}) \quad x \notin \text{dom}(\sigma)$	
CONG $\frac{e, \sigma, \text{CT} \longrightarrow_l (\nu \vec{u})(e', \sigma', \text{CT}')}{E[e], \sigma, \text{CT} \longrightarrow_l (\nu \vec{u})(E[e'], \sigma', \text{CT}')} \quad \vec{u} \notin \text{fnv}(E)$	

Fig. 4: Local expressions

(Field lookup)	
$\text{fields}(\text{Object}) = \bullet$	$\frac{\text{CSig}(C) = D \quad \vec{T}\vec{f} \{m_i : C_i \rightarrow U_i\} \quad \text{fields}(D) = \vec{T}'\vec{f}'}{\text{fields}(C) = \vec{T}'\vec{f}', \vec{T}\vec{f}}$
(Method body lookup)	
$\frac{\text{CT}(C) = \text{class } C \text{ extends } D \{ \vec{T}\vec{f}; K \vec{M} \} \quad U \ m(C \ x)\{e\} \in \vec{M}}{\text{mbody}(m, C, \text{CT}) = (x, e)}$	$\frac{\text{CT}(C) = \text{class } C \text{ extends } D \{ \vec{T}\vec{f}; K \vec{M} \} \quad U \ m(C \ x)\{e\} \notin \vec{M}}{\text{mbody}(m, C, \text{CT}) = \text{mbody}(m, D, \text{CT})}$
(Method signature lookup)	
$\frac{\text{CSig}(C) = [\text{remote}] \ C \text{ extends } D \quad \vec{T}\vec{f} \{m_i : C_i \rightarrow U_i\}}{\text{mtype}(m_i, C) = \vec{T}_i' \rightarrow U_i'}$	$\frac{\text{CSig}(C) = [\text{remote}] \ C \text{ extends } D \quad \vec{T}\vec{f} \{m_i : C_i \rightarrow U_i\} \quad m \notin \{\vec{m}\}}{\text{mtype}(m, C) = \text{mtype}(m, D)}$

Fig. 5: Lookup functions

$E ::= [] \mid \text{if } (E) \{e\} \text{ else } \{e\} \mid E.f \mid E;e \mid x = E \mid E.f = e \mid o.f = E$
 $\mid \text{new } C(\vec{v}, E, \vec{e}) \mid E.m(e) \mid o.m(E) \mid T x = E \mid \text{defrost}(e; E) \mid \text{defrost}(E; v)$
 $\mid \text{sync } (E) \{e\} \mid E.\text{wait} \mid E.\text{notify} \mid E.\text{notifyAll} \mid \text{sandbox } \{E\} \mid \text{insync } o \{E\}$
 $\mid \text{forked } E \mid \text{go } E \text{ with } c \mid E \text{ with } c \mid \text{go } E \text{ to } c \mid \text{return}(c) E$

Fig. 6: Evaluation contexts

RESOLVE
$\frac{\text{CT}(C_i) = \text{class } C_i \text{ extends } D_i \{ \vec{T}\vec{f}; K \vec{M} \}}{\text{resolve } \vec{C} \text{ from } l' \text{ in } e, \sigma, \text{CT} \longrightarrow_l \text{download } \vec{D} \text{ from } l' \text{ in } e, \sigma, \text{CT}}$
DOWNLOAD
$\frac{\{\vec{D}\} = \{\vec{C}\} \setminus \text{dom}(\text{CT}_1) \quad \{\vec{F}\} = \text{fcl}(\text{CT}_2(\vec{D})) \quad \text{CT}' = \text{CT}_2(\vec{D})[\vec{F}^{l_2}/\vec{F}]}{l_1[E[\text{download } \vec{C} \text{ from } l_2 \text{ in } e] \mid P, \sigma_1, \text{CT}_1] \mid l_2[P_2, \sigma_2, \text{CT}_2] \longrightarrow l_1[E[\text{resolve } \vec{D} \text{ from } l_2 \text{ in } e] \mid P, \sigma_1, \text{CT}_1 \cup \text{CT}'] \mid l_2[P_2, \sigma_2, \text{CT}_2]}$
DNOTHING
$\text{download } \vec{C} \text{ from } l' \text{ in } e, \sigma, \text{CT} \longrightarrow_l e, \sigma, \text{CT} \quad C_i \in \text{dom}(\text{CT})$
ERR-CLASSNOTFOUND
$\frac{\exists C_i \in \vec{C}. C_i \notin \text{dom}(\text{CT}_1) \cup \text{dom}(\text{CT}_2)}{l_1[E[\text{download } \vec{C} \text{ from } l_2 \text{ in } e] \mid P, \sigma_1, \text{CT}_1] \mid l_2[P_2, \sigma_2, \text{CT}_2] \longrightarrow l_1[E[\text{Error}] \mid P, \sigma_1, \text{CT}_1] \mid l_2[P_2, \sigma_2, \text{CT}_2]}$

Fig. 7: Class resolution and downloading

class that implements the interface can be passed to a remote consumer and type-safety will be preserved. However this only works if sites are able to dynamically acquire class files from one another. This hidden behaviour is omitted from known sequential formalisms, as it is not required in the single-location setting, and so the formalisation of class downloading is one of the key contributions of DJ.

The rules for class downloading in DJ are given in Figure 7 and approximately model the lazy downloading mechanism found in JDK 1.3 without verification [Drossopoulou and Eisenbach 2002]. The *download* expression is responsible for the transfer of class table entries from a remote site. DOWNLOAD defines the semantics for this operation. For a download request *download* $\vec{C}$ from l in e we first produce $\vec{D}$ by removing the names of any classes locally available (and thus eliminating duplication). We then compute vector $\vec{F}$ from all the free class names mentioned in the bodies of the classes in $\vec{D}$. Finally, the classes named in $\vec{D}$ are downloaded and added to the local class table. Any occurrence of a member of $\vec{F}$ in a newly downloaded class body is tagged with the name of the remote site (l_2 in this case). *Resolution*, defined by RESOLVE, is the process of examining classes for unmet dependencies and scheduling the download of missing classes. Informally this amounts to downloading immediate superclasses.

The DOWNLOAD and RESOLVE rules work together to iteratively resolve all class dependencies for a given object. Once all dependencies have been met, normal execution continues after DNOTHING.

We model a failure in this process by the last rule. The rule ERR-CLASSNOTFOUND approximates `ClassNotFoundException` that would occur in the case of the site l_2

not possessing some class requested by l_1 . In this case, the code attempting the download will reduce to the **Error** expression.

In this paper we chose the option of class loading without verification as it allows significantly simpler presentation. However, our formalisation of class downloading is intended to be modular: it is possible to model different class loading mechanisms by adjusting the reduction rules for downloading and resolution and the class dependency algorithm introduced in Algorithm 4.2. For example, in rule RESOLVE the vector $\vec{D}$ is constructed from the direct superclasses of the classes being resolved. One aspect of Java verification is that it checks subtypes for method arguments. By inspecting the body of methods in the classes being resolved, we could extend $\vec{D}$ to reflect these checks as a first approximation.

Following on from this we observe that, with verification *on*, the overhead induced by Java’s lazy class loading policy is increased—since verifying a class typically requires the loading of more classes than just the direct superclass—making an even stronger case for eager code passing.

4.3 Serialisation and deserialisation

One of the contributions of DJ is a precise formalisation of the semantics of serialisation using the frozen expressions which are detailed in § 4.6 (for the encoding, see § 3.1). Serialisation occurs in two instances. In the first, the expressions **serialize**(v) and **deserialize**(e) allow explicit flattening and re-inflation of objects by the programmer, whereas the second instance occurs when values must be transported across the network.


serialize(v) and **deserialize**(e) must appear automatically as runtime expressions to serialise parameters and return values of remote method invocations. This is because instances of local classes—those classes without the **remote** keyword in their signature—are incapable of remote method invocation, and so cannot be passed by reference as parameters or as return values to remote methods. Should this occur, the remote party would receive the identifier of an unreachable object. Avoiding this problem involves making a deep clone of the local object, and we see this in action in § 4.4.

4.4 Method invocation

Unlike sequential formalisms, DJ describes *remote method invocation*. To accommodate RMI, the rules for method call take a novel form employing concepts from the π -calculus, representing the context of a call by a local linear channel. While this technique is well-known in the π -calculus [Milner et al. 1992], DJ may be the first to use it to faithfully capture the semantics of RMI in a Java-like language. Among other benefits, it allows us to define the semantics of local and remote method calls concisely and uniformly: a method call is local when the receiver is co-located with the caller; whereas it becomes remote when the receiver is located elsewhere. Remote calls also differ from local ones because of the need for parameter and return value serialisation, which is reflected as several extra reduction steps.

We summarise the general picture of a remote method invocation in Figure 8(a), which starts from dispatch of a remote method and ends with delivery of its return value. The corresponding formal rules are given in Figure 8(b).

We start our illustration from local method calls. For a method call $o.m(v)$, if



(a) Evaluation steps for a remote call

METHLOCAL $E[o.m(v)] \mid P, \sigma, \text{CT} \longrightarrow_l (\nu c)(E[\text{await } c] \mid o.m(v) \text{ with } c \mid P, \sigma, \text{CT}) \quad c \text{ fresh}, o \in \text{dom}(\sigma)$	
METHREMOTE $E[o.m(v)] \mid P, \sigma, \text{CT} \longrightarrow_l (\nu c)(E[\text{await } c] \mid \text{go } o.m(\text{serialize}(v)) \text{ with } c \mid P, \sigma, \text{CT})$ $c \text{ fresh}, o \notin \text{dom}(\sigma)$	
METHINVOKE $\frac{\sigma(o) = (C, \dots) \quad \text{mbody}(m, C, \text{CT}) = (x, e)}{o.m(v) \text{ with } c, \sigma, \text{CT} \longrightarrow_l (\nu x)(e[o/\text{this}][\text{return}(c)/\text{return}], \sigma \cdot [x \mapsto v], \text{CT})}$	
AWAIT $E[\text{await } c] \mid \text{return}(c) v, \sigma, \text{CT} \longrightarrow_l E[v], \sigma, \text{CT}$	
SERRETURN $l[\text{return}(c) v \mid P, \sigma, \text{CT}] \longrightarrow_l l[\text{go } \text{serialize}(v) \text{ to } c \mid P, \sigma, \text{CT}] \quad c \notin \text{fn}(P)$	
LEAVE $l_1[\text{go } o.m(v) \text{ with } c \mid P_1, \sigma_1, \text{CT}_1] \mid l_2[P_2, \sigma_2, \text{CT}_2] \longrightarrow_l l_1[P_1, \sigma_1, \text{CT}_1] \mid l_2[o.m(\text{deserialize}(v)) \text{ with } c \mid P_2, \sigma_2, \text{CT}_2] \quad o \in \text{dom}(\sigma_2)$	
RETURN $l_1[\text{go } v \text{ to } c \mid P_1, \sigma_1, \text{CT}_1] \mid l_2[P_2, \sigma_2, \text{CT}_2] \longrightarrow_l l_1[P_1, \sigma_1, \text{CT}_1] \mid l_2[\text{return}(c) \text{ deserialize}(v) \mid P_2, \sigma_2, \text{CT}_2] \quad c \in \text{fn}(P_2)$	
ERR-LOSTCALL $\text{go } o.m(v) \text{ with } c, \sigma, \text{CT} \longrightarrow_l \text{Error}, \sigma, \text{CT}$	
ERR-LOSTRETURN $\text{go } v \text{ to } c, \sigma, \text{CT} \longrightarrow_l \text{Error}, \sigma, \text{CT}$	

(b) Reduction rules

Fig. 8: Remote method invocation

$o \in \text{dom}(\sigma)$ then the rule METHLOCAL is applied. A new channel c is created to carry the return value of the method; the return point of the method call is replaced with the term **await** c corresponding to a receiver waiting for the return value supplied on channel c . The method call itself is spawned in a new thread as $o.m(v)$ **with** c carrying channel c .

The next stage is the application of the method invocation rule METHINVOKE. Both remote and local invocations apply this rule. The auxiliary function $\text{mbody}(m, C, \text{CT})$ is given in Figure 5, and is responsible for looking up the correct method body in the local class table. It returns a pair of the method code and the formal parameter

name. The receiver is substituted $[o/\text{this}]$ and a new store entry x is allocated for the formal parameter v . We apply the substitution $e[\text{return}(c)/\text{return}]$ to indicate that the return value of the method must be sent along channel c . The rule **AWAIT** is used to communicate the return value to its caller.

When $o \notin \text{dom}(\sigma)$ the method invocation is *remote*. The rule **METHREMOTE** is applied, with care being taken to automatically serialise the parameter v if it is a local object identifier. We note that frozen values are also transferred to the remote location without modification (like base values).

After serialisation, we are left with a thread of the form **go** $o.m(w)$ **with** c where w is the serialised representation of the original parameter v . At this point, the network level rule **LEAVE** triggers the migration of the calling thread to the location that holds the receiving object in its local store. After transfer over the network, the parameter is automatically deserialised and **METHINVOKE** applied. Again, the return value must be automatically serialised using **SERRETURN**. Then it crosses the network by application of **RETURN**. After returning to the caller site, it is again deserialised.

The last two rules present instances of network failure. In the case of **ERR-LOSTCALL**, the network becomes partitioned such that a remote method call attempting to reach its destination cannot. Likewise, in **ERR-LOSTRETURN**, the return value from a remote method call is lost. Both cases reduce to **Error**.

4.5 Multi-threaded programs

DJ contains several concurrency primitives that should be familiar to Java programmers. The reduction rules are given in Figure 9. We shall focus on the most important rules.

The rule **FORK** defines a simple command for creating a new thread. When evaluated, a new thread in the current location is started and begins executing an expression.

The rule **SYNC** defines a basic monitor construct. When executing $E[\text{sync}(o) \{e\}]$, we are attempting to acquire the lock on the object identified by o . To determine whether a lock is taken, the function $\text{getLock}(\sigma, o)$ returns the number of entries to the monitor on object o in store σ . If this count is non-zero, then the predicate $\text{insync}(o, E)$ is used to determine whether it is the current thread that owns the monitor (since Java allows re-entrant monitors). If this is the case then execution proceeds by incrementing the entry count using the function $\text{setLock}(\sigma, o, n')$ with $n' = n + 1$, otherwise execution blocks. The predicates are formally defined as follows.

Definition 4.1. (Lock functions and predicates)

$$\text{insync}(o, E) \iff \exists E_1, E_2 \text{ such that } E = E_1[\text{insync } o \{E_2[\]\}]$$

Suppose $\sigma(o) = (C, \vec{f} : \vec{v}, n, \{\vec{c}\})$. Then we define:

$$\begin{aligned} \text{getLock}(\sigma, o) &= n & \text{blocked}(\sigma, o) &= \{\vec{c}\} \\ \text{setLock}(\sigma, o, n') &= \sigma[o \mapsto (C, \vec{f} : \vec{v}, n', \{\vec{c}\})] \\ \text{block}(\sigma, o, c) &= \sigma[o \mapsto (C, \vec{f} : \vec{v}, n, \{\vec{c}\} \cup \{c\})] \\ \text{unblock}(\sigma, o, \vec{c}) &= \sigma[o \mapsto (C, \vec{f} : \vec{v}, n, \{\vec{c}\} \setminus \{\vec{c}\})] \end{aligned}$$

FORK $E[\text{fork}(e)], \sigma, \text{CT} \longrightarrow_l E[\epsilon] \mid \text{forked } e, \sigma, \text{CT}$	THREADDEATH $\text{forked } v, \sigma, \text{CT} \longrightarrow_l \mathbf{0}, \sigma, \text{CT}$
SYNC $\frac{\text{getLock}(\sigma, o) = \begin{cases} 0 & \text{setLock}(\sigma, o, 1) = \sigma' \\ n > 0 & \text{insync}(o, E) \implies \text{setLock}(\sigma, o, n+1) = \sigma' \end{cases}}{E[\text{sync } (o) \{e\}], \sigma, \text{CT} \longrightarrow_l E[\text{insync } o \{e\}], \sigma', \text{CT}}$	
WAIT $\frac{\text{insync}(o, E) \quad \text{getLock}(\sigma, o) = n \quad \text{setLock}(\sigma, o, 0) = \sigma'' \quad \text{block}(\sigma'', o, c) = \sigma'}{E[o.\text{wait}] \mid P, \sigma, \text{CT} \longrightarrow_l (\nu c)(E[\text{waiting}(c) \ n] \mid P, \sigma', \text{CT})}$	
NOTIFY $\frac{\text{insync}(o, E) \quad c \in \text{blocked}(\sigma, o) \quad \text{unblock}(\sigma, o, c) = \sigma'}{E[o.\text{notify}] \mid E_1[\text{waiting}(c) \ n], \sigma, \text{CT} \longrightarrow_l E[\epsilon] \mid E_1[\text{ready } o \ n], \sigma', \text{CT}}$	
NOTIFYALL $\frac{\text{insync}(o, E) \quad \text{blocked}(\sigma, o) = \{\vec{c}\} \quad m \geq 0 \quad \text{unblock}(\sigma, o, \vec{c}) = \sigma'}{E[o.\text{notifyAll}] \mid E_1[\text{waiting}(c_1) \ n_1] \mid \dots \mid E_m[\text{waiting}(c_m) \ n_m], \sigma, \text{CT} \longrightarrow_l E[\epsilon] \mid E_1[\text{ready } o \ n_1] \mid \dots \mid E_m[\text{ready } o \ n_m], \sigma', \text{CT}}$	
NOTIFYNONE $\frac{\text{insync}(o, E) \quad \text{blocked}(\sigma, o) = \emptyset}{E[o.\text{notify}], \sigma, \text{CT} \longrightarrow_l E[\epsilon], \sigma, \text{CT}}$	READY $\frac{\text{getLock}(\sigma, o) = 0 \quad \text{setLock}(\sigma, o, n) = \sigma'}{\text{ready } o \ n, \sigma, \text{CT} \longrightarrow_l \epsilon, \sigma', \text{CT}}$
LEAVECRITICAL $\frac{\text{getLock}(\sigma, o) = n \quad \text{setLock}(\sigma, o, n-1) = \sigma'}{\text{insync } o \{v\}, \sigma, \text{CT} \longrightarrow_l v, \sigma', \text{CT}}$ $\text{insync } o \{\text{return}(c) \ v\}, \sigma, \text{CT} \longrightarrow_l \text{return}(c) \ v, \sigma', \text{CT}$	

Fig. 9: Concurrency primitives

To temporarily release a lock held on an object o , the command $o.\text{wait}$ can be used with semantics as in **WAIT**. First, a new channel c is created and its name is added to the blocked set for the object o by application of the function $\text{block}(\sigma, o, c)$. The currently executing thread then enters a sleeping state, written $\text{waiting}(c) \ n$ where n indicates the number of times the thread had entered the monitor on o before now. The lock count for o is then set to 0.

To wake sleeping threads, the commands $o.\text{notify}$ and $o.\text{notifyAll}$ are provided. They differ in that the former non-deterministically wakes only one of the threads waiting on o , whereas the latter wakes them all. We shall focus on the rule **NOTIFY**. When notifying a thread, that thread must be waiting on some channel c which is held in the blocked set for o . This channel is then removed from the blocked set by the function $\text{unblock}(\sigma, o, c)$. The woken thread then moves to the state of being $\text{ready } o \ n$, which means it is ready to re-acquire the lock on o , n times. It cannot *immediately* acquire this lock, since necessarily the thread that performs the notification is still within its critical section. However, as soon as that thread leaves its critical section the newly woken party can compete to acquire the lock.

4.6 Direct code mobility

Frozen expressions offer a direct way to manipulate code and data. They permit the storing of unevaluated terms that can, for example, be shipped to remote loca-

FREEZE	
$\{\vec{y}\} = \text{fv}(e) \setminus \{x\}$	$\sigma_y = \bigcup \sigma(y_i) \quad \sigma' = \text{og}(\sigma, \text{fn}(e) \cup \text{fn}(\sigma_y)) \cup \sigma_y \quad \{\vec{u}\} = \text{dom}(\sigma')$
$\text{CT}' = \begin{cases} \text{cg}(\text{CT}, \text{fcl}(e) \cup \text{fcl}(\sigma')) & t = \text{eager} \\ \text{cg}(\text{CT}, \vec{C}) & t = \vec{C} \\ \emptyset & t = \text{lazy} \end{cases}$	$\{\vec{F}\} = \text{fcl}(e) \setminus \text{dom}(\text{CT}')$
$\text{freeze}[t](T\ x)\{e\}, \sigma, \text{CT} \longrightarrow_l \lambda(T\ x).(\nu \vec{u})(l, e[\vec{F}^l / \vec{F}], \sigma', \text{CT}'), \sigma, \text{CT}$	
DEFROST	
$\{\vec{F}\} = \text{fcl}(\sigma') \setminus \text{dom}(\text{CT}')$	
$\text{defrost}(v; \lambda(T\ x).(\nu \vec{u})(m, e, \sigma', \text{CT}')), \sigma, \text{CT}$	
$\longrightarrow_l (\nu x \vec{u})(\text{download } \vec{F} \text{ from } m \text{ in sandbox } \{e\}, \sigma \cup \sigma' \cdot [x \mapsto v], \text{CT} \cup \text{CT}')$	
LEAVESANDBOX	
$\text{sandbox } \{v\}, \sigma, \text{CT} \longrightarrow_l v, \sigma, \text{CT}$	

Fig. 10: Creating and executing frozen expressions

tions for evaluation or merely saved for future use. As we have seen in § 3.1, our formulation of the primitives subsumes the serialisation operations found in Java that were explained in § 4.3.

As introduced in Figure 3, there are two operations associated with frozen values—for their creation and use—called *freezing* and *defrosting* respectively. Their rules are given in Figure 10.

Freezing is given by **FREEZE**, and has modes **lazy**, **eager**, and *user-specified*. Its operation is divided into two steps. The first step in any mode is to determine the store locations used by the expression e . We do this by examining the expression for any free variables, excluding the formal parameter x . The store entries for each variable are copied, σ_y . Next, we search for all the free object identifiers in e , written $\text{fn}(e)$. Because variables may hold references to objects, we must then examine the store fragment σ_y for any object identifiers held in the co-domain of variable mappings. Finally, objects have internal structure, so we apply the object graph function given in Algorithm 4.1 to copy all local objects transitively referenced by e or its variables, resulting in σ' . Base values stored in variables are copied “as-is”.

In the second step the freezing mode matters because it directly affects the amount of class information included in CT' . For the **lazy** case, no extra classes are provided with the expression, so the result of applying **FREEZE** is a value of the form $\lambda(T\ x).(\nu \vec{u})(l, e, \sigma, \emptyset)$.

When the case is **eager**, the creator of the frozen expression takes responsibility for including *all* classes that e depends upon. In the case that the user specifies a list of classes $\vec{C}$, only those classes and their dependencies are included. In either situation, we must use the class dependency algorithm in Algorithm 4.2 to determine the classes that the expression (or the user-specified classes) rely upon.

Finally Figure 11 lists the reduction rules for threads and networks. They are standard using contexts and the structure rules.

4.7 Algorithms

4.7.1 Object graph. An object graph is defined as the set of all mappings from an object identifier to store object for every local object transitively referenced by local object identifier v . The lock count, n , for each object is reset to zero when copied and the blocked set $\vec{c}$ emptied to preserve linearity. If the value v refers to a remote object, or a base value such as a boolean, then the object graph is empty.

ALGORITHM 4.1. (Object graph calculation) The function $\text{og}(\sigma, v)$ which computes the object graph of value v in store σ is defined as follows.

$$\text{og}(\sigma, v) = \begin{cases} \emptyset & \text{if } v \notin \text{dom}(\sigma) \vee \text{remote}(C) \\ [v \mapsto (C, \vec{f}: \vec{v}, 0, \emptyset)] \cup \text{og}(\sigma_i, o_i) & \text{otherwise} \end{cases}$$

where $\sigma(v) = (C, \vec{f}: \vec{v}, n, \{\vec{c}\})$, $\{\vec{o}\} = \text{fn}(\vec{v})$, $\sigma_1 = \sigma \setminus \{v\}$ and $\sigma_{i+1} = \sigma_i \setminus \text{dom}(\text{og}(\sigma_i, o_i))$.

See [Ahern and Yoshida 2005a, Example 4.9] for an example of the algorithm. In the full Java language, fields may be marked **transient**. Such fields are never serialised (for example they may contain a value that can be derived from other fields, or a reference to a non-serialisable object). Similarly, the Emerald language [Hutchinson et al. 1991] supports a qualifier called “**attached**” that indicates which of an object’s fields should be brought along it when it is copied. To support these extra features in DJ would involve the straightforward extension of syntax and a trivial modification to the object graph algorithm.

4.7.2 Class dependencies. An expression e directly depends upon a class C when $C \in \text{fcl}(e)$. e indirectly depends upon a class C when $\exists D \in \text{fcl}(e)$ and D is a subclass of C , or C appears free in the body of a method declared in D . Informally, dependency occurs when execution of an expression may at some point trigger instantiation of a class.

In order to calculate sets of dependencies we define an algorithm as follows:

ALGORITHM 4.2. (Class dependency set calculation)

$$\begin{aligned} \text{cg}(\text{CT}, \vec{M}) &= \bigcup \text{cg}(\text{CT}, \text{fcl}(e_i)) \text{ with } M_i = U_i \ m_i(C_i \ x_i) \{e_i\} \\ \text{cg}(\text{CT}, C) &= \begin{cases} \emptyset & \text{if } C \notin \text{dom}(\text{CT}) \vee C \in \text{dom}(\text{FCT}) \\ \text{cg}(\text{CT}, \text{CT}(C)) & \text{otherwise} \end{cases} \\ \text{cg}(\text{CT}, \vec{C}) &= \bigcup \text{cg}(\text{CT}, C_i) \\ \text{cg}(\text{CT}, L) &= \text{cg}(\text{CT} \setminus C, D) \cup \text{cg}(\text{CT} \setminus C, \vec{M}) \cup [C \mapsto L] \\ &\text{where } L = \text{class } C \text{ extends } D \ \{\vec{T}\vec{f}; K \vec{M}\} \end{aligned}$$

4.7.3 Correctness of algorithms. In this subsection we show one of the key results that are used in the proof in § 6, the correctness of the graph algorithms.

First we define the predicate $\text{reachable}(\sigma, o, o')$ to hold if there exists a path in store σ from the object with identifier o to the object with identifier o' . This can be an immediate link (when o' is stored in a field of o), or it can be via the fields of one or more intermediaries. This is defined as follows, where $\sigma(o) = (C, \vec{f}: \vec{v})$:

$$\text{reachable}(\sigma, o, o') \iff (o' \in \text{fn}(\vec{v}) \vee \exists o'' \in \text{fn}(\vec{v}). \text{reachable}(\sigma, o'', o'))$$

We extend this predicate to arbitrary thread P :

$$\text{reachable}(\sigma, P, o) \iff \exists o' \in \{\vec{o}\}.\text{reachable}(\sigma, o', o)$$

where $\{\vec{x}\} = \text{fv}(P)$ and $\{\vec{o}\} = \bigcup \text{fn}(\sigma(x_i)) \cup \text{fn}(P)$.

With this predicate we define the relation $RCH(\sigma)$ which contains all reachable pairs of objects in a store σ as follows:

$$RCH(\sigma) = \{(o, o') \mid \forall o, o' \in \text{dom}(\sigma). o \neq o' \wedge \text{reachable}(\sigma, o, o')\}$$

Our object graph algorithm must, to be correct, preserve the tree structure of the store when copying objects, hence it must preserve this reachability relation.

For a store σ and an object graph σ_g computed from that store, the predicate $\text{ogcomp}(\sigma, \sigma_g)$ (*completeness of an object graph*) holds if the computed graph preserves the reachability relation for all object identifiers in its domain. Given $RCH(\sigma)$ and $RCH(\sigma_g)$, we define:

$$\text{ogcomp}(\sigma, \sigma_g) \stackrel{\text{def}}{=} \forall o \in \text{dom}(\sigma) \cap \text{dom}(\sigma_g). (o, o') \in RCH(\sigma) \iff (o, o') \in RCH(\sigma_g)$$

This property ensures all links are correctly copied to the graph σ_g , and no new links are created.

4.7.4 Class dependency algorithm. The correctness of the class dependency algorithm relies upon the definition of the following predicate:

$$\text{comp}(C, \text{CT}) \stackrel{\text{def}}{=} \forall D \ C <: D. D \in \text{dom}(\text{CT})$$

which is read: class table CT is *complete with respect to* class C . When C is actually used, the class table CT at that location should be complete w.r.t. C . We extend the notion of completeness to entire class tables: we say a class table CT is *complete* if the following predicate holds:

$$\text{ctcomp}(\text{CT}) \stackrel{\text{def}}{=} \forall D \in \text{dom}(\text{CT}). \text{comp}(D, \text{CT})$$

This means for every class $D \in \text{dom}(\text{CT})$, every superclass of D is also available in CT .

With these preliminaries dealt with, we have the following lemma:

LEMMA 4.1. (Correctness of algorithms)

- (1) $\Gamma; \Delta \vdash \sigma : \text{ok}$ and $\sigma' = \text{og}(\sigma, o)$ implies $\Gamma; \emptyset \vdash \sigma' : \text{ok}$; for all $o' \in \text{dom}(\sigma')$, we have $\sigma'(o') = (C, \dots)$ with $\text{local}(C)$; and for all $o' \in (\text{fn}(\sigma') \setminus \text{dom}(\sigma'))$, $\Gamma \vdash o' : C$ with $\text{remote}(C)$.
- (2) $\sigma' = \text{og}(\sigma, v)$ implies $\text{ogcomp}(\sigma, \sigma')$.
- (3) Suppose $\vdash \text{CT} : \text{ok}$ with $C_i \in \text{dom}(\text{CSig})$. Then we have $\vdash \text{cg}(\text{CT}, \vec{C}) : \text{ok}$.
- (4) $\text{ctcomp}(\text{CT})$ and $\text{CT}' = \text{cg}(\text{CT}, C)$ imply $\text{ctcomp}(\text{CT}' \cup \text{FCT})$.

PROOF. See Appendix C. $\square$

4.8 Examples of executions

This subsection gives three small examples of the dynamic semantics, focusing on distributed primitives. Since no concurrent threads exist, we omit the counter 0 and an empty queue $\emptyset$ from store entries $[o \mapsto (C, \dots, 0, \emptyset)]$.

$\frac{\text{RC-PAR} \quad \frac{P_1, \sigma, \text{CT} \longrightarrow_l (\nu \vec{u})(P'_1, \sigma', \text{CT}')}{P_1 \mid P_2, \sigma, \text{CT} \longrightarrow_l (\nu \vec{u})(P'_1 \mid P_2, \sigma', \text{CT}')} \quad \vec{u} \notin \text{fnv}(P_2)}{\quad}$		$\frac{\text{RC-STR} \quad F \equiv F_0 \longrightarrow_l F'_0 \equiv F'}{F \longrightarrow_l F'}$
$\frac{\text{RC-RES} \quad \frac{(\nu \vec{u})(P, \sigma, \text{CT}) \longrightarrow_l (\nu \vec{u}')(P', \sigma', \text{CT}')}{(\nu u\vec{u})(P, \sigma, \text{CT}) \longrightarrow_l (\nu u\vec{u}')(P', \sigma', \text{CT}')}}{\quad}$		
$\frac{\text{RN-CONF} \quad F \longrightarrow_l F'}{l[F] \longrightarrow l[F']}$	$\frac{\text{RN-PAR} \quad N \longrightarrow N'}{N \mid N_0 \longrightarrow N' \mid N_0}$	$\frac{\text{RN-RES} \quad N \longrightarrow N'}{(\nu u)N \longrightarrow (\nu u)N'}$
		$\frac{\text{RN-STR} \quad N \equiv N_0 \longrightarrow N'_0 \equiv N'}{N \longrightarrow N'}$

Fig. 11: Network and thread

Freeze and Defrost. First we demonstrate freeze and defrost. After executing the program in Listing 9, at location l , we should obtain a frozen expression of the form:

$$\lambda(\text{int } x).(\nu o_1, o_2, y)(l, x + y + o_1.f, \sigma_1, \text{CT}_1)$$

where $\sigma_1 = [o_1 \mapsto (A, f : 5, g : o_2)] \cdot [o_2 \mapsto (B, \dots)] \cdot [y \mapsto 6]$ and $\text{CT}_1 = [B \mapsto \dots]$

```

1  class A {
2    int f; B g;
3    A(int f, B g) { this.f = f; this.g = g; }
4  }
5  class B { }
6  // Program
7  int y = 6; A o1 = new A(5, new B());
8  freeze[B] int x){x + y + o1.f};

```

Listing 9: Example of a program using freeze

To defrost a frozen value $\lambda(T x).(\nu \vec{u})(m, e, \sigma_1, \text{CT}_1)$ we apply DEFROST. Firstly, any classes supplied with the frozen value are appended to the current class table. Any class names appearing free in e are tagged with their originating location: **new** $C(\vec{e})$ becomes **new** $C^l(\vec{e})$. During execution of the newly defrosted code, when an expression such as the above **new** $C^l(\vec{v})$ is encountered then NEWR is applied if the body of C has not been downloaded to the execution location.

The second stage is to merge the data shipped with the value, σ_1 , into the local store. It is not possible to merely append this to the local store, since this could cause a name clash (for example two entries for variable x in the same scope). Therefore we create new memory locations for the formal parameter of the frozen expression, as well as for every element in the domain of the accompanying store entries. This is written $(\nu x\vec{u})$. It is then safe to append the new store and allocate space for the formal parameter. We write the new store at the location as $\sigma \cup \sigma_1 \cdot [x \mapsto v]$.

The final aspect of the defrost rule is to download the classes for all the objects added to the store in the previous step, because we may have added instances of classes not present at this location. This means instead of immediately evaluating e we call **download $\vec{F}$ from m in sandbox $\{e\}$** . This accurately mimics the mechanism employed by the `RMIClassLoader` class used in Java RMI. When sending marshaled

objects, RMI implementations annotate the data stream for classes with a codebase URL. This is a pointer to a remote directory that the `RMIClassLoader` can refer to download classes that are not available at the current location.

After class downloading has completed, we are left with an expression of the form `sandbox {e}`. Execution inside the sandbox then proceeds until a value is computed, which is then propagated to the enclosing scope according to the rule `LEAVESANDBOX`. Take the frozen expression computed in the example previously and call it t . We now give another example of defrosting this time at a location l , where it is important to notice that a variable y is already in scope: ν -operator will be used to avoid collision of bound variables and names. We abbreviate `download` to `dl` and `sandbox` to `sb` in the following:

$$\begin{aligned}
& \text{defrost}(5; t), [y \mapsto \text{true}], \text{CT} \\
\longrightarrow_m & (\nu x, o_1, o_2, y_2)(\text{dl } A \text{ from } l \text{ in sb } \{x + y_2 + o_1.f\}, \sigma_2, \text{CT}_2) \\
& \text{with } \sigma_2 = [y \mapsto \text{true}] \cdot [o_1 \mapsto (A, f : 5, g : o_2)] \cdot [o_2 \mapsto (B, \epsilon)] \cdot [y_1 \mapsto 6] \cdot [x \mapsto 5] \\
& \text{and } \text{CT}_2 = \text{CT} \cdot [B \mapsto \dots] \\
\longrightarrow_m & (\nu x, o_1, o_2, y_1)(\text{resolve } A \text{ from } l \text{ in sb } \{x + y_1 + o_1.f\}, \sigma_2, \text{CT}_3) \\
& \text{with } \text{CT}_3 = \text{CT}_2 \cdot [A \mapsto \dots]
\end{aligned}$$

Assuming that the superclass of A is *Object*, this should already be present in the local class table.

$$\begin{aligned}
& \longrightarrow_m (\nu x, o_1, o_2, y_1)(\text{dl } \text{Object} \text{ from } l \text{ in sb } \{x + y_1 + o_1.f\}, \sigma_2, \text{CT}_3) \\
& \longrightarrow_m (\nu x, o_1, o_2, y_1)(\text{sb } \{x + y_1 + o_1.f\}, \sigma_2, \text{CT}_3) \\
& \longrightarrow_m \text{sb } \{16\}, [y \mapsto \text{true}], \text{CT}_3 \quad \longrightarrow_m 16, [y \mapsto \text{true}], \text{CT}_3
\end{aligned}$$

In the final steps, we garbage-collect the store entries added by the frozen expression since they are now no longer required.

Class downloading. To illustrate the different class loading mechanisms, we change the above example as follows and investigate the cases when we change B in `freeze` to `eager` or `lazy`.

```

1 class A extends C{ ...}
2 class B { }
3 class C {D m(){return new D()}}
4 class D { }

```

Listing 10: Example of eager and lazy class downloading

- In the case of `eager`, the frozen expression ships all classes (A, B, C, D). Hence there is no downloading required after `defrost`.
- In the case of `lazy`, the frozen expression ships no classes. When defrosting, it downloads A and B . When resolving them at the next step, A 's superclass C is called to be downloaded. After C is downloaded, the final class table becomes $\text{CT}_5 = \text{CT}_3 \cdot [C \mapsto \text{class } C \{D \text{ d}() \{ \text{return new } D^l() \} \}]$. Note that D is *not* downloaded: hence it is renamed to D^l so that if D requires instantiation, `NEW $R$` will be applied and D downloaded from l .

Remote method invocation. The last example is RMI. We replace class B in location l in Listing 9 with the one below. We assume location m has the remote class R and classes A and B are local.

```

1 // location l
2 class B { }
3 // Program
4 A o1 = new A(5, new B()); return r.f(o1);
5 // location m
6 class R { Integer f(A x){ return x.f +1}}
```

Listing 11: Example of remote method invocation

After the execution at the location l , we obtain the method invocation of the form:

$$(\nu o_1, o_2)(E[\text{return } r.f(o_1)], \sigma'_1, \text{CT}'_1)$$

with $\text{CT}'_1 = [A \mapsto \dots] \cdot [B \mapsto \dots]$ and $\sigma_2 = [o_1 \mapsto (A, f : 5, g : o_2)] \cdot [o_2 \mapsto (B, \epsilon)]$. This reduces to:

$$\begin{aligned} &\longrightarrow_l (\nu o_1, o_2, c)(E[\text{await } c] \mid \text{go } r.f(\text{serialize}(o_1)) \text{ with } c, \sigma'_1, \text{CT}'_1) \\ &\longrightarrow_l (\nu o_1, o_2, c)(E[\text{await } c] \mid \text{go } r.f(v) \text{ with } c, \sigma'_1, \text{CT}'_1) \end{aligned}$$

with $v = \lambda(\text{unit } x).(\nu o_1, o_2)(o_1, \sigma'_1)$. When the remote method invocation happens, the argument o_1 is serialised and frozen value v is created. Now $\text{go } r.f(v)$ with c moves to the location m by LEAVE, opening the scope of c . After the message reaches to the location m , v is deserialised, starting downloading classes A and B as follows. Below we assume $\text{CT}'_2 = [R \mapsto \dots]$ and $\sigma'_2 = [r \mapsto (R, \emptyset)]$:

$$\begin{aligned} &r.f(\text{deserialize}(v)) \text{ with } c, \sigma'_2, \text{CT}'_2 \\ &\longrightarrow_l (\nu o_1, o_2)(r.f(o_1) \text{ with } c, \sigma'_1 \cdot \sigma'_2, \text{CT}'_2) \\ &\longrightarrow_l (\nu o_1, o_2)(r.f(\text{dl } A, B \text{ from } l \text{ in sb } \{o_1\}) \text{ with } c, \sigma'_1 \cdot \sigma'_2, \text{CT}'_2) \\ &\longrightarrow_l (\nu o_1, o_2)(r.f(o_1) \text{ with } c, \sigma'_1 \cdot \sigma'_2, \text{CT}'_2 \cdot \text{CT}'_1) \\ &\longrightarrow_l (\nu o_1, o_2, x)(\text{return}(c) \ x.f(o_1) + 1, \sigma'_1 \cdot \sigma'_2 \cdot [x \mapsto o_1], \text{CT}'_2 \cdot \text{CT}'_1) \\ &\longrightarrow_l \text{return}(c) \ 6, \sigma'_2, \text{CT}'_2 \cdot \text{CT}'_1 \longrightarrow_l \text{go } 6 \text{ to } c, \sigma'_2, \text{CT}'_2 \cdot \text{CT}'_1 \end{aligned}$$

Next “go 6 to c ” can safely return to the location l (since there exists only one $\text{await } c$) changing its form to “ $\text{return}(c) \ 6$ ” by RETURN. Finally we get $E[6]$ by AWAIT, as required.

5. TYPING SYSTEM

This section presents the key typing rules for DJ, focusing on the linear channel types and the use of invariants for typing runtime expressions and the new primitives. First we introduce the syntax of types and environments in Figure 12.

T represents expression types: booleans, class names, frozen expressions that take a parameter of type T and return elements of type U and the **unit** type. The metavariable U ranges over the same types as T but is augmented with the special type **void** with the usual empty meaning. We write $C <: D$ when class C is a subtype of class D . Our notion of subtyping is mostly standard (we assume $<:$ causes no cycle as in [Igarashi et al. 2001; Bierman et al. 2003]). It is given in Figure 12. The arrow type is standard.

$T ::= \text{boolean} \mid \text{unit} \mid C \mid T \rightarrow U$	(Types)
$U ::= \text{void} \mid T$	(Returnable types)
Extended types not appearing in program text:	
$S ::= U \mid \text{ret}(U)$	(Return types)
$\tau ::= \text{chan} \mid \text{chanI}(U) \mid \text{chanO}(U)$	(Channel types)
$\Gamma ::= \emptyset \mid \Gamma, x : T \mid \Gamma, o : C \mid \Gamma, \text{this} : C$	(Expression environment)
$\Delta ::= \emptyset \mid \Delta, c : \tau$	(Channel environment)
$T <: T$	
$\frac{C <: D \quad D <: E}{C <: E}$	
$\frac{U'_i <: U_i \quad 0 \leq i < n}{\vec{U}' <: \vec{U}}$	
$\frac{T' <: T \quad U <: U'}{T \rightarrow U <: T' \rightarrow U'}$	
$\frac{U' <: U}{\text{ret}(U') <: \text{ret}(U)}$	
$\frac{\text{CSig}(C) = C \text{ extends } D \vec{T} \vec{f} \{m_i : C_i \rightarrow U_i\}}{C <: D}$	

Fig. 12: Types and subtyping

Two runtime types (which do not appear in programs) are newly introduced. *Return types* are ranged over by S are used to denote the type of value returned by a method invocation ($U \ m(C \ x)\{e\}$ is well-typed if e has the type $\text{ret}(U)$). *Channel types* are ranged over by metavariable τ and represents the types for channels used in method calls, which is explained in the next subsection.

There are two different kinds of environment. The environment for typing expressions, written Γ , is a finite map from variables, o-ids and **this** to types ranged over by T . Δ is a finite map from channel names to channel types, and appears in judgements for method calls and those involving multiple threads and locations. We often omit empty environments from judgements for clarity of presentation.

5.1 Linear channel types

One of the key tasks of the typing rules is to ensure *linear* use of channels. This means that for every channel c there is exactly one process waiting to input from c and one to output to c . In terms of DJ, this ensures that a method receiver always returns its value (if ever) to the correct caller, and that a returned value always finds the initial caller waiting for it. In Figure 12, $\text{chanI}(U)$ is *linear input* of a value of type U ; $\text{chanO}(U)$ is the opponent called *linear output*. The type chan is given to channels that have matched input and output types. $\text{chanI}(U)$ is assigned to **await**, while $\text{chanO}(U)$ is to threads with/to c (either **return**(c) e , e **with/to** c , or **go** e **with/to** c).

To see the use of linear types, consider the following network; the return expression cannot determine the original location if we have two **awaits** at the same channel c , violating the linearity of c .

$$l_1[E_1[\text{await } c], \sigma_1, \text{CT}_1] \mid l_2[E_2[\text{await } c], \sigma_2, \text{CT}_2] \mid l_3[\text{go } v \text{ to } c, \sigma_3, \text{CT}_3] \quad (1)$$

The uniqueness of the returned answer is also lost if return channel c appears twice.

$$l_1[\text{return}(c) \ e_1, \sigma_1, \text{CT}_1] \mid l_2[\text{return}(c) \ e_2, \sigma_2, \text{CT}_2] \quad (2)$$

The aim of introducing linear channels is to avoid these situations during execution of runtime method invocations. The following binary operation $\asymp$ is used for controlling the composition of threads and networks.

Definition 5.1. The commutative, partial, binary composition operator on channel types, $\odot$, is defined as $\mathbf{chanI}(U) \odot \mathbf{chanO}(U) \stackrel{\text{def}}{=} \mathbf{chan}$. Then we define the composition of two channel environments $\Delta_1 \odot \Delta_2$ as:

$$\{\Delta_1(c) \odot \Delta_2(c) \mid c \in \text{dom}(\Delta_1) \cap \text{dom}(\Delta_2)\} \cup \Delta_1 \setminus \text{dom}(\Delta_2) \cup \Delta_2 \setminus \text{dom}(\Delta_1)$$

Two channel types, τ and τ' are *composable* iff their composition is defined: $\tau \asymp \tau' \iff \tau \odot \tau'$ is defined. Similarly for $\Delta_1 \asymp \Delta_2$.

Note that $\odot$ and $\asymp$ are partial operators. Hence the composition of other combinations is not allowed. Once we compose linear input and output types, then it is typed by $\mathbf{chan}$, hence it becomes uncomposable because $\mathbf{chan} \not\asymp \tau$ for any τ . Intuitively if P is typed by environment Δ_1 and Q by Δ_2 , and if $\Delta_1 \asymp \Delta_2$, then we can compose P and Q as $P|Q$ safely, preserving channel linearity. Hence (1) is untypable because of $\mathbf{chanI}(U) \not\asymp \mathbf{chanI}(U)$ at c . (2) is too by $\mathbf{chanO}(U) \not\asymp \mathbf{chanO}(U)$ at c .

5.2 Well-formedness

Well-formedness is defined for types, environments, stores and class tables in Figure 13. There are six kinds of judgement, and all are interrelated. In the following we assume α ranges over either τ, S, U or C **extends** $D \vec{T} \vec{f} \{\mathbf{m}_i : C_i \rightarrow U_i\}$.

$\Gamma; \Delta \vdash \text{Env}$	$\Gamma; \Delta$ are well-formed environments.
$\vdash \alpha : \text{tp}$	α is a well-formed type.
$\Gamma; \Delta \vdash \sigma : \text{ok}$	σ is a well-formed store in $\Gamma; \Delta$.
$\vdash \text{CSig} : \text{ok}$	CSig is a well-formed signature.
$\vdash \text{CT} : \text{ok}$	CT is a well-formed class table.

The judgements are standard. Note that CSig only contains well-formed types; and C is well-formed if its CSig entry is so.

5.3 Value and expression typing

Types are assigned to values and expressions using only the expression environment Γ . They have judgements of the form:

$$\Gamma \vdash e : \alpha \quad e \text{ has type } \alpha \text{ in expression environment } \Gamma.$$

where α ranges over T, U and S , and the typing rules are given in Figure 14. The typing judgement is lightweight or local in the sense that it does not require knowledge about method bodies held at other locations, requiring only the declared signature of the method. This is possible by the use of the class signatures and invariants as explained later.

First we focus on the key typing rule for frozen expressions, TV-FROZEN. In order for such a value to be well-typed we must ensure that the store σ and CT are well-typed, and that the expression e computes a result of the expected type when supplied its formal parameter. The simplicity of this rule comes from the

(Environments)	
$\emptyset \vdash \text{Env}$	$\frac{\Gamma \vdash \text{Env} \quad \vdash T : \text{tp} \quad x \notin \text{dom}(\Gamma)}{\Gamma, x : T \vdash \text{Env}}$
$\frac{\Gamma \vdash \text{Env} \quad \vdash C : \text{tp} \quad \text{this} \notin \text{dom}(\Gamma)}{\Gamma, \text{this} : C \vdash \text{Env}}$	$\frac{\Gamma \vdash \text{Env} \quad \vdash C : \text{tp} \quad o \notin \text{dom}(\Gamma)}{\Gamma, o : C \vdash \text{Env}}$
$\frac{\Gamma \vdash \text{Env}}{\Gamma; \emptyset \vdash \text{Env}}$	$\frac{\Gamma; \Delta \vdash \text{Env} \quad \vdash \tau : \text{tp} \quad c \notin \text{dom}(\Delta)}{\Gamma; \Delta, c : \tau \vdash \text{Env}}$
(Types)	
$\vdash \text{boolean} : \text{tp}$ $\vdash \text{void} : \text{tp}$ $\vdash \text{chan} : \text{tp}$	$\frac{\vdash U : \text{tp} \vee U \in \text{CSig}}{\vdash \text{chanI}(U) : \text{tp}}$ $\vdash \text{chan0}(U) : \text{tp}$ $\vdash \text{ret}(U) : \text{tp}$
$\frac{\vdash \text{CSig}(C) : \text{tp}}{\vdash C : \text{tp}}$	$\frac{\forall C \in \text{dom}(\text{CSig}) \vdash C : \text{tp}}{\vdash \text{CSig} : \text{ok}}$
$\frac{\vdash D : \text{tp} \quad \forall S \in \{\vec{C}, \vec{U}\} \vdash S : \text{tp} \vee S \in \text{dom}(\text{CSig}) \quad \text{fields}(C) = \vec{T}\vec{f} \quad m_i \in \text{CSig}(D) \implies \text{mtype}(m_i, C) <: \text{mtype}(m_i, D)}{\vdash [\text{remote}] C \text{ extends } D \vec{T}\vec{f} \{m_i : C_i \rightarrow U_i\} : \text{tp}}$	
(Stores)	
$\frac{\Gamma; \Delta \vdash \text{Env}}{\Gamma; \Delta \vdash \emptyset : \text{ok}}$	$\frac{\Gamma; \Delta \vdash \sigma : \text{ok} \quad \Gamma \vdash x : T \quad x \notin \text{dom}(\sigma) \quad \Gamma \vdash v : T' \quad T' <: T}{\Gamma; \Delta \vdash \sigma \cdot [x \mapsto v] : \text{ok}}$ $\frac{\Gamma \vdash o : C \quad \Gamma; \Delta \vdash \sigma : \text{ok} \quad o \notin \text{dom}(\sigma)}{\Gamma; \Delta \vdash (C, \vec{f} : \vec{v}, n, \{\vec{c}\}) : \text{ok}}$ $\frac{\Gamma; \Delta \vdash \sigma \cdot [o \mapsto (C, \vec{f} : \vec{v}, n, \{\vec{c}\})] : \text{ok}}$
$\frac{\Gamma \vdash \vec{v} : \vec{T}' \quad \vec{T}' <: \vec{T} \quad \text{fields}(C) = \vec{T}\vec{f} \quad n \geq 0 \quad \Gamma; \Delta \vdash c_i : \text{chan0}(\text{void})}{\Gamma; \Delta \vdash (C, \vec{f} : \vec{v}, n, \{\vec{c}\}) : \text{ok}}$	
(Class tables)	
$\frac{\text{this} : C, x : D \vdash e : \text{ret}(U') \quad \text{mtype}(m, C) = D \rightarrow U \quad U' <: U}{\text{this} : C \vdash U \text{ m}(D \ x)\{e\} : \text{ok in } C}$	$\frac{\text{this} : C \vdash \vec{M} : \text{ok in } C \quad \text{fields}(D) = \vec{T}'\vec{f}' \quad \text{fields}(C) = \vec{T}\vec{f} \quad K = C (\vec{T}'\vec{f}', \vec{T}\vec{f})\{\text{super}(\vec{f}'); \text{this}.\vec{f} = \vec{f}\}}{\vdash \text{class } C \text{ extends } D \{\vec{T}\vec{f}; K \vec{M}\} : \text{ok}}$
$\vdash \emptyset : \text{ok}$	$\frac{L = \text{class } C \text{ extends } D \{\vec{T}\vec{f}; K \vec{M}\} \quad \vdash L : \text{ok} \quad \vdash \text{CT}' : \text{ok}}{\vdash \text{CT}' \cdot [C \mapsto L] : \text{ok}}$

Fig. 13: Well-formedness

assumption that runtime values are created under the invariants defined in § 6. By combining with the invariants, we shall see:

- Instances of remote classes are not contained in σ , i.e. if $o \in \text{dom}(\sigma)$, then we have $\sigma(o) = (C, \dots)$ with $\text{local}(C)$. This is guaranteed by the combination of invariants from $\text{Inv}(4)$ to $\text{Inv}(8)$ in § 6.1.2.
- The closure contains no free variables and no local object-identifiers: for exam-

TV-BASIC $\frac{}{\Gamma \vdash \mathbf{Env}}$ $\Gamma \vdash \mathbf{true} : \mathbf{boolean}$ $\Gamma \vdash \mathbf{false} : \mathbf{boolean}$ $\Gamma \vdash () : \mathbf{unit}$ $\Gamma \vdash \epsilon : \mathbf{void}$	TV-NULL $\frac{\Gamma \vdash \mathbf{Env} \quad \vdash C : \mathbf{tp}}{\Gamma \vdash \mathbf{null} : C}$	TV-oid $\frac{\Gamma, o : C, \Gamma' \vdash \mathbf{Env}}{\Gamma, o : C, \Gamma' \vdash o : C}$
	TV-FROZEN $\frac{\Gamma, x : T, \vec{u} : \vec{T} \vdash e : U \quad \Gamma, \vec{u} : \vec{T}; \emptyset \vdash \sigma : \mathbf{ok} \quad \vdash \mathbf{CT} : \mathbf{ok}}{\Gamma \vdash \lambda(T x).(\nu \vec{u})(l, e, \sigma, \mathbf{CT}) : T \rightarrow U}$	
TE-VAR $\frac{\Gamma, x : T, \Gamma' \vdash \mathbf{Env}}{\Gamma, x : T, \Gamma' \vdash x : T}$	TE-THIS $\frac{\Gamma, \mathbf{this} : C, \Gamma' \vdash \mathbf{Env}}{\Gamma, \mathbf{this} : C, \Gamma' \vdash \mathbf{this} : C}$	
TE-COND $\frac{\exists S : S_1 <: S \wedge S_2 <: S \quad \Gamma \vdash e : \mathbf{boolean} \quad \Gamma \vdash e_1 : S_1 \quad \Gamma \vdash e_2 : S_2}{\Gamma \vdash \mathbf{if} (e) \{e_1\} \mathbf{else} \{e_2\} : S}$	TE-WHILE $\frac{\Gamma \vdash e_1 : \mathbf{boolean} \quad \Gamma \vdash e_2 : \mathbf{void}}{\Gamma \vdash \mathbf{while} (e_1) \{e_2\} : \mathbf{void}}$	TE-FLD $\frac{\Gamma \vdash e : C \quad \vdash C : \mathbf{tp} \quad e \neq \mathbf{this}, o \implies \mathbf{local}(C) \quad \mathbf{fields}(C) = \vec{T}\vec{f}}{\Gamma \vdash e.f_i : T_i}$
TE-SEQ $\frac{\Gamma \vdash e_1 : U \quad \Gamma \vdash e_2 : S}{\Gamma \vdash e_1; e_2 : S}$	TE-ASS $\frac{\Gamma \vdash e : T' \quad T' <: T \quad \Gamma \vdash x : T}{\Gamma \vdash x = e : T'}$	TE-FLDASS $\frac{\Gamma \vdash e.f : T \quad T' <: T \quad \Gamma \vdash e' : T'}{\Gamma \vdash e.f = e' : T'}$
TE-NEW $\frac{\mathbf{fields}(C) = \vec{T}\vec{f} \quad T'_i <: T_i \quad \Gamma \vdash e_i : T'_i \quad \vdash C : \mathbf{tp}}{\Gamma \vdash \mathbf{new} C(\vec{e}) : C}$	TE-METH $\frac{\mathbf{mtype}(m, C) = D \rightarrow U \quad \Gamma \vdash e_0 : C \quad \Gamma \vdash e : D' \quad D' <: D}{\Gamma \vdash e_0.m(e) : U}$	TE-DEC $\frac{\Gamma \vdash e : T \quad T <: T' \quad \Gamma, x : T \vdash e_0 : S}{\Gamma \vdash T' x = e; e_0 : S}$
TE-RETURN $\frac{\Gamma \vdash e : U}{\Gamma \vdash \mathbf{return} e : \mathbf{ret}(U)}$	TE-RETURNVOID $\frac{}{\Gamma \vdash \mathbf{return} : \mathbf{ret}(\mathbf{void})}$	
TE-FREEZE $\frac{\Gamma, x : T \vdash e : U}{\Gamma \vdash \mathbf{freeze}[t](T x)\{e\} : T \rightarrow U}$	TE-DEFROST $\frac{\Gamma \vdash e_0 : T' \quad T' <: T \quad \Gamma \vdash e : T \rightarrow U}{\Gamma \vdash \mathbf{defrost}(e_0; e) : U}$	
TE-FORK $\frac{\Gamma \vdash e : S}{\Gamma \vdash \mathbf{fork}(e) : \mathbf{void}}$	TE-SYNC $\frac{e \neq \mathbf{this}, o \implies \mathbf{local}(C) \quad \Gamma \vdash e_1 : C \quad \Gamma \vdash e_2 : S}{\Gamma \vdash \mathbf{sync} (e_1) \{e_2\} : S}$	TE-MONITOR $\frac{e \neq \mathbf{this}, o \implies \mathbf{local}(C) \quad \Gamma \vdash e : C}{\Gamma \vdash e.\mathbf{wait} : \mathbf{void} \quad e.\mathbf{notify} : \mathbf{void} \quad e.\mathbf{notifyAll} : \mathbf{void}}$
TE-CLASSLOAD $\frac{\Gamma \vdash e : U \quad \vdash \vec{C} : \mathbf{tp}}{\Gamma \vdash \mathbf{download} \vec{C} \mathbf{from} l \mathbf{in} e : U \quad \mathbf{resolve} \vec{C} \mathbf{from} l \mathbf{in} e : U}$	TE-INSYNC $\frac{\Gamma \vdash o : C \quad \Gamma \vdash e : S}{\Gamma \vdash \mathbf{insync} o \{e\} : S}$	TE-SANDBOX $\frac{\Gamma \vdash e : U}{\Gamma \vdash \mathbf{sandbox} \{e\} : U}$
TE-READY $\frac{\Gamma \vdash o : C \quad n > 0}{\Gamma \vdash \mathbf{ready} o n : \mathbf{void}}$	TE-HOLE $\frac{\vdash U : \mathbf{tp}}{\Gamma \vdash []^U : U}$	

Fig. 14: Value and expression typing rules

ple, by the combination of the invariants from $\text{Inv}(4)$ to $\text{Inv}(14)$ in § 6.1.4, we know $\sigma(o_i) = v_i$ is closed so that we can ensure that the resulting frozen value is closed again.

The assumption for the initial class table is more complicated as shall be explained in the next section.

5.3.1 Locality for field and thread synchronisation. There are two important restrictions which we should impose in correspondence with the current Java implementation. The first constraint is to disallow field access and assignment to a remote object in a different location. Hence the following should be prohibited even if class C is remote.

$$l[E[o.f]]P, \sigma_1, \text{CT}_1 \mid m[Q, \sigma_2 \cdot [o \mapsto (C, \dots)], \text{CT}_2] \quad (3)$$

However we wish to allow to type the following with class C remote:

$$l[E[o.f]]P, \sigma_1 \cdot [o \mapsto (C, \dots)], \text{CT}_1 \mid m[Q, \sigma_2, \text{CT}_2] \quad (4)$$

An early version of the work simply replaced the typing rule for field access with one that prevented it on any instance of a remote class. While safe this was overly restrictive, since even at the location where the remote object was held in store, no update to any of its fields could ever take place, hence (4) above was untypable.

In order to propose a typing rule to prevent remote field access statically but allow field access on remote objects locally, we require a combination of the locality invariants in § 6.1.2, the rule TE-FLD and also the initial conditions explained in Definition 6.3. The rule TE-FLD restricts field accesses only for local classes if e is neither **this** or o . The special expression **this** is allowed to have a remote class because **this** is always instantiated by an object identifier o that is present in the local store (see METHINVOKE). This constraint, together with our initial conditions guarantees that field access is always local.

The second restriction with respect to Java implementation is on thread synchronisation: performing thread synchronisation on a remote object is undefined behaviour. In Java it is possible to synchronise on the *stub* to a remote object, but this is not the same as synchronising on the actual remote object, since it does not acquire the lock on the underlying object held at the remote site and does not prevent other clients in the network from accessing that resource.

```

1  // Client 1 in Location 2
2  // ... import reference to r via RMI registry
3  synchronized (r) {
4      r.set(1);
5      return r.get();
6  }
7  // Client 2 in Location 3
8  // ... import reference to r via RMI registry
9  synchronized (r) {
10     r.set(2);
11     return r.get();
12 }
```

Listing 12: Incorrect synchronisation program

In the above, suppose we have the remote class which contains synchronised methods `set` and `get` in location 1 and two clients in locations 2 and 3. In this example the clients happen to be aware that their server is providing a shared resource, so they try to guarantee a “transaction” by “locking” the remote object. However this only locks the local stub objects, and does not prevent interleaving of operations: hence it is possible for client 1 to return 2 and client 2 to return 1. To avoid this situation by type-checking, we can just put the same condition as the field access as defined in TE-SYNC. Combining the invariants of locality, then we can now detect the above situation.

To implement a server-side locking solution would require engineering effort and an agreed protocol between clients. For instance, we consider a semaphore-style arrangement to guarantee the atomicity of a “transaction” in the following example:

```

1  // Client 1 in Location 2
2  // ... import reference to r via RMI registry
3  r.down();
4  r.set(1);
5  int v = r.get();
6  r.up();
7  return v;
8
9  // Client 2 in Location 3
10 // ... import reference to r via RMI registry
11 r.down();
12 r.set(2);
13 int v = r.get();
14 r.up();
15 return v;
```

Listing 13: Correct synchronisation program

This would require synchronised `down()` and `up()` methods to be installed in the remote object `r`, and would be very fragile since it relies on the good behaviour of clients to correctly signal the semaphore upon leaving the critical section. This option would be typable by our system, since it does not require synchronisation on the remote object `r`.

5.4 Thread and network typing

Threads, configurations and networks are assigned types under both the expression environment Γ and the channel environment Δ . The judgements take the following forms:

$\Gamma; \Delta \vdash P : \mathbf{thread}$	P is a well-typed thread in environment $\Gamma; \Delta$.
$\Gamma; \Delta \vdash F : \mathbf{conf}$	F is a wt. configuration in environment $\Gamma; \Delta$.
$\Gamma; \Delta \vdash N : \mathbf{net}$	N is a wt. network in environment $\Gamma; \Delta$.

The typing rules are given in Figure 15. The most important rule for threads is TT-PAR; we type a parallel compositions of threads if a composition of their respective channel environments preserves the linearity of channels. This is checked by $\Delta_1 \asymp \Delta_2$.

$\frac{\text{TT-NIL} \quad \Gamma; \emptyset \vdash \mathbf{Env}}{\Gamma; \emptyset \vdash \mathbf{0} : \text{thread}}$	$\frac{\text{TT-PAR} \quad \Gamma; \Delta_i \vdash P_i : \text{thread} \quad \Delta_1 \asymp \Delta_2}{\Gamma; \Delta_1 \odot \Delta_2 \vdash P_1 \mid P_2 : \text{thread}}$	$\frac{\text{TT-AWAIT} \quad \Gamma; \Delta \vdash E[]^U : \text{thread} \quad c \notin \text{dom}(\Delta)}{\Gamma; \Delta, c : \text{chanI}(U) \vdash E[\text{await } c]^U : \text{thread}}$
$\frac{\text{TT-RES} \quad \Gamma; \Delta, c : \text{chan} \vdash P : \text{thread}}{\Gamma; \Delta \vdash (\nu c)P : \text{thread}}$	$\frac{\text{TT-RETURN} \quad \Gamma \vdash e : \text{ret}(U') \quad U' <: U}{\Gamma; c : \text{chan0}(U) \vdash e[\text{return}(c)/\text{return}] : \text{thread}}$	
$\frac{\text{TT-WAITING} \quad \Gamma; \Delta \vdash E[]^{\text{void}} : \text{thread} \quad c \notin \text{dom}(\Delta) \quad n > 0}{\Gamma; \Delta, c : \text{chanI}(\text{void}) \vdash E[\text{waiting}(c) n]^{\text{void}} : \text{thread}}$	$\frac{\text{TT-FORKED} \quad \Gamma \vdash e : S}{\Gamma; \emptyset \vdash \text{forked } e : \text{thread}}$	
$\frac{\text{TT-GOSER} \quad \Gamma \vdash o : C \quad \text{remote}(C) \quad \Gamma \vdash o.m(v) : U}{\Gamma; c : \text{chan0}(U) \vdash \text{go } o.m(\text{serialize}(v)) \text{ with } c : \text{thread}}$	$\frac{\text{TT-METHWITH} \quad \Gamma \vdash o.m(v) : U}{\Gamma; c : \text{chan0}(U) \vdash o.m(v) \text{ with } c : \text{thread}}$	
$\frac{\text{TT-DESERWITH} \quad \Gamma \vdash v : \text{unit} \rightarrow D' \quad \Gamma \vdash o : C \quad D' <: D \quad \text{remote}(C) \quad \text{mtype}(m, C) = D \rightarrow U}{\Gamma; c : \text{chan0}(U) \vdash o.m(\text{deserialize}(v)) \text{ with } c : \text{thread} \quad \text{go } o.m(v) \text{ with } c : \text{thread}}$		
$\frac{\text{TT-VALTO} \quad \Gamma \vdash v : U' \quad U' <: U}{\Gamma; c : \text{chan0}(U) \vdash \text{go } \text{serialize}(v) \text{ to } c : \text{thread} \quad \text{go } v \text{ to } c : \text{thread}}$	$\frac{\text{TC-WEAK} \quad \Gamma; \Delta \vdash F : \text{conf} \quad c \notin \text{dom}(\Delta)}{\Gamma; \Delta, c : \text{chan} \vdash F : \text{conf}}$	
$\frac{\text{TC-RESC} \quad \Gamma; \Delta, c : \text{chan} \vdash F : \text{conf}}{\Gamma; \Delta \vdash (\nu c)F : \text{conf}}$	$\frac{\text{TC-RESID} \quad \Gamma, u : T; \Delta \vdash F : \text{conf} \quad u \in \text{dom}(F)}{\Gamma; \Delta \vdash (\nu u)F : \text{conf}}$	$\frac{\text{TC-CONF} \quad \Gamma; \Delta_1 \vdash P : \text{thread} \quad \Gamma; \Delta_2 \vdash \sigma : \text{ok} \quad \vdash \text{CT} : \text{ok} \quad \text{FCT} \subseteq \text{CT} \quad \Delta_1 \asymp \Delta_2}{\Gamma; \Delta_1 \odot \Delta_2 \vdash P, \sigma, \text{CT} : \text{conf}}$
$\frac{\text{TN-NIL} \quad \Gamma; \emptyset \vdash \mathbf{Env}}{\Gamma; \emptyset \vdash \mathbf{0} : \text{net}}$	$\frac{\text{TN-CONF} \quad \Gamma; \Delta \vdash F : \text{conf}}{\Gamma; \Delta \vdash l[F] : \text{net}}$	$\frac{\text{TN-PAR} \quad \Gamma; \Delta_i \vdash N_i : \text{net} \quad \text{dom}(N_1) \cap \text{dom}(N_2) = \emptyset \quad \Delta_1 \asymp \Delta_2 \quad \text{loc}(N_1) \cap \text{loc}(N_2) = \emptyset}{\Gamma; \Delta_1 \odot \Delta_2 \vdash N_1 \mid N_2 : \text{net}}$
$\frac{\text{TN-WEAK} \quad \Gamma; \Delta \vdash N : \text{net} \quad c \notin \text{dom}(\Delta)}{\Gamma; \Delta, c : \text{chan} \vdash N : \text{net}}$	$\frac{\text{TN-RESID} \quad \Gamma, u : T; \Delta \vdash N : \text{net} \quad u \in \text{dom}(N)}{\Gamma; \Delta \vdash (\nu u)N : \text{net}}$	$\frac{\text{TN-RESC} \quad \Gamma; \Delta, c : \text{chan} \vdash N : \text{net}}{\Gamma; \Delta \vdash (\nu c)N : \text{net}}$

Fig. 15: Thread and network typing rules

We must make a similar check in TC-CONF, since the blocked queue of threads waiting for locks requires the use of a channel environment to type the store σ . A configuration is then well-typed in an environment $\Gamma; \Delta_1 \odot \Delta_2$ if its threads, P , are well typed in the environment $\Gamma; \Delta_1$ and its store σ is well-typed under $\Gamma; \Delta_2$ with $\Delta_1 \asymp \Delta_2$. The class table must also be well-formed, and must contain a copy of the foundation classes FCT. The rule TN-CONF promotes configurations to the network level. For the rule TN-PAR, we use the set of location names in a network N are given by the function $\text{loc}(N)$, defined as: $\text{loc}(\mathbf{0}) = \emptyset$, $\text{loc}(l[F]) = \{l\}$, $\text{loc}(N_1 \mid N_2) = \text{loc}(N_1) \cup \text{loc}(N_2)$ and $\text{loc}((\nu u)N) = \text{loc}(N)$. We also use $\Delta_1 \asymp \Delta_2$ as TC-CONF

to check the composability.

6. NETWORK INVARIANTS AND TYPE SOUNDNESS

This section presents the main technical results of the present paper. We first introduce several runtime invariants and show that if an initial network satisfies certain conditions then reductions always preserve these runtime invariants. Next we establish subject reduction by the use of invariants. Finally combining subject reduction and invariants, we derive progress and other safety guarantees.

6.1 Network invariants

We start from the definition of a property over networks, given in Definition 6.1.

Definition 6.1. (Properties) Let ψ denote a property over networks (i.e. ψ is a subset of networks). We write $N \models \psi$ if N satisfies ψ (i.e. if $N \in \psi$); we also write $N \not\models \psi$ if N does not satisfy ψ . We define the error property **Err** as the set of the networks which contain **Error** as subexpression, i.e. $\text{Err} = \{N \mid N \equiv (\nu \vec{u})(l[E[\text{Error}] \mid P, \sigma, \text{CT}] \mid N')\}$. We say ψ is a *network invariant with an initial property* ψ_0 if $\psi = \{N \mid \exists N_0. (N_0 \models \psi_0, N_0 \rightarrow N, N \not\models \text{Err})\}$

The following lemma is needed to formulate the network invariants of DJ. This concerns the *canonical forms*: every typable network can be written in such a form. Intuitively, a canonical form is one in which all restricted identifiers are moved out to the network level.

LEMMA 6.1. (Canonical forms) *Suppose that $\Gamma; \Delta \vdash N : \text{net}$ then we have $N \equiv (\nu \vec{u})(\prod_{0 \leq i < n} l_i[P_i, \sigma_i, \text{CT}_i])$ where n denotes the number of locations in N .*

PROOF. By induction on the number of networks in parallel, n . $\square$

The following lemma states that the typability is preserved under the structure rules. By this and the above lemma, we only have to consider the canonical forms for defining the network invariants.

LEMMA 6.2. (Structural equivalence preserves typability)

- (1) *If $\Gamma; \Delta \vdash F : \text{conf}$ and $F \equiv F'$ then $\Gamma; \Delta \vdash F' : \text{conf}$.*
- (2) *Assume $\Gamma; \Delta \vdash P : \text{thread}$ and $P \equiv P'$, then we have $\Gamma; \Delta \vdash P' : \text{thread}$.*
- (3) *If $\Gamma; \Delta \vdash N : \text{net}$ and $N \equiv N'$ then $\Gamma; \Delta \vdash N' : \text{net}$.*

PROOF. By induction on typing derivations paying attention to the last rule applied. See Appendix D.3. $\square$

In order to ensure the correct execution of networks and the preservation of safety, we require certain properties to remain invariant.

Definition 6.2. (Network invariants) Given network $N \equiv (\nu \vec{u})(\prod_{0 \leq i < n} l_i[F_i])$ with $F_i = (P_i, \sigma_i, \text{CT}_i)$, and assuming $0 \leq j < n$, $i \neq j$ where required, we define property $\text{Inv}(r)$ as a set of networks which satisfy the condition r (with $1 \leq r \leq 17$) as defined below.

The majority of these properties fall into one of three important categories: *class availability*, *locality* and *linearity*. Each invariant has a clear operational (and arguably engineering) meaning.

6.1.1 Class availability

Inv(1) $\text{FCT} \subseteq \text{CT}_i$

Inv(2) $P_i \equiv E[\text{new } C(\vec{v})] \mid Q_i \implies \text{comp}(C, \text{CT}_i)$

Inv(3) $C \in \text{dom}(\text{CT}_i) \cap \text{dom}(\text{CT}_j) \implies$
 $\text{CT}_i(C) = \text{CT}_j(C) \vee \text{CT}_i(C) = \text{CT}_j(C)[\vec{D}^{l_i} / \vec{D}]$ with $\text{fcl}(\text{CT}_i(C)) = \{\vec{D}\}$

Key invariant properties in the presence of distribution are those of *class availability*. For example when a class is needed, it and all its superclasses must be present in the local class table. This requirement eliminates erroneous networks containing locations such as: $l[E[\text{new } C(\vec{v})], \sigma, \emptyset]$ where class C is not present in l 's empty class table, so the initial step of execution will cause a crash. Note that even if C is present, if its superclass D is not then this is also an unexpected state. For example, in our system Inv(2) says that if we attempt to instantiate C , we need to have all its superclasses.

Inv(3) models the strict default class version control of the Java serialisation API. For example suppose we serialise an instance of the following class:

```

1 class A implements java.io.Serializable {
2   private int i;
3   private int j = 0;
4   A(int i) { this.i = i; }
5 }
```

If we then pass this to a remote consumer who has also has a class A, then deserialisation is not guaranteed to succeed, even if they have a binary compatible copy of the class:

```

1 class A implements java.io.Serializable {
2   private int i;
3   A(int i) { this.i = i; }
4 }
```

This is because it is impossible to recreate the original A at the new site without special low level programming. Moreover the `serialVersionUID`—a long integer hash value computed from the structure of a class file—will differ between the serialised object and the version of A held by the consumer [Greanier 2005].¹

6.1.2 Locality

Inv(4) $\text{fv}(P_i) \subseteq \text{dom}(\sigma_i) \subseteq \{\vec{u}\}$

Inv(5) $\text{dom}(\sigma_i) \cap \text{dom}(\sigma_j) = \emptyset$

Inv(6) $o \in \text{fn}(F_i) \cap \text{fn}(F_j) \implies \exists! k. \sigma_k(o) = (C, \dots) \wedge \text{remote}(C)$

Inv(7) $o \in \text{fn}(F_i) \wedge \exists k. \sigma_k(o) = (C, \dots) \wedge \text{local}(C) \implies k = i$

Inv(8) $o \in \text{fn}(F_i) \implies \exists k \ 1 \leq k \leq n. o \in \text{dom}(\sigma_k)$

¹It is possible to override this value at the programmer level, however we do not consider such advanced techniques for versioning serialised objects.

Inv(9) Suppose

$$R_i \in \{ o.m(e) \text{ with } c, E[o.f], E[o.f = e], E[\text{sync } (o) \{e\}], \\ E[\text{insync } o \{e\}], E[o.\text{notify}], E[o.\text{notifyAll}], E[o.\text{wait}], E[\text{ready } o \ n] \} \\ \text{Then } P_i \equiv Q_i \mid R_i \implies \sigma_i(o) = (C, \dots) \wedge \text{comp}(C, \text{CT}_i)$$

An important property in the system is the locality of store entries such as local variables and object identifiers, captured by these invariants. For instance, combining Inv(4) and Inv(5), we can derive $\text{fv}(P_i) \cap \text{fv}(P_j) = \emptyset$, which ensures that local variables are not shared between threads at different locations. In Inv(9) we ensure that non-remote operations like field access and thread synchronisation are not attempted on remote object references. This particular situation highlights the necessity of the invariants, since we cannot guarantee this property alone in the typing system as we discussed in § 5.3.

6.1.3 Linearity invariants. Below we say *thread P inputs at c* if $P \equiv E[\text{await } c] \mid R$ or $P \equiv E[\text{waiting}(c) \ n] \mid R$ for some E and R ; dually *thread P outputs at c* if $P \equiv R \mid Q$ with $R \equiv \text{return}(c) \ e$ or $R \equiv \text{go } e/e \text{ with/to } c$ for some Q and e .

Inv(10) $P_i \equiv Q_i \mid R_i$ and Q_i inputs at $c \implies$ neither R_i nor P_j inputs at c .

Inv(11) $P_i \equiv Q_i \mid R_i$ and Q_i outputs at $c \implies$ neither R_i nor P_j outputs at c .

Linearity of channel usage ensures the determinacy of method calls and returns and also the notification of blocked threads. This is ensured by the linear type checking.

6.1.4 Closure and lock invariants

Closures

Inv(12) $P_i \equiv E[v] \mid Q_i$ then $\text{fv}(v) = \emptyset$

Inv(13) $\sigma_i(x) = v \implies \text{fv}(v) = \emptyset$

Inv(14) $\sigma_i(o) = (C, \vec{f} : \vec{v}) \implies \text{fv}(v_j) = \emptyset$

Inv(15) $P_i \equiv E[\lambda(T \ x).(\nu \vec{u})(l, e, \sigma, \text{CT})] \mid Q_i$ and $\text{fn}(\lambda(T \ x).(\nu \vec{u})(l, e, \sigma, \text{CT})) = \{\vec{u}'\}$ implies $\exists k. \sigma_k(u'_j) = (C_j, \dots)$ with $\text{remote}(C_j)$.

Locks

Inv(16) $P_i \equiv E[\text{ready } o \ n] \mid Q_i \implies \text{insync}(o, E) \wedge n > 0$

Inv(17) $P_i \equiv E[\text{waiting}(c) \ n] \mid Q_i \implies \exists ! o.c \in \text{blocked}(\sigma_i, o) \wedge \text{insync}(o, E) \wedge n > 0$

The *closure* invariants ensure that values and store entries do not contain any unbound variables. This is important to guarantee that newly created frozen expressions are similarly closed.

The *lock* invariants ensure the correct behaviour of the locking primitives at runtime. Inv(16) ensures that a thread that is ready to reacquire a lock will set that lock's count to a non-zero number. Inv(17) ensures that a thread does not wait for a non-existent lock.

6.2 Initial network

Before proving the network invariant, we define the initial network configurations. Roughly speaking an initial configuration contains no runtime values and expressions except o-ids. It can, however, contain parallel threads distributed among

locations; these have been generated by compiling multiple user-defined main programs. Definition 6.3 states these conditions formally.

Definition 6.3. (Initial network) We call network $N \equiv (\nu \vec{u})(\prod_{0 \leq i < n} l_i[P_i, \sigma_i, \mathbf{CT}_i])$ an *initial network* if it satisfies the following conditions (called *initial properties*):

- it contains no runtime expressions or values except o-ids and parallel compositions of $\mathbf{return}(c)$ e ; and $\mathbf{freeze}[t](T \ x)\{e\}$ does not contain free o-ids, i.e. $\mathbf{fn}(e) = \emptyset$.
- it satisfies all properties $\mathbf{Inv}(i)$ except $\mathbf{Inv}(2)$, which is replaced by:
 - (a) $\mathbf{fcl}(P_i) \subseteq \mathbf{dom}(\mathbf{CT}_i)$,
 - (b) $C \in \mathbf{fcl}(\mathbf{CT}_i) \cup \mathbf{dom}(\mathbf{CT}_i) \implies \mathbf{comp}(C, \mathbf{CT}_i)$ and
 - (c) $\sigma_i(o) = (C, \dots) \implies \mathbf{comp}(C, \mathbf{CT}_i)$.
- we also strengthen the $\mathbf{Inv}(9)$ by replacing the reduction context E by the arbitrary context $\mathcal{C}$ in R_i .

We denote the set of networks satisfying these conditions by $\mathbf{Init}$.

The extra requirement states that all initial class tables are complete w.r.t. classes in the program and stores. For example, suppose

```
new A().m(),  $\emptyset$ , CT
with CT(A) = class A extends B {; void m(){new C();return}}
```

First A should be defined in $\mathbf{CT}$ (this is ensured by (a) in $\mathbf{Inv}(2')$); secondly B should be also defined in $\mathbf{CT}$ (this is ensured by (a) and (b): since $A \in \mathbf{dom}(\mathbf{CT})$, we have $\mathbf{comp}(A, \mathbf{CT})$, which implies $B \in \mathbf{dom}(\mathbf{CT})$); and thirdly, C should be defined in $\mathbf{CT}$ too since $\mathbf{new} C()$ appears after the method invocation at $\mathbf{m}$. This condition is ensured by (b) since $C \in \mathbf{fcl}(\mathbf{CT})$. The condition (c) is similarly understood. We also note that during runs of programs, the initial properties may *not* be satisfied since classes can be downloaded lazily. Later we formalise this situation in Lemma 6.3 and prove the invariant $\mathbf{Inv}(2)$. The initial condition of $\mathbf{Inv}(9)$ is similarly understood as (c).

6.3 Type soundness and progress properties

To prove some cases of the subject reduction theorem, we require some invariants to hold in the assumptions. Therefore the proof routine for type soundness is divided into the following three steps:

Step 1 We prove one step invariant property for a typed network starting from the initial properties. This step has two sub-cases:

(i) Assume $\Gamma; \Delta \vdash N_0 : \mathbf{net}$ and N_0 satisfies the initial properties. Then $N_0 \longrightarrow N_1$ implies $N_1 \models \mathbf{Inv}(r)$ for each $1 \leq r \leq 17$ if $N_1 \not\models \mathbf{Err}$.

(ii) Assume $\Gamma; \Delta \vdash N_m : \mathbf{net}$ ($m \geq 1$) and $N_m \models \mathbf{Inv}(r)$ for all $1 \leq r \leq 17$. Then $N_m \longrightarrow N_{m+1}$ implies $N_{m+1} \models \mathbf{Inv}(r)$ for each $1 \leq r \leq 17$ if $N_{m+1} \not\models \mathbf{Err}$.

Step 2 We prove the subject reduction theorem using Step 1, i.e. $\Gamma; \Delta \vdash N : \mathbf{net}$ and $N \longrightarrow N'$ implies $\Gamma; \Delta \vdash N' : \mathbf{net}$.

Step 3 Then invariant of $\mathbf{Inv}(r)$ is a corollary of Steps 1 and 2.

The proof of **Step 1** is given in the next subsection. Then assuming this holds, the proof of **Step 2** proceeds by induction on the derivation of reduction with a case analysis on the final typing rule applied. It is given in § 6.5.

6.4 Proofs of network invariants

This subsection lists the key additional invariants related to dynamic downloading of classes and synchronisations which are used for the main proofs of class invariants, lock invariants and progress, respectively. We shall use the notation P_{im} to denote the threads at location i after m reduction steps. For proofs, see Appendix E.

LEMMA 6.3. (Class table properties) *Assume:*

$\Gamma; \Delta \vdash N_k : \mathbf{net}$ for $0 \leq k \leq m$, $N_0 \in \mathbf{Init}$, $N_k \in \mathbf{Inv}(r)$ for $k > 0$, $1 \leq r \leq 17$
 $N_0 \rightarrow N_m \rightarrow N_{m+1} \equiv (\nu \vec{u}_{m+1})(\prod_{0 \leq i < n} l_i[P_{im+1}, \sigma_{im+1}, \mathbf{CT}_{im+1}])$ with $m > 0$

Then we have:

- (1) $\mathbf{CT}_{im} \subseteq \mathbf{CT}_{im+1}$.
- (2) $C \in \mathbf{fd}(P_{im+1})$ implies $C \in \mathbf{dom}(\mathbf{CT}_{im+1})$.
- (3) Assume $\mathbf{reachable}(\sigma_{im+1}, P_{im+1}, o)$ and $\sigma_{im+1}(o) = (C, \dots)$. Then we have either
 - (a) $\mathbf{comp}(C, \mathbf{CT}_{im+1})$ or
 - (b) either $P_{im+1} \equiv E[\mathbf{download} \vec{C} \text{ from } l_j \text{ in } e] \mid Q_{im+1}$ or
 $P_{im+1} \equiv E[\mathbf{resolve} \vec{C} \text{ from } l_j \text{ in } e] \mid Q_{im+1}$
 where $\exists D \in \vec{C}. C <: D$ and $\neg \mathbf{reachable}(\sigma_{im+1}, Q_{im+1}, o)$
- (4) $\exists N_k \equiv (\nu \vec{u}_k)(\prod_{0 \leq i < n} l_i[P_{ik}, \sigma_{ik}, \mathbf{CT}_{ik}])$ with $P_{ik} \equiv E[\mathbf{download} \vec{C} \text{ from } l_j \text{ in } e] \mid Q_{ik}$
 and $0 \leq j < n$ implies

$$\forall C_z \in \{\vec{C}\}. \forall C'. C_z <: C'. \exists N_m. N_k \rightarrow N_m \text{ and}$$

$$P_{im} \equiv E[\mathbf{resolve} \vec{D} \text{ from } l_j \text{ in } e] \mid Q_{im} \text{ with } C' \in \{\vec{D}\} \text{ and } C' \in \mathbf{dom}(\mathbf{CT}_{im})$$

Lemma 6.3 says (1) the class table at each location always increases; (2) if a class name appears free in a thread, then it is always in the domain of the class table; (3) if a free name or variable in P_{im+1} is reachable to o through store σ_{im+1} , then the class of o is complete otherwise it is in the middle of downloading. (4) all superclasses will be eventually downloaded if no error occurs.

The next lemma states that the number of entries by a thread to an object's monitor is correctly accounted by said object.

LEMMA 6.4. (Lock coherence) *Assume $\Gamma; \Delta \vdash N_k : \mathbf{net}$ ($0 \leq k \leq m$), N_0 satisfies the initial network conditions and $\mathbf{Inv}(r) \models N_k$ for the invariants indexed over by r . Assume $N_0 \rightarrow N_1 \rightarrow \dots \rightarrow N_{m+1} \equiv (\nu \vec{u}_{m+1})(\prod_{0 \leq i < n} l_i[P_{im+1}, \sigma_{im+1}, \mathbf{CT}_{im+1}])$ with $\mathbf{Err} \not\models N_k$.*

- (1) If $P_{im+1} \equiv E_1[\mathbf{insync} o \{ \dots E_p[\mathbf{insync} o \{e\}] \dots \}] \mid Q_{im+1} \wedge e \neq E'[\mathbf{insync} o \{e'\}]$ then:
 - (a) $e \neq E[\mathbf{waiting}(c) n']$ with $c \in \mathbf{blocked}(o, \sigma_{im+1})$ and $e \neq E[\mathbf{ready} o \dots]$.
 implies $\mathbf{getLock}(\sigma_{im+1}, o) = p$,
 - (b) $e = E[\mathbf{ready} o n'] \implies p = n'$,

(c) $e = E[\text{waiting}(c) \ n']$ with $c \in \text{blocked}(o, \sigma_{im+1})$ implies $p = n'$.

(2) Suppose $\text{getLock}(\sigma_{im+1}, o) = p$ and $p > 0$. Then:

$$P_{im+1} \equiv E_1[\text{insync } o \ \{ \dots E_p[\text{insync } o \ \{e\}] \dots \}] \mid Q_{im+1} \wedge e \neq E'[\text{insync } o \ \{e'\}]$$

Summary of the one-step invariant proof. We summarise the proofs of **Step 1** for each invariant. We use the induction on the number of reduction steps, examining the last applied reduction rule. The proof requires a careful case analysis since several invariants and typing rules are mutually related. Below “we use $\text{Inv}(r)$ ” means that “we assume $\text{Inv}(r)$ holds at the inductive step m ”; and “the case of the rule (r) ” means that “the case when the last applied rule is (r) ”.

$\text{Inv}(1)$ and $\text{Inv}(2)$ use Lemma 6.3 (1). For $\text{Inv}(3)$, we analyse two rules, **DOWNLOAD** and **DEFROST**, which changes the class table. $\text{Inv}(4)$ requires a case analysis on the three rules, **DEC**, **METHINVOKE** and **DEFROST**, with which the set of free variables of a term changes. For all cases, we use $\text{Inv}(12)$. $\text{Inv}(5)$ only requires examination of the case of **DEFROST**. For $\text{Inv}(6)$, we analyse **METHREMOTE** and **RETURN**, assuming $\text{Inv}(8)$, $\text{Inv}(5)$ and $\text{Inv}(15)$. The interesting case for $\text{Inv}(7)$ is when o newly appears at the $m + 1$ -step. We have four such cases, **NEW**, **DEFROST**, **LEAVE** and **METHREMOTE**. For all cases, we use $\text{Inv}(15)$. $\text{Inv}(8)$ is mechanical by examination of the rules for structural equivalence. $\text{Inv}(9)$ is one of the most non-trivial invariants. We derive it from Lemma 6.3 (1, 3), assuming $\text{Inv}(2)$, $\text{Inv}(8)$ and $\text{Inv}(7)$ hold at the m th-step. $\text{Inv}(10)$ and $\text{Inv}(11)$ are straightforward by the definition of $\Delta_1 \asymp \Delta_2$. $\text{Inv}(12)$ requires investigation of the cases where a value comes into a redex position. We have five cases, and use $\text{Inv}(13)$ and $\text{Inv}(14)$. For $\text{Inv}(13)$, we check the cases where new variable mappings are added to the store, or when an existing mapping is changed. We have three cases, **DEC**, **DEFROST** and **ASS**, and all use $\text{Inv}(12)$. $\text{Inv}(14)$ needs to check **NEW**, **DEFROST** and **FLDASS**. All are straightforward by application of $\text{Inv}(12)$. For $\text{Inv}(15)$, the only interesting case is **FREEZE**, and we use Lemma 4.1. For $\text{Inv}(16)$ and $\text{Inv}(17)$, we use Lemma 6.4.

6.5 Proofs of type soundness

We first prove the following standard substitution lemma. Below α denotes either U or T .

LEMMA 6.5. (Substitution and context lemma)

- (1) Assume $\Gamma, x : T \vdash e : \alpha$ and $\Gamma \vdash v : T'$. Suppose that e does not contain $x = e'$ or $T \ x = e'$ as its subterm. Then we have $\Gamma \vdash e[v/x] : \alpha'$ for some $\alpha' <: \alpha$.
- (2) $\Gamma, \text{this} : C \vdash e : \alpha$ and $\Gamma \vdash o : C'$ with $C' <: C$ imply $\Gamma \vdash e[o/\text{this}] : \alpha'$ for some $\alpha' <: \alpha$.
- (3) $\Gamma \vdash E[\]^U : \alpha$ and $\Gamma \vdash e : U'$ with $U' <: U$ iff $\Gamma \vdash E[e]^U : \alpha$.

PROOF. (1,2) By induction on the structure of the expression e using Lemma 6.2. See Appendix F.1. (3) is by induction on the structure on E . All proofs are mechanical. $\square$

Now we achieve the main theorem.

THEOREM 6.1. (**Subject reduction**)

- (1) Assume $\Gamma, \vec{u} : \vec{T} \vdash e : \alpha$, $\Gamma, \vec{u} : \vec{T} \vdash \sigma : \text{ok}$ and $\vdash \text{CT} : \text{ok}$. Suppose $(\nu \vec{u})(e, \sigma, \text{CT}) \longrightarrow_l (\nu \vec{u}')(e', \sigma', \text{CT}')$ and $e' \not\equiv \text{Err}$. Then we have $\Gamma, \vec{u}' : \vec{T}' \vdash e' : \alpha'$ for some $\alpha' < \alpha$, $\Gamma, \vec{u}' : \vec{T}' \vdash \sigma' : \text{ok}$ and $\vdash \text{CT}' : \text{ok}$.
- (2) Assume $\Gamma; \Delta \vdash F : \text{conf}$, $F \longrightarrow_l F'$ and $F' \not\equiv \text{Err}$. Then we have $\Gamma; \Delta \vdash F' : \text{conf}$.
- (3) Assume $\Gamma; \Delta \vdash N : \text{net}$, $N \longrightarrow N'$ and $N' \not\equiv \text{Err}$. Then we have $\Gamma; \Delta \vdash N' : \text{net}$.

Note that the above theorem guarantees type safety: if there is neither a null pointer error nor an unavoidable network error (i.e. $N' \not\equiv \text{Err}$), then the typability ensures that an execution does not go wrong.

As a corollary we derive:

COROLLARY 6.1. (*Network invariant*) $\wedge_{1 \leq r \leq 17} \text{Inv}(r)$ is a network invariant with the initial network properties Init defined in Definition 6.3.

6.6 Progress and Linearity Properties

Finally we can derive the following advanced progress and linearity properties.

Definition 6.4. (Progress invariants) Given network $N \equiv (\nu \vec{u})(\prod_{0 \leq i < n} l_i[P_i, \sigma_i, \text{CT}_i])$, and assuming $0 \leq k < n$, we define property $\text{Prog}(r)$ as a set which satisfy the following conditions.

- Prog(1)** $P_i \equiv E[\text{new } C(\vec{v})] \mid Q_i \implies C \in \text{dom}(\text{CT}_i)$
Classes are always available for instantiation.
- Prog(2)** $P_i \equiv E[\text{download } \vec{C} \text{ from } l_k \text{ in } e] \mid Q_i \implies \vec{C} \in \text{dom}(\text{CT}_i) \cup \text{dom}(\text{CT}_k)$
Download operations always succeed in retrieving the required classes from the specified location.
- Prog(3)** $P_i \equiv E[\text{resolve } \vec{C} \text{ from } m \text{ in } e] \mid Q_i \implies \vec{C} \in \text{dom}(\text{CT}_i)$
No attempt is made to resolve classes that are not available in the local class table.
- Prog(4)** $P_i \equiv E[o.f_j] \mid Q_i \implies [o \mapsto (C, \dots)] \in \sigma_i \wedge \text{fields}(C) = \vec{T}\vec{f}$
No attempt is made to invoke a field access on the store if the class of the store does not provide that field.
- Prog(5)** $P_i \equiv E[o.f_j := v] \mid Q_i \implies [o \mapsto (C, \dots)] \in \sigma_i \wedge \text{fields}(C) = \vec{T}\vec{f}$
No attempt is made to invoke a field access on the store if the class of the store does not provide that field.
- Prog(6)** $P_i \equiv E[x] \mid Q_i \implies x \in \text{dom}(\sigma_i)$
Expressions only access variables they are local to.
- Prog(7)** $P_i \equiv E[x := v] \mid Q_i \implies x \in \text{dom}(\sigma_i)$
Expressions only assign to variables they are local to.
- Prog(8)** $P_i \equiv o.m(v) \text{ with } c \mid Q_i \wedge \sigma_i(o) = (C, \dots) \implies \text{mbody}(m, C, \text{CT}_i) \text{ defined}$
No attempt is made to invoke a method on an object of a given class if that class does not provide that method.
- Prog(9)** $P_i \equiv \text{go } o.m(v) \text{ with } c \mid Q_i \implies \exists! k. o \in \text{dom}(\text{CT}_k)$
Remote method invocations always refer to a unique live location in the network.

Prog(10) $P_i \equiv \text{go } v \text{ to } c \mid Q_i \wedge c \in \{\bar{u}\} \implies \exists!k. P_k \equiv E[\text{await } c] \mid Q_k$

If a method return exists, there must be exactly one location waiting for it on that channel.

THEOREM 6.2. (Progress, locality and linearity)

$\bigwedge_{1 \leq r \leq 10} \text{Prog}(r)$ is a network invariant with the initial network properties Init defined in Definition 6.3.

PROOF. Immediately Prog(1) is derived from Inv(2). Prog(2) is by the monotonicity of the class tables proved in Lemma 6.3 (1). Prog(3) is obvious by DOWNLOAD. Prog(4) and Prog(5) are proved by Inv(9). Prog(6) and Prog(7) are obvious by Inv(4). Prog(8) is derived from Inv(9). Prog(9) is by combining Inv(8) and Inv(5). Prog(10) is straightforward by combining Inv(10) and Inv(11). $\square$

6.7 Progress with synchronisation and normal forms

In this last subsection, we first investigate a simple progress property in the presence of the synchronisation primitives. Then we will show the normal form of DJ—the form of a whole network when computation terminates. We start from the first proposition which says that one monitor is held by only one thread.

PROPOSITION 6.1. (Mutual exclusion) For a location $l[P, \sigma, \text{CT}]$, suppose

if $P \equiv E_1[\text{insync } o \{e_1\}] \mid \dots \mid E_n[\text{insync } o \{e_n\}] \mid Q$
 then $\forall j. 1 \leq j \leq n. (e_j = E'_j[\text{waiting}(c) \dots] \vee e_j = E'_j[\text{ready } o \dots])$
 or $\exists! j. 1 \leq j \leq n. (e_j \neq E'_j[\text{waiting}(c) \dots] \wedge e_j \neq E'_j[\text{ready } o \dots])$

with $c \in \text{blocked}(o, \sigma)$.

PROOF. See Appendix E.4. $\square$

Below we list a simple progress property. (1) states if expression e which holds a monitor is neither error nor a synchronisation expression, then e can always progress; and (2) says that a thread can exit the monitor if the monitor is only hold by threads who are going to exit.

PROPOSITION 6.2. (Progress with synchronisation) For a location $l[P, \sigma, \text{CT}]$,

(1) Suppose $P \equiv E[\text{insync } o \{E'[e]\}] \mid Q_i, e \mid Q, \sigma, \text{CT} \longrightarrow e' \mid Q', \sigma', \text{CT}'$, and $e \notin \{\text{insync } o' \{e'\}, \text{waiting}(c) n, \text{ready } o n, \text{sync } (o') \{e'\}, \text{Error}\}$. Then we have $E[\text{insync } o \{E'[e]\}] \mid Q, \sigma, \text{CT} \longrightarrow E[\text{insync } o \{E'[e']\}] \mid Q', \sigma', \text{CT}'$.

(2) Suppose $P \equiv E[\text{ready } o n] \mid Q$. Assume if $Q \equiv E[\text{insync } o \{e'\}] \mid R_i$ then $e' \in \{E'[\text{ready } o n'], E'[\text{waiting}(c) n']\}$. Then we have: $E[\text{ready } o n] \mid Q, \sigma, \text{CT} \longrightarrow E[\epsilon] \mid Q, \sigma, \text{CT}$

PROOF. (1) If P_i satisfies the assumption, then by Proposition 6.1, $E[\text{insync } o \{E'[e]\}]$ is only the thread which holds the monitor o . Hence progress is obvious by the definition of $\longrightarrow$. (2) is by $\text{getLock}(\sigma, o) = 0$. $\square$

In the presence of synchronisation, there would be no progress in the program even if it is well-typed and does not reach an error state. For example, threads may

deadlock naturally by requesting monitors in a certain order, stopping them from proceeding forever: a simple example is

$E[\text{insync } o \{ \text{sync } (o) \{ e \} \}]$. Also waiting processes $\text{waiting}(c) n$ may not proceed forever because of a lack of “notify” (i.e. lost-wakeup). Then, as its consequence, $\text{ready } o n$ may never exit. We can define the states of deadlock and liveness, and prove a general progress property under a certain kind of scheduling. We shall leave this topic to a forthcoming exposition. We now concludes the normal form theorem.

THEOREM 6.3. (Normal forms) *Assume $N_0 \models \text{Init}$ and $N_0 \rightarrow N \not\rightarrow$ and $N \not\models \text{Err}$. Then we have $N \equiv (\nu \vec{u})(\prod_{0 \leq i < n} l_i(\prod_{0 \leq j_i < n_i} P_{j_i}, \sigma_i, \text{CT}_i))$ with P_{j_i} is either $\text{go } v \text{ to } c$ or $E[\text{insync } o \{ e \}]$ with $e \in \{ \text{waiting}(o) n, \text{ready } o n, \text{sync } (o') \{ e' \} \}$.*

PROOF. By induction on N . By the initial condition Init , we can set $\Delta = \vec{c}' : \text{chan} \cup \vec{c}_i : \text{chan0}(\vec{U}_i)$. The proof is direct from the progress properties. We only investigate the cases that the reduction happens across different networks. Suppose, for example, by contradiction, that $N \not\rightarrow$ but there exists P_i such that $P_i \equiv o.m(\vec{v})$ with $c \mid Q_i$. If o is the local object id, then $N \rightarrow N'$ by **Prog(8)**. Assume that o is a remote o-id and $o \notin \text{dom}(\sigma_i)$. This time by **METHREMOTE**, $N \rightarrow N'$, contradiction. Next suppose there exists P_i such that $P_i \equiv \text{go } v \text{ to } c \mid Q_i$ with $c \in \{ \vec{u} \}$ or $c : \text{chan} \in \Delta$. Then by **Prog(10)**, there exists k such that $P_k \equiv E[\text{await } c] \mid Q_k$. Then we can apply **RETURN**, hence a contradiction. The unicity of $\text{go } v_{j_i} \text{ to } c_{j_i}$ is derived by **Inv(11)**. $\square$

7. JUSTIFICATION OF OPTIMISATION

We prove the correctness of the optimised code in § 2 using sound syntactic transformation rules over programs and runtime. The key idea is a use of the following *noninterference property* [Jones 1993; Reynolds 1978] to justify the correctness of these rules. Let us write $N \mapsto N'$ for a transformation rule of the optimisation from N to N' . Once we check $N \mapsto N'$ is type-preserving and satisfies the following noninterference property then N and N' are immediately observationally equal, hence the transformation is semantics-preserving.

if $N \rightarrow N_1$ and $N \mapsto^* N_2$, then $N_1 \equiv N_2$ or there exists N' such that $N_1 \mapsto^* N'$ and $N_2 \rightarrow N'$.

For tractable reasoning, we introduce several syntactic transformation rules satisfying this property. By the use of these equational laws, which come from those of the linear types of mobile processes [Kobayashi et al. 1996; Yoshida et al. 2001], justifications are carried out syntactically using $\mapsto$.

7.1 Observational congruence

We define an observational congruence over the typed language and runtime by applying the equational theory of process algebra [Honda and Yoshida 1995]. Hereafter we assume all networks are typed and started executing from the initial condition Init in Definition 6.3.

Definition 7.1. (Typed relations) A relation $\mathcal{R}$ over networks is *typed* when $\Gamma_1; \Delta_1 \vdash N_1 \mathcal{R} \Gamma_2; \Delta_2 \vdash N_2$ implies $\Gamma_1 = \Gamma_2$ and $\Delta_1 = \Delta_2$. We write $\Gamma; \Delta \vdash$

$N_1 \mathcal{R} N_2$ when $\Gamma; \Delta \vdash N_1$ and $\Gamma; \Delta \vdash N_2$ are related by a typed relation $\mathcal{R}$. A *typed congruence* is a typed relation $\mathcal{R}$ which is an equivalence closed under all typed contexts and the structure rules, i.e. $\equiv \subseteq \mathcal{R}$.

By the subject reduction theorem, we immediately know $\longrightarrow$ is a typed relation.

The formulation of behavioural equality is based on two conditions: reduction-closedness and an observational predicate. In the distributed setting, terms can effectively change meaning (for example by side-effecting a store), so we define “equality” to mean that two equated programs go to an equated state again. The second condition comes from the concept of observation in mobile process theory [Honda and Yoshida 1995]. For an observation, we take the output (“go”) to channel c .

Definition 7.2. (Reduction-closedness and the observational predicate)

- A typed congruence $\mathcal{R}$ on networks is *reduction-closed* whenever $\Gamma; \Delta \vdash N_1 \mathcal{R} N_2$, $N_1 \longrightarrow N'_1 \not\models \text{Err}$ implies, for some N'_2 , $N_2 \longrightarrow N'_2$ with $\Gamma; \Delta \vdash N'_1 \mathcal{R} N'_2$; and its symmetric case.
- We define the observational predicate $\downarrow_c$ and $\Downarrow_c$ as follows.

$$\begin{aligned} N \downarrow_c & \text{ if } N \equiv (\nu \vec{u})(l[\text{go } v \text{ to } c \mid P, \sigma, \text{CT}] \mid N') & \text{ with } c \notin \{\vec{u}\} \\ N \Downarrow_c & \text{ if } \exists N'. (N \longrightarrow N' \wedge N' \downarrow_c) \end{aligned}$$

We say $\mathcal{R}$ respects the observational predicate if $\Gamma; \Delta \vdash N_1 \mathcal{R} N_2$ with $c : \text{chan0}(U) \in \Delta$ implies $N_1 \downarrow_c$ iff $N_2 \downarrow_c$.

Now we define the observational congruence.

Definition 7.3. (Observational congruence) A typed congruence $\mathcal{R}$ is *sound* if it is reduction-closed and respects the observational predicate.

- We write $\cong$ for the maximum sound equality over a network invariant, i.e. $\cong$ is defined over a set which excludes the error states $\{N \mid \exists N_0. (N_0 \models \text{Init}, N_0 \longrightarrow N, N \not\models \text{Err})\}$.
- We write $\cong^\bullet$ for the maximum sound equality over untyped networks which include error states.

7.2 Transformation

We introduce a set of tractable conversion rules which can quickly check the equivalence of distributed networks. First we formally define the noninterference property.

Definition 7.4. Let us assume $\mapsto$ is a typed relation closed under name restriction, parallel composition and the structure rules. We say $\mapsto$ satisfies a *noninterference property*, i.e. if $N \longrightarrow N_1$ and $N \mapsto^* N_2$, then $N_1 \equiv N_2$ or there exists N' such that $N_1 \mapsto^* N'$ and $N_2 \longrightarrow N'$.

LEMMA 7.1. *Suppose $\mapsto$ satisfies a noninterference property and $\mapsto$ respects the observational predicate. Then $N_1 \mapsto N_2 \not\models \text{Err}$ implies $N_1 \cong N_2$.*

PROOF. By taking $\mathcal{R} = \{(N, M) \mid N \mapsto^* M \not\models \text{Err}\}$ and showing it is sound. By the assumption, we know $\mapsto$ respects the action predicate and so does $\mapsto^*$. Then it remains to show that $\mathcal{R}$ is reduction closed. Suppose $N \mathcal{R} M$ and $N \longrightarrow N'$ then by

the definition of noninterference, $N' \equiv M$ or $M \rightarrow M'$ such that $N' \mapsto^* M$. Since $\equiv \subset \mapsto^* \subset \mathcal{R}$, this completes the case. Now suppose $N\mathcal{R}M$ and $M \rightarrow M'$ then by Definition 7.3 we have that $N\mathcal{R}M \rightarrow M'$ and so $N \rightarrow N' \equiv M'$. $\square$

Code that can move safely. The transformation rules should reduce the number of communications and class downloadings, while preserving meaning. For this, we need to identify what kinds of code and programs can safely move from one location to another. Below predicate $\text{Mobile}_\Gamma(e)$ is true if $\text{fv}(e) = \emptyset$ and $o \in \text{fn}(e)$ implies $\Gamma \vdash o : C$ with $\text{remote}(C)$; i.e. e does not contain any free variables or local o-ids under environment Γ ; in addition it does not contain any of the following terms as a subterm (since they break the locality invariants, see § 5.3.1 and § 6.1.2).

$\{o.f, o.f = v, \text{sync } (o) \{e'\}, \text{insync } o \{e'\}, o.\text{notify}, o.\text{notifyAll}, o.\text{wait}, \text{ready } o \ n \}$

If $\text{Mobile}_\Gamma(e)$, e can *move* from one location to another preserving its meaning.

Transformation Rules. We define the key translation rules below, assuming that the right hand side is typed under $\Gamma; \Delta$. We omit surrounding context where it is unnecessary. Assume the right hand side is typed under $\Gamma; \Delta$. We omit the surrounding context where it is unnecessary.

Linearity

(l1) $\text{return}(c) \ E[\text{sandbox } \{e_1; \dots; e_n\}] \mapsto e_1; \dots; \text{return}(c) \ E[e_n]$

(l2) $E[\text{await } c] \mid e[\text{return}(c)/\text{return}] \mapsto E[\text{sandbox } \{e[e'/\text{return } e']\}]$

(l1) is standard. (l2) means that a method body e can be evaluated inline. This is ensured by linearity of channel c .

Class

(cm) $l[P, \sigma, \text{CT}] \mapsto l[P, \sigma, \text{CT} \cup \text{CT}'] \quad \text{ctcomp}(\text{CT}'), \vdash \text{CT}' : \text{ok}$

(cm) says that a complete class table can always move.

Closed

(cr) $(\nu x)(E[x] \mid P, \sigma \cdot [x \mapsto v]) \mapsto (E[v], \sigma) \quad \text{when } \{x := e\} \notin P \cup E[\]$

(fr) $(\nu \vec{u})(E[\text{freeze}[t](T \ x)\{e\}] \mid P, \sigma, \text{CT}) \mapsto (\nu \vec{u})(E[\lambda(T \ x).(\nu \vec{u})(l, e, \sigma', \text{CT}')] \mid P, \sigma, \text{CT})$

where in (fr), $\text{dom}(\sigma') \cap (\text{fnv}(P) \cup \text{fnv}(E)) = \emptyset$ and $\text{fnv}(\sigma') \subseteq \text{dom}(\sigma')$ and σ' and CT' are given following FREEZE. (cr) says that the timing of reading a value is unimportant if x only does not write in its scope. (fr) means that the timing of freezing an expression is ignored, provided it shares no information with other parties. Note “ νx ” in (cr) and “ $\nu \vec{u}$ ” in (fr) ensure x and u_i are not shared.

Method Invocation

(mi) $l[E[o.m(v)], \sigma, \text{CT}] \mid m[Q, \sigma', \text{CT}']$

$\mapsto l[E[\text{defrost}(v; \lambda(T \ x).(m, e[o/\text{this}], \emptyset, \emptyset))], \sigma, \text{CT}] \mid m[Q, \sigma', \text{CT}']$

where $\text{Mobile}_\Gamma(e[v, o/x, \text{this}])$, $\text{Mobile}_\Gamma(v)$, $[o \mapsto (C, \dots)] \in \sigma'$, $\text{mtype}(m, C) = T \rightarrow U$, and $\text{mbody}(m, C, \text{CT}') = (x, e)$. This rule means we can fetch a closure of the mobile method body from the remote site safely.

Mobile Method Bodies

(rm) $l[E[\text{await } c] \mid P, \sigma_l] \mid m[(\nu \vec{x})(R \mid Q, \sigma_m \cdot [\vec{x} \mapsto \vec{v}])]$

$\mapsto l[(\nu \vec{x})(E[\text{await } c] \mid R \mid P, \sigma_l \cdot [\vec{x} \mapsto \vec{v}], \text{CT}_l]) \mid m[Q, \sigma_m]$

where $\text{Mobile}_\Gamma(e[\vec{v}/\vec{x}])$, $\text{Mobile}_\Gamma(v_i)$, $R \equiv e[\text{return}(c)/\text{return}]$ and $\vec{x} \notin \text{fv}(Q)$. This

rule says a mobile thread with a store can move from the remote site safely.

Synchronisation

Suppose P and E does not include **notify**, **notifyAll**, **wait** or **ready**. Then:

$$\begin{aligned} (\text{sync}) \quad & l[P \mid E[\text{insync } o \{e\}] \mid P, \sigma_2, \text{CT}] \mapsto (\nu \bar{u})l[E[\text{insync } o \{e'\}] \mid P, \sigma'_2, \text{CT}'] \\ & \text{if } l[E[e] \mid P, \sigma_1, \text{CT}] \longrightarrow (\nu \bar{u})l[E[e'] \mid P, \sigma'_1, \text{CT}'], \\ & \sigma_2 = \text{setLock}(\sigma_1, o, \text{getLock}(\sigma, o) + 1), \sigma'_2 = \text{setLock}(\sigma'_1, o, \text{getLock}(\sigma, o) + 1) \end{aligned}$$

The reduction under synchronisation is deterministic if there is no **wait**, **notify** and **notifyall** in the program.

Deterministic Rule

$$(\text{ni}) \quad N \longrightarrow N' \implies N \mapsto N'$$

if the last reduction rule applied was not generated by variable read (VAR), variable assignment (ASS), field access (FLD), field assignment (FLDASS), freezing (FREEZE), the signal to a waiting thread (NOTIFY), monitor entry (SYNC) or the reduction for reacquiring a monitor (READY). The transformation rule $N \mapsto N'$ is defined as a binary relation generated by the above rules closed under parallel composition, restriction and structure rules (as in Figure 11).

THEOREM 7.1.

- (1) **(noninterference)** $\mapsto$ satisfies a noninterference property and respects the observational predicate under a network invariant.
- (2) **(type preservation)** Assume $\Gamma; \Delta \vdash N : \text{net}$ and $N \not\equiv \text{Err}$. Then $N \mapsto N'$ implies $\Gamma; \Delta \vdash N' : \text{net}$.
- (3) **(semantic preservation)** $N \mapsto N'$ implies $N \cong N'$.

PROOF. (1) is mechanical by investigating each of the above rules in turn. The most interesting rule is (mi), which is derived by (ni, rm, l2). (sync) is by Proposition 6.2 (1). (2) is straightforward by noting that $\longrightarrow$ is a typed relation, and that transformed terms can be related by it. (3) uses (1) and (2) together with Lemma 7.1. $\square$

Note that by Definition 6.1, Theorem 7.1 excludes the error statement. This is because the transformation is *not* sound *if an error occurs during execution*, as we shall discuss in the next subsection. More formally, $N \mapsto N'$ does not always imply $N \cong^\bullet N'$.

PROPOSITION 7.1.

- (1) $\text{freeze}[t](T \ x)\{e\} \cong \text{freeze}[t'](T \ x)\{e\}$.
- (2) There is a fully abstract embedding $\llbracket N \rrbracket$ of networks N that contain methods $\mathbf{m}(\vec{e})$ and frozen expressions $\text{freeze}[t](\vec{T} \ \vec{x})\{e\}$ with multiple parameters into networks with methods and frozen expressions with only single parameters.

PROOF. (1) Use Lemma 7.1 and (cm). (2) A translation of **freeze** is standard by currying. We encode methods with multiple parameters into those with just a single parameter in the most intuitive manner. Each method, instead of taking a vector $\vec{T} \ \vec{x}$ of parameters, takes a single parameter of a newly created class C . C contains fields $T_1 \ f_1; \dots; T_n \ f_n$; where field f_i corresponds to the i th parameter

of the original method definition. Then, all call sites for a particular method are replaced with a constructor call to an instance of the correct “parameter class”, so $o.m(\vec{v})$ becomes $o.m(\text{new } C(\vec{v}))$ for some C . We then prove that $N \cong \llbracket N \rrbracket$. See Appendix G for the encoding and proofs. $\square$

Call-backs. Before proving the main theorem, we formalise the notion of call-backs between two locations.

Definition 7.5. Let us write $c \rightsquigarrow c'$ if there exists a chain of channels such that

$$\text{return}(c) \ E_1[\text{await } c_1] \mid \text{return}(c_1) \ E_2[\text{await } c_2] \cdots \mid \text{return}(c_n) \ E_n[\text{await } c']$$

Suppose

$$N \rightarrow (\nu \vec{u})(l_1[\text{go } o.m(\vec{v}) \text{ with } c \mid E[\text{await } c] \mid P, \sigma_1, \text{CT}_1] \mid l_2[Q, \sigma_2, \text{CT}_2] \mid M) \stackrel{\text{def}}{=} N'$$

with $o \in \text{dom}(\sigma_2)$. We say there exists a *call-back* from l_1 to l_2 if we have:

$$N' \rightarrow (\nu \vec{u})(l_1[E[\text{await } c] \mid P', \sigma'_1, \text{CT}'_1] \mid l_2[\text{go } o'.m'(\vec{v}') \text{ with } c' \mid Q', \sigma'_2, \text{CT}'_2] \mid M')$$

where $o' \in \text{dom}(\sigma'_1)$ and there exists $c \rightsquigarrow c'$ in Q' .

The definition means: suppose “ $E[\text{await } c]$ ” is created in l_1 by the remote method invocation “ $o.m$ ” to l_2 . Then the computation of $E[\text{await } c]$ is blocked until some method invocation “ $o'.m'$ ” returned from l_2 . The equation between the third RMI program (RMI3) in Listing 5 in § 2 and the third optimal program (Opt3) in Listing 6 holds if there is no call-back as explained in § 2. Our framework can also justify the incorrectness of the optimisation between (RMI3) and (Op3) in the presence of call-back. Note that there is no guarantee that the caller and the call-back can be synchronised correctly in a naïve program, as they cannot hold the same lock (see § 5.3.1). For this reason, since most RMI programs do not use call-backs, we do not investigate them.

7.3 Correctness of the optimisations

We now prove the correctness of the optimised programs in § 2. We transform one program to another using the transformation rules defined above.

We first demonstrate how to transform the optimised program 1 (Opt1) in Listing 1 to the original program 1 (RMI1) in Listing 2. Let us assume e is a program from line 2 to 4 in (RMI1). We omit the surrounding context as there is no class loading in this example. After the method invocation by $o.m\text{Opt1}(r, n)$ with c ,

$$(\nu a)(\text{thunk}(\text{int}) \ t = \text{freeze}[\text{lazy}]\{e; z\}; \text{return}(c) \ r.\text{run}(t), [a \mapsto n])$$

Let $v = \lambda(\text{unit } x).(\nu a)(l, e; z, [a \mapsto n], \emptyset)$. Then the above configuration is transformed to:

$$\begin{aligned} &\mapsto (\nu a)(\text{thunk}(\text{int}) \ t = v; \text{return}(c) \ r.\text{run}(t), [a \mapsto n]) && (\text{fr}) \\ &\mapsto^+ (\nu t)(\text{return}(c) \ r.\text{run}(t), [t \mapsto v]) && (\text{ni}) \\ &\mapsto \text{return}(c) \ r.\text{run}(v), \emptyset && (\text{cr}) \\ &\mapsto \text{return}(c) \ \text{defrost}(v; \lambda(T \ x).(l, \text{defrost}(x), \emptyset, \emptyset)), \emptyset && (\text{mi}) \\ &\mapsto^+ (\nu a)(\text{return}(c) \ \text{sandbox } \{e; z\}, [a \mapsto n]) && (\text{ni}) \\ &\mapsto (\nu a)(e; \text{return}(c) \ z, [a \mapsto n]) && (\text{l1}) \end{aligned}$$

The last line is identical to (RMI1) after the method invocation by $o.m1(r, n)$ with c . Note that **defrost** and **sandbox** do not affect other parties, so that the reduction

(ni) satisfies a noninterference property, hence this reduction preserves the semantics. Because we have $\text{Mobile}_0(v)$, we can apply (mi) in the forth line. Hence (Opt1) is transformed to (RMI1).

The correctness of (Opt2) in Listing 4 is also straightforward by repeating the same routine twice.

We show (RMI3) in Listing 5 is equivalent with (Opt3) in Listing 6 under the assumption there is no call-back. Then the body of (Opt3) is equivalent to `return r.run(freeze[eager](T x){e[ $\vec{e}$ ]/ $\vec{b}$ ];z})` and $e'_i = \text{deserialize}(v_i)$ where $v_i = \lambda(\text{unit } x).(\nu \vec{u})(l, a, \sigma_i)$ is a serialised value at line i ($3 \leq i \leq 5$) in (Opt3). Then we apply a similar transformation with the above to derive (RMI3). See Appendix H for the detailed proofs.

Note that our freezing preserves sharing between objects (Point 1 in (Opt3) in § 2), hence we can prove the following equation:

$$x.f = y; r.h(x, y) \cong x.f = y; r.\text{run}(\text{freeze}\{r.h(x, y)\}).$$

Finally by Proposition 7.1 (1), we can derive (Opt4) from (Opt3), hence (Opt4) is equivalent to (RMI3). Not all equations are valid if a network error occurs during executions. For example, `eager` and `lazy` are not equal in the presence of `ERR-CLASSNOTFOUND`, hence Proposition 7.1 (1) is not applicable. To summarise, we have:

THEOREM 7.2. (Correctness of the Optimisations)

- (1) (RMI1) and (Opt1) are equivalent up to $\cong$.
- (2) (RMI2) and (Opt2) are equivalent up to $\cong$.
- (3) (RMI3) and (Opt3) are equivalent up to $\cong$ without call-back.
- (4) (Opt3) and (Opt4) are equivalent up to $\cong$, hence (RMI3) and (Opt4) are equivalent up to $\cong$ without call-back.
- (5) None of them are equivalent up to $\cong^\bullet$.

8. RELATED WORK

Class loading and downloading. Class loading and downloading are crucial to many useful Java RMI applications, offering a convenient mechanism for distributing code to remote consumers. The class verification and maintenance of type safety during linking are studied in [Liang and Bracha 1998; Qian et al. 2000]. Our formulation of class downloading is modular, so it is adaptable to model other linking strategies [Drossopoulou and Eisenbach 2002; Drossopoulou et al. 2003], see § 4.2. We set the class invariant `Inv(3)` in Definition 6.2. This is because the Java serialisation API imposes the strict default class version control discussed in § 6.1.1. Another solution is to explicitly model the Java exception `InvalidClassException` to check for mismatch between downloaded and existing classes. This dynamic approach leads to the same invariant to prove the subject reduction theorem.

Most of the literature surrounding class loading in practice takes the lazy approach. As we discussed earlier, in the setting of remote method invocation laziness can be expensive due to delay involved in retrieving a large class hierarchy over the network. Krintz et al [Krintz et al. 1999] propose a class splitting and pre-fetching algorithm to reduce this. Their specific example is applet loading: if

the time spent in an interactive portion of an applet is used to download classes that may be needed in future, we can fetch them ahead of time so that the user does not encounter a large delay, sharing the motivation for our (eager) code mobility primitive. The partly eager class loading in their approach is implicit, but requires control flow information about the program in question to determine where to insert instructions to trigger ahead-of-time fetching. This framework may be difficult to apply in a general distributed setting, since clients may not have access to the code of a remote server. Also their approach merely mitigates the effect of network delay rather than removing it; it still requires the sequential request of a hierarchy of superclasses. We believe an explicit thunk primitive as we proposed in the present work may offer an effective alternative in such situations.

Distributed objects. Obliq [Cardelli 1994] is a distributed object-based, lexically scoped language proposed by Cardelli. One key feature of the language is that methods are stored within objects—there is no hierarchy of tables to inspect as in most class-based languages. Merro et al [Merro et al. 2002] encode a core part of Obliq into the untyped π -calculus. They use their encoding to show a flaw in part of the original migration semantics and propose a repair. Later Nestmann et al [Nestmann et al. 2002] formalised a typing system for a core Obliq calculus and studied different kinds of object aliasing. Briaïs and Nestmann [Briaïs and Nestmann 2002] then strengthened the safety result in [Merro et al. 2002] by directly developing the must equivalence at the language level (without using the translation into the π -calculus). They also apply a noninterference property to show the two terms (with and without surrogation) are must-equivalent. DJ models two important concerns in distributed class-based object-oriented languages missing from Obliq, that is object serialisation and dynamic class downloading associated with inheritance in Java (note that the same term “serialisation” used in [Cardelli 1994] refers to one in the sense of transaction theory). These features require a consistent formulation of dynamic deep copying of object/class graphs. As we have seen in § 7, detailed analysis of these features is required to justify the correctness of the optimisation examples in § 2. The proof method using syntactic transformations in § 7 is also new.

Emerald [Hutchinson et al. 1991] is another example of a distributed object-based language. It supports classes represented as objects, however there is no concept of class loading as in DJ—information about inheritance hierarchies is discarded at compile-time. Objects in Emerald may be *active* in that they are permitted their own internal thread of control that runs concurrently with method invocations on that object. Such objects may explicitly move themselves to other locations by making a library call. In DJ the fundamental unit of mobility is arbitrary higher-order expressions: this general code freezing primitive can represent object mobility similar to Emerald when it is combined with standard Java RMI. Finally, there has been no study of the formal semantics of Emerald.

Gordon and Hankin [Gordon and Hankin 1999] extend the object calculus [Abadi and Cardelli 1996] with explicit concurrency primitives from the π -calculus. Their focus is synchronisation primitives (such as fork and join) rather than distribution, so they only use a single location. Jeffrey [Jeffrey 2000] treats an extension of [Gordon and Hankin 1999] for the study of locality with static and dynamic type

checking. The concurrent object calculus is not class-based, hence neither work treats dynamic class loading or serialisation (though [Jeffrey 2000] treats transactional serialisation as in [Cardelli 1994]), which are among the key elements for analysis of RMI and code mobility in Java.

Scope and runtime formalisms for Java. Zhao et al [Zhao et al. 2004] propose a calculus with primitives for explicit memory management, called SJ, for a study of containment in real-time Java. The SJ calculus proposes a typing discipline based on the idea of *scoped types*—memory in real-time applications is allocated in a strict hierarchy of scopes. Using the existing Java package structure to divide such scopes, their typing system statically prevents some scope invariants being broken. Their focus is on real-time concurrency in a single location, while ours is on dynamic distribution of code in multiple locations. DJ also guarantees similar scoping properties by invariants, for example $\text{Inv}(6)$ in Definition 6.2 ensures that identifiers for local objects do not leak to other locations in the presence of synchronisation primitives.

The representation of object-oriented runtime in formal semantics is not limited to distributed programs, as found in study of execution models of the .NET CLR by Gordon and Syme [Gordon and Syme 2001] and Yu et al [Yu et al. 2004].

The JavaSeal [Vitek et al. 1998] project is an implementation of the Seal calculus for Java. It is realised as an API and run-time system inside the JVM, targeted as a programming framework for building multi-agent systems. The semantics of these APIs depend on distributed primitives in the implementation language, which are precisely the target of the formal analysis in the present paper. JavaSeal may offer a suggestion for the implementation and security treatment of higher-order code passing proposed in the present paper.

Functions with marshaling primitives. Ohori and Kato [Ohori and Kato 1993] extend a purely functional part of ML with two primitives for remote higher-order code evaluation via channels, and show that the type system of this language is sound with respect to a low-level calculus. The low-level calculus is equipped with runtime primitives such as closures of functions and creation of names. Their focus is pure polymorphic functions, hence they treat neither side-effects nor (distributed) object-oriented features. Acute [Acute 2005] is an extension of OCaml equipped with type-safe marshaling and distributed primitives. By using flags called marks, the user can control dynamic loading of a sequence of modules when marshaling his code. This facility is similar to our lazy and eager class loading. The language also provides more flexible way to rebind local resources and modules. An extension of our freeze operator for fine-grained rebinding is an interesting topic, though as we discussed in § 6.1.1, it is not suitable in practice due to the Java serialisation API.

Staged computation and meta-programming. Taha and Sheard [Taha and Sheard 1997] give a dialect of ML containing staging annotations to generate code at run-time, and to control evaluation order of programs. The authors give a formal semantics of their language, called MetaML, and prove that the code a well-typed program generates will itself be type-safe.

The `freeze` and `defrost` primitives in DJ can be thought of as staging annotations, and also guarantee that frozen expressions should be well-typed in any

context. However we study distribution and concurrency in an imperative setting, with strong emphasis on runtime features. These features are not discussed in MetaML as it is a functional language, nor the problems associated with classloading we address.

Kamin et al [Kamin et al. 2003] extend the syntax of Java with staging annotations and provide a compiler for a language called Jumbo. They allow creation of classes at runtime, focusing on single-location performance optimisation: there is no discussion of use in distributed applications, a main focal point of our work. They give no static guarantees about type safety of generated code, nor do they allow code to be generated in fragments smaller than an entire class. They do not consider higher-order quotation, permitting only one level of quotation and anti-quotation.

Zook et al [Zook et al. 2004] propose Meta-AspectJ as a meta-programming tool for an aspect-oriented language. They implement a compiler that takes code templates—containing quoted Aspect-J code—and turns them into aspect declarations that can be applied as normal to Java programs. Their system is more focused on compile-time code generation, and offers weaker static guarantees: well-typed generators do not guarantee type safety of the generated aspects.

9. CONCLUSIONS AND FURTHER WORK

This paper introduced a Java-like core language for RMI with higher-order code mobility. It models runtime for distributed computation including dynamic class downloading and object serialisation. Using the new primitives for code mobility, we subsumed the existing serialisation mechanism of Java and were able to precisely describe examples of communication-based optimisations for RMI programs on a formal foundation. We established type soundness and safety properties of the language using distributed invariants. Finally, by the behavioural theory developed in § 7, we were able to systematically prove the correctness of the examples in § 2.

Explicit code mobility as a language primitive gives powerful control over code distribution strategies in object-oriented distributed applications. This is demonstrated in the examples in § 2. In [Bogle and Liskov 1994; Yeung and Kelly 2003; Yeung 2004], these optimisations are informally described as implementation details. Not only is source-level presentation necessary for their semantic justification, but also explicit treatment of code mobility gives programmers fine-grained control over the evaluation order and location of executing code. It also opens the potential for source-level verification methodologies for access control, secrecy and other security concerns, as briefly discussed below. Note current customised class downloading mechanisms do not offer active code mobility and algorithmic control of code distribution (as in the last example of § 2).

Further, the fine-grained control of code mobility has a direct practical significance: the optimisation strategy in [Yeung and Kelly 2003; Yeung 2004] cannot aggregate code in which new object generation is inserted, such as:

```

1  int m3(RemoteObject r, MyObj a) {
2      int x = r.f(a);
3      int y = r.g(new MyObj(x));
4      int z = r.h(a, y);

```

```

5 |   return z;
6 | }

```

where `MyObj` is a local class in the client. This is because we need active class code delivery if this code is to be executed in a remote server. In contrast, the freeze primitive in our language can straightforwardly handle aggregation of this code. We also believe that, in comparison with direct, byte-code level implementation in [Yeung and Kelly 2003; Yeung 2004], the use of our high-level primitives may not jeopardise efficiency but rather can even enhance it by e.g. allowing more flexible inter-procedure optimisation.

The complexity of the third program optimisation poses the question of whether the original copying semantics of Java RMI are themselves correct in the first place: making a remote call can entail subtly different invocation semantics to calling a local method. Our code freezing primitive allows us to make the call semantics explicit, and also allows us to support more traditional ideas about object mobility [Hutchinson et al. 1991; Cardelli 1994], such as side-effects in calls at the server side.

The class-based language considered in the present work does not include such language features as casting [Igarashi et al. 2001; Bierman et al. 2003], exceptions [Ancona et al. 2002] and parametric polymorphism [Igarashi et al. 2001]; although these features can be represented by extension of the present syntax and types, their precise interplay with distributed language constructs requires examination.

An important future topic is enrichment of the invariants and type structures to strengthen safety properties (e.g. for security). Here we identify two orthogonal directions. The first concerns mobility. As can be seen in the second example in § 2, the current type structure of higher-order code (e.g. `thunk<int>`) tells the consumer little about the behaviour of the code he is about to execute, which can be dangerous [McGraw and Morrisett 2000; Bogle and Liskov 1994]. In Java, the `RMISecurityManager` can be used with an appropriate policy file to ensure that code downloaded from remote sites has restricted capability. By extending DJ with principals, we can examine the originator of a piece of code to determine suitable privileges prior to execution [Wallach et al. 1997]. To ensure the integrity of resources we can dynamically check invariants when code arrives (e.g. by adding constraints in DEFROST), or we could allow static checking by adding more fine-grained information about the accessibility of methods in class signatures, along the lines of [Yoshida 2004].

The second direction is to extend the syntax and operational semantics to allow complex, structured, communications. For this purpose we have been studying *session types* [Honda et al. 1998; Vasconcelos et al. 2004] for ensuring correct pattern matching of sequences of socket communications, incorporating a new class of channels at the user syntax level. Our operational semantics for RMI is smoothly extensible to model advanced communication protocols. Session types are designed using class signatures, and safety is proved together with the same invariance properties developed in this paper.

Study of the semantics of failure and recovery in our framework is an important topic. So far we have incorporated the possibility of failures in class downloading and remote invocation due to network partition (defined by **Err**-rules in § 4). When

a message is lost, some notion of time-out is generally used to determine whether to re-transmit or fail. Such error recovery can be investigated by defining different invocation semantics (for example at-most-once [Microsystems Inc. 2005]) and adding runtime extensions to DJ. This point is also relevant when we consider socket-based communication instead of RMI.

We have implemented an initial version of our new primitives for code mobility [Tejani 2005]. This takes the form of a source-to-source translator, compiling the `freeze` and `defrost` operations into standard Java source. Eager class loading via RMI requires modification to the class loading mechanism, which is achieved by installing a custom class loader working in conjunction with our translated source. This approach has the advantage that we can use an ordinary Java compiler and existing tools, and that the JVM would not need modification. However a more direct approach (for example extending the virtual machine) may yield better performance.

The examples in § 2 and the transformation rules in § 7 lead to the question of how to automatically translate from RMI source programs to programs exploiting code mobility for added efficiency. Developing a general theory and an integrated tool is non-trivial due to an interplay between inter node and procedure optimisations. Furthermore we need to formalise a cost theory for distributed communication with respect to the distance of the locations and the size of code and class tables transferred. DJ can be used as a reference model to define efficiency since it exposes distributed runtime explicitly by means of syntax and reduction rules. For example, we can add marshaling costs to the FREEZE rule with respect to the size of the frozen expression; we can investigate the cost of class downloading with respect to the size of a downloaded class table CT' and a distance between location l_1 and location l_2 , using rule DOWNLOAD. An interesting further topic is an application to DJ of the cost-preorder theory developed for process algebra [Arun-Kumar and Hennessy 1992] to compare program performance.

Acknowledgements. We deeply thank Paul Kelly for his discussions about the Veneer vJVM and the RMI optimisations in [Yeung and Kelly 2003; Yeung 2004]. Comments from the anonymous reviewers helped to revise the submission. We thank Luca Cardelli, Susan Eisenbach, Kohei Honda, Andrew Kennedy and members of the SLURP and ToCS groups at Imperial College London for their discussions, Farzana Tejani and Karen Osmond for their work on implementation. We thank Mariangiola Dezani, Sophia Drossopoulou and Uwe Nestmann for their comments on an earlier version of this paper. The first author is partially supported by an EPSRC PhD studentship and the second author is partially supported by EPSRC Advanced Fellowship (GR/T03208/01), EPSRC GR/S55538/01 and EPSRC GR/T04724/01.

REFERENCES

- ABADI, M. AND CARDELLI, L. 1996. *A Theory of Objects*. Springer-Verlag.
- ACUTE. 2005. Acute home page. <http://www.cl.cam.ac.uk/users/pes20/acute>.
- AHERN, A. AND YOSHIDA, N. 2005a. Formal Analysis of a Distributed Object-Oriented Language and Runtime. Tech. Rep. 2005/01, Department of Computing, Imperial College London: Available at: <http://www.doc.ic.ac.uk/~aja/dcbl.html>.
- AHERN, A. AND YOSHIDA, N. 2005b. Formalising java rmi with explicit code mobility. In *OOP-*

- SLA '05 (to appear), the 20th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages and Applications*. ACM Press.
- ANCONA, D., LAGORIO, G., AND ZUCCA, E. 2002. Simplifying types in a calculus for Java exceptions. Tech. rep., DISI - Università di Genova.
- ARUN-KUMAR, S. AND HENNESSY, M. 1992. An efficiency preorder for processes. *Acta Inf.* 29, 8, 737–760.
- BIERMAN, G., PARKINSON, M., AND PITTS, A. 2003. MJ: An imperative core calculus for Java and Java with effects. Tech. Rep. 563, University of Cambridge Computer Laboratory. April.
- BOGLE, P. AND LISKOV, B. 1994. Reducing cross domain call overhead using batched futures. In *OOPSLA '94*. ACM Press, 341–354.
- BRIAIS, S. AND NESTMANN, U. 2002. Mobile objects “must” move safely. In *FMOODS 2002*. Kluwer Academic Publishers, 129–146.
- CARDELLI, L. 1994. Obliq: A language with distributed scope. Tech. Rep. 122, Systems Research Center, Digital Equipment Corporation.
- CHRIST, R. 2000. San Francisco Performance: A case study in performance of large-scale Java applications. *IBM Systems Journal* 39, 1.
- DROSSOPOULOU, S. AND EISENBACH, S. 2002. Manifestations of Dynamic Linking. In *The First Workshop on Unanticipated Software Evolution (USE 2002)*. <http://joint.org/use2002/proceedings.html>, Málaga, Spain.
- DROSSOPOULOU, S., LAGORIO, G., AND EISENBACH, S. 2003. Flexible Models for Dynamic Linking. In *12th European Symposium on Programming*, P. Degano, Ed. Lecture Notes in Computer Science, vol. 2618. Springer-Verlag, 38–53.
- FLANAGAN, D. 2000. *Java Examples in a Nutshell*. O'Reilly UK.
- GORDON, A. D. AND HANKIN, P. D. 1999. A concurrent object calculus: Reduction and typing. Tech. Rep. 457, University of Cambridge Computer Laboratory. February.
- GORDON, A. D. AND SYME, D. 2001. Typing a multi-language intermediate code. In *Proceedings of the 28th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. ACM Press, 248–260.
- GREANIER, T. 2005. Discover the secrets of the Java Serialization API. <http://java.sun.com/developer/technicalArticles/Programming/serialization/>.
- HONDA, K. 1996. Composing processes. In *Proceedings of POPL'96*. 344–357.
- HONDA, K. AND TOKORO, M. 1991. An object calculus for asynchronous communication. In *Proceedings of ECOOP'91*.
- HONDA, K., VASCONCELOS, V. T., AND KUBO, M. 1998. Language primitives and type disciplines for structured communication-based programming. In *ESOP'98*. Lecture Notes in Computer Science, vol. 1381. 22–138.
- HONDA, K. AND YOSHIDA, N. 1995. On reduction-based process semantics. *TCS* 151, 2, 385–435.
- HUTCHINSON, N. C., RAJ, R. K., BLACK, A. P., LEVY, H. M., AND JUL, E. 1991. The Emerald Programming Language. Tech. rep., Department of Computer Science, University of British Columbia, Vancouver BC, Canada V6T 1Z2. October.
- IGARASHI, A., PIERCE, B. C., AND WADLER, P. 2001. Featherweight Java: a minimal core calculus for Java and GJ. *ACM Transactions on Programming Languages and Systems* 23, 3, 396–450.
- JEFFREY, A. 2000. A Distributed Object Calculus. In *FOOL*. ACM Press.
- JONES, C. 1993. Constraining interference in an object-based design method. In *Proceedings of TAPSOFT'93*. LNCS 668. Springer Verlag, 136–150.
- KAMIN, S., CLAUSEN, L., AND JARVIS, A. 2003. Jumbo:run-time code generation for java and its applications. In *CGO03*. IEEE.
- KOBAYASHI, N., PIERCE, B., AND TURNER, D. 1996. Linear types and π -calculus. In *Proceedings of POPL'96*. 358–371.
- KRINTZ, C., CALDER, B., AND HÖLZLE, U. 1999. Reducing transfer delay using Java class file splitting and prefetching. In *Proceedings of the 14th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*. ACM Press, 276–291.

- LIANG, S. AND BRACHA, G. 1998. Dynamic class loading in the Java virtual machine. In *Proceedings of the 13th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*. ACM Press, 36–44.
- MCGRAW, G. AND MORRISETT, G. 2000. Attacking malicious code: a report to the infosec research council. *IEEE Software* 17, 5, 33–44.
- MERRO, M., KLEIST, J., AND NESTMANN, U. 2002. Mobile objects as mobile processes. *Information and Computation* 177, 2, 195–241.
- MICROSYSTEMS INC., S. 2005. Java Remote Method Invocation (RMI). <http://java.sun.com/products/jdk/rmi/>.
- MILNER, R., PARROW, J., AND WALKER, D. 1992. A calculus of mobile processes, parts I and II. *Info. & Comp.* 100, 1.
- NESTMANN, U., HÜTTEL, H., KLEIST, J., AND MERRO, M. 2002. Aliasing models for mobile objects. *Inf. Comput.* 177, 2, 195–241.
- OHORI, A. AND KATO, K. 1993. Semantics for communication primitives in a polymorphic language. In *POPL '93: Proceedings of the 20th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. ACM Press, 99–112.
- PIERCE, B. C. 2002. *Types and Programming Languages*. MIT Press.
- QIAN, Z., GOLDBERG, A., AND COGLIO, A. 2000. A formal specification of Java class loading. In *Proceedings of the 15th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*. ACM Press, 325–336.
- REYNOLDS, J. C. 1978. Syntactic control of interference. In *Proceedings of the 5th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*. ACM Press, 39–46.
- SANGIORGI, D. 1992. Expressing mobility in process algebras: First-order and higher order paradigms. Ph.D. thesis, University of Edinburgh.
- STAMOS, J. W. AND GIFFORD, D. K. 1990. Implementing remote evaluation. *IEEE Trans. Softw. Eng.* 16, 7, 710–722.
- TAHA, W. AND SHEARD, T. 1997. Multi-stage programming with explicit annotations. In *PEPM '97: Proceedings of the 1997 ACM SIGPLAN symposium on Partial evaluation and semantics-based program manipulation*. ACM Press, 203–217.
- TEJANI, F. MEng. Computing Final Year Project, Imperial College London, to appear in June, 2005. Implementation of a distributed mobile java. M.S. thesis, Imperial College London.
- VASCONCELOS, V. T., RAVARA, A., AND GAY, S. 2004. Session types for functional multithreading. In *CONCUR'04. Lecture Notes in Computer Science*, vol. 3170. 497–511.
- VITEK, J., BRYCE, C., AND BINDER, W. 1998. Designing JavaSeal or how to make Java safe for agents. *Electronic Commerce Objects*.
- WALLACH, D. S., BALFANZ, D., DEAN, D., AND FELTEN, E. W. 1997. Extensible security architectures for Java. In *Proceedings of the sixteenth ACM symposium on Operating systems principles*. ACM Press, 116–128.
- YEUNG, K. 2004. Dynamic performance optimisation of distributed java applications. Ph.D. thesis, Imperial College London.
- YEUNG, K. AND KELLY, P. 2003. Optimizing Java RMI programs by communication restructuring. In *Middleware'03. Lecture Notes in Computer Science*, vol. 2672. 324–343.
- YOSHIDA, N. 2004. Channel dependency types for higher-order mobile processes. In *POPL '04, Conference Record of the 31st Annual Symposium on Principles of Programming Languages*. ACM Press, 147–160. Full version available at www.doc.ic.ac.uk/~yoshida.
- YOSHIDA, N., BERGER, M., AND HONDA, K. 2001. Strong Normalisation in the π -Calculus. In *Proc. LICS'01. IEEE*, 311–322. The full version in *Journal of Inf. & Comp.*, 191 (2004) 145–202, Elsevier.
- YU, D., KENNEDY, A., AND SYME, D. 2004. Formalization of generics for the .net common language runtime. In *POPL '04: Proceedings of the 31st ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. ACM Press, 39–51.
- ZHAO, T., NOBLE, J., AND VITEK, J. 2004. Scoped types for real-time Java. In *Proceedings of the 25th Annual IEEE Symposium on Real-time Systems*.

ZOOK, D., HUANG, S. S., AND SMARAGDAKIS, Y. 2004. Generating AspectJ Programs with Meta-AspectJ. In *Proceedings of the Third International Conference on Generative Programming and Component Engineering*. Lecture Notes in Computer Science, vol. 3286. Springer-Verlag GmbH, 1–19.

A. AUXILIARY DEFINITIONS

This appendix contains the full definitions of some of the functions used in the main sections.

A.1 Domains

The function **dom** returns the domain of a mapping. It is defined over stores, class tables, class signatures and configurations, and is given as follows.

$$\begin{aligned}
 \text{dom}(\emptyset) &= \emptyset, \quad \text{dom}(\sigma \cdot [x \mapsto \dots]) = \text{dom}(\sigma) \cup \{x\}, \quad \text{dom}(\sigma \cdot [o \mapsto \dots]) = \text{dom}(\sigma) \cup \{o\} \\
 \text{dom}(\text{CT} \cdot [C \mapsto \dots]) &= \text{dom}(\text{CT}) \cup \{C\} \\
 \text{dom}(\text{CSig} \cdot [C \mapsto \dots]) &= \text{dom}(\text{CSig}) \cup \{C\} \\
 \text{dom}((\nu \vec{u})(P, \sigma, \text{CT})) &= \text{dom}(\sigma) \setminus \{\vec{u}\} \\
 \text{dom}(\mathbf{0}) &= \emptyset, \quad \text{dom}(l[F]) = \text{dom}(F), \quad \text{dom}(N_1 \mid N_2) = \text{dom}(N_1) \cup \text{dom}(N_2) \\
 \text{dom}((\nu u)N) &= \text{dom}(N) \setminus \{u\}
 \end{aligned}$$

A.2 Free variables and names

The functions for determining free variables **fv** and free names **fn** are defined as follows. For classes and methods:

$$\begin{array}{lll}
 \text{class } C \text{ extends } D \{ \vec{T} \vec{f}; K \vec{M} \} & \text{fv} & \text{fn} \\
 U \ m(C \ x)\{e\} & \bigcup \text{fv}(M_i) & \bigcup \text{fn}(M_i) \\
 & \text{fv}(e) \setminus \{x\} & \text{fn}(e)
 \end{array}$$

For values:

$$\begin{array}{lll}
 \text{true, false, } (), \text{null, } \epsilon & \text{fv} & \text{fn} \\
 o & \emptyset & \emptyset \\
 \lambda(T \ x).(\nu \vec{u})(l, e, \sigma, \text{CT}) & \emptyset & \{o\} \\
 \lambda(T \ x).(\nu \vec{u})(l, e, \sigma, \text{CT}) & ((\text{fv}(e) \setminus \{x\}) \cup \text{fv}(\sigma) \cup \text{fv}(\text{CT})) \setminus \{\vec{u}\} & (\text{fn}(e) \cup \text{fn}(\sigma) \cup \text{fn}(\text{CT})) \setminus \{\vec{u}\}
 \end{array}$$

For expressions we omit the cases where the free variables (resp. names) of a term are merely the union of the free variables of its subterms.

$$\begin{array}{lll}
 x & \text{fv} & \text{fn} \\
 \text{this} & \{x\} & \emptyset \\
 x = e & \emptyset & \emptyset \\
 T \ x = e_0; e_1 & \{x\} \cup \text{fv}(e) & \text{fn}(e) \\
 \text{return} & \text{fv}(e_0) \cup (\text{fv}(e_1) \setminus \{x\}) & \bigcup \text{fn}(e_i) \\
 \text{freeze}[t](T \ x)\{e\} & \emptyset & \emptyset \\
 \text{await } c & \text{fv}(e) \setminus \{x\} & \text{fn}(e) \\
 \text{insync } o \ \{e\} & \emptyset & \{c\} \\
 \text{ready } o \ n & \emptyset & \{o\} \cup \text{fn}(e) \\
 \text{waiting}(c) \ n & \emptyset & \{o\} \\
 \text{Error} & \emptyset & \{c\} \\
 & & \emptyset
 \end{array}$$

(Configurations)		
$(\nu u)P, \sigma, \text{CT} \equiv (\nu u)(P, \sigma, \text{CT})$	$u \notin \text{fn}(\sigma) \cup \text{fn}(\text{CT})$	
$(\nu u)(\nu u')F \equiv (\nu u')(\nu u)F$		
$(\nu x)(P, \sigma \cdot [x \mapsto v], \text{CT}) \equiv P, \sigma, \text{CT}$	$x \notin \text{fv}(P)$	
$(\nu o)(P, \sigma \cdot [o \mapsto (C, \dots)], \text{CT}) \equiv P, \sigma, \text{CT}$	$o \notin \text{fn}(P) \cup \text{fn}(\sigma)$	
(Threads)		(Networks)
$P \mid \mathbf{0} \equiv P$		$N \mid \mathbf{0} \equiv N$
$P \mid P_0 \equiv P_0 \mid P$		$N \mid N_0 \equiv N_0 \mid N$
$P \mid (P_0 \mid P_1) \equiv (P \mid P_0) \mid P_1$		$N \mid (N_0 \mid N_1) \equiv (N \mid N_0) \mid N_1$
$(\nu u)(P \mid P_0) \equiv (\nu u)P \mid P_0 \quad u \notin \text{fn}(P_0)$		$(\nu u)(N \mid N_0) \equiv (\nu u)N \mid N_0 \quad u \notin \text{fnv}(N_0)$
$(\nu c)\mathbf{0} \equiv \mathbf{0}$		$(\nu c)\mathbf{0} \equiv \mathbf{0}$
$(\nu u)(\nu u')P \equiv (\nu u')(\nu u)P$		$(\nu u)(\nu u')N \equiv (\nu u')(\nu u)N$
$\text{return}(d) \epsilon \equiv \text{return}(d)$		$l[(\nu u)(F)] \equiv (\nu u)l[F]$
$\epsilon; e \equiv e$		
$\text{return } \epsilon \equiv \text{return}$		

Fig. 16: Structural equivalence

For threads:

$\mathbf{0}$	$\emptyset$	$\emptyset$
$P_1 \mid P_2$	$\bigcup \text{fv}(P_i)$	$\bigcup \text{fn}(P_i)$
$(\nu u)P$	$\text{fv}(P) \setminus \{u\}$	$\text{fn}(P) \setminus \{u\}$
$\text{forked } e$	$\text{fv}(e)$	$\text{fn}(e)$
$[\text{go}] e \text{ with/to } c$	$\text{fv}(e)$	$\{c\} \cup \text{fn}(e)$
$\text{return}(c) e$	$\text{fv}(e)$	$\{c\} \cup \text{fn}(e)$
$(\nu \vec{u})(P, \sigma, \text{CT})$	$(\text{fv}(P) \cup \text{fv}(\sigma) \cup \text{fv}(\text{CT})) \setminus \{\vec{u}\}$	$(\text{fn}(P) \cup \text{fn}(\sigma) \cup \text{fn}(\text{CT})) \setminus \{\vec{u}\}$

For networks:

$\mathbf{0}$	$\emptyset$	$\emptyset$
$l[F]$	$\text{fv}(F)$	$\text{fn}(F)$
$N_1 \mid N_2$	$\bigcup \text{fv}(N_i)$	$\bigcup \text{fn}(N_i)$
$(\nu u)N$	$\text{fv}(N) \setminus \{u\}$	$\text{fn}(N) \setminus \{u\}$
$\emptyset$	$\emptyset$	$\emptyset$
$\sigma \cdot [x \mapsto v]$	$\{x\} \cup \text{fv}(v) \cup \text{fv}(\sigma)$	$\text{fn}(v) \cup \text{fn}(\sigma)$
$\sigma \cdot [o \mapsto (C, \vec{f} : \vec{v}, n, \{\vec{c}\})]$	$\text{fv}(\vec{v}) \cup \text{fv}(\sigma)$	$\{o, \vec{c}\} \cup \text{fn}(\vec{v}) \cup \text{fn}(\sigma)$
$\emptyset$	$\emptyset$	$\emptyset$
$\text{CT} \cdot [C \mapsto L]$	$\text{fv}(L) \cup \text{fv}(\text{CT})$	$\text{fn}(L) \cup \text{fn}(\text{CT})$

B. STRUCTURAL EQUIVALENCE

This section defines the structural equivalences for DJ. They are defined for threads, networks and configurations in Figure 16. Formally, $\equiv$ is an equivalence relation which includes α -conversion and is generated by the equations in Figure 16.

The last two rules for configurations define garbage collection of useless store entries, while the last three rules for threads are used to erase runtime value ϵ of the void type. Others rules, including scope opening, are inherited from those of the π -calculus [Milner et al. 1992], and so are standard.

C. PROOFS FOR THE CORRECTNESS OF THE ALGORITHMS

This section lists the proofs for Lemma 4.1. We show (1, 2) by induction on the length of σ , and (3, 4) by induction on the length of CT .

In the subsequent proofs, we use *length functions* for stores and class tables defined by:

$$\begin{aligned} \text{length}(\emptyset) &= 0 & \text{length}(\sigma \cdot [o \mapsto (C, \dots)]) &= \text{length}(\sigma \cdot [x \mapsto v]) = \text{length}(\sigma) + 1 \\ \text{length}(\emptyset) &= 0 & \text{length}(\text{CT} \cdot [C \mapsto \text{class } C \text{ extends } D \{ \vec{T} \vec{f}; K \vec{M} \}]) &= \text{length}(\text{CT}) + 1 \end{aligned}$$

(1) By induction on $\text{length}(\sigma)$. For the base case assume $\text{length}(\sigma) = 0$. Then $\sigma = \emptyset$ and $\Gamma; \emptyset \vdash \emptyset : \text{ok}$ immediately. The two conjunctions hold vacuously, therefore this case is complete.

For the inductive case, we assume that the property holds for $\text{length}(\sigma) < n$ and prove for $\text{length}(\sigma) = n$. Examining the algorithm, there are two sub-cases. If $v \notin \text{dom}(\sigma)$ or $v \in \text{dom}(\sigma)$ and v is an instance of a class C such that $\text{remote}(C)$, then $\sigma' = \emptyset$. Then, as in the base case, the property holds straightforwardly.

Now for the second sub-case, i.e. $\text{og}(\sigma, o) = [v \mapsto (C, \vec{f} : \vec{v}, 0, \emptyset)] \cup \text{og}(\sigma_i, o_i)$, examining Algorithm 4.1, we have:

$$\sigma(v) = (C, \vec{f} : \vec{v}, n, \{\vec{c}\}) \text{ with } \text{local}(C), \{\vec{c}\} = \text{fn}(\vec{v}) \quad (5)$$

$$\sigma_1 = \sigma \setminus \{v\} \quad (6)$$

$$\sigma_{i+1} = \sigma_i \setminus \text{dom}(\text{og}(\sigma_i, o_i)) \quad (7)$$

From because $v \in \text{dom}(\sigma)$ and (6) we have that $\text{length}(\sigma_1) < n$. Examining (7), $\text{length}(\sigma_i) \leq \text{length}(\sigma_1) < \text{length}(\sigma)$ and so by the inductive hypothesis:

$$\begin{aligned} \Gamma; \emptyset \vdash (\sigma'_i =) \text{og}(\sigma_i, o_i) : \text{ok}, \quad \forall o' \in \text{dom}(\sigma'_i). \sigma'_i(o') &= (C, \dots) \text{ with } \text{local}(C) \\ \forall o' \in (\text{fn}(\sigma'_i) \setminus \text{dom}(\sigma'_i)). \Gamma \vdash o' : C \text{ with } \text{remote}(C) \end{aligned}$$

By applying Lemma D.2 (7), and by set-union we have:

$$\begin{aligned} \Gamma; \emptyset \vdash (\sigma'' =) \bigcup \text{og}(\sigma_i, o_i) : \text{ok}, \quad \forall o' \in \text{dom}(\sigma''). \sigma''(o') &= (C, \dots) \text{ with } \text{local}(C) \\ \forall o' \in (\text{fn}(\sigma'') \setminus \text{dom}(\sigma'')). \Gamma \vdash o' : C \text{ with } \text{remote}(C) \end{aligned}$$

By typability of σ , we have that $\Gamma; \Delta \vdash (C, \vec{f} : \vec{v}, n, \{\vec{c}\}) : \text{ok}$. Therefore we can straightforwardly deduce: $\Gamma; \emptyset \vdash (C, \vec{f} : \vec{v}, 0, \emptyset) : \text{ok}$, and by examination of (5), $\{\vec{c}\} = \text{fn}(\vec{v})$ and any identifiers in $\vec{o}$ that point to instances of local classes will have been gathered by recursive application of the algorithm, and hence in $\text{dom}(\sigma'')$. Thus we finish the case.

(2) Assume $\sigma' = \text{og}(\sigma, v)$. The base case is where $\text{length}(\sigma) = 0$ i.e. $\sigma = \emptyset$. Examining the algorithm we see that $\sigma' = \emptyset$. Therefore trivially $\text{ogcomp}(\sigma, \sigma')$.

For the inductive step, assume that $\sigma' = \text{og}(\sigma, v)$ and $\text{ogcomp}(\sigma, \sigma')$ for $\text{length}(\sigma) < n$. Now setting $\text{length}(\sigma) = n$ there are two sub-cases:

- (a) $\sigma' = \emptyset$. Trivially, there are no pairs in the reachability relation $\text{RCH}(\sigma')$ and so we have $\text{ogcomp}(\sigma, \sigma')$ vacuously.
- (b) $\sigma' = [v \mapsto \sigma(v)] \cup \text{og}(\sigma_i, o_i)$, where $\sigma(v) = (C, \vec{f} : \vec{v})$, $\{\vec{o}\} = \text{fn}(\vec{v})$, $\sigma_1 = \sigma \setminus \{o\}$ and $\sigma_{i+1} = \sigma_i \setminus \text{dom}(\text{og}(\sigma_i, o_i))$.

Clearly, $\text{length}(\sigma_i) < n$ by the initial removal of o from σ_1 . Write $\sigma'_i = \text{og}(\sigma_i, o_i)$. Further examination of the algorithm shows that σ'_{i+1} is computed from σ_i less the elements collected in σ'_i . Therefore $\text{dom}(\sigma'_i) \cap \text{dom}(\sigma'_{i+1}) = \emptyset$ so by the inductive hypothesis $\text{ogcomp}(\sigma_1, \bigcup \text{og}(\sigma_i, o_i))$. Recall that $\sigma_1 = \sigma \setminus \{v\}$. By adding $[v \mapsto \sigma(v)]$ to each side, we add the same number of reachable states and therefore $\text{ogcomp}(\sigma, \sigma')$.

(3) By induction on $\text{length}(\text{CT})$.

(4) Suppose $\text{length}(\text{CT}_0) = 0$ then by definition, CT' is complete. Now, take $\text{length}(\text{CT}_n) = n$. Given $\text{CT}' = \text{cg}(\text{CT}_n, C)$ for some C is complete by assumption, we either have that $C \in \text{dom}(\text{CT}_n)$ and $\text{CT}' \neq \emptyset$, or $C \notin \text{dom}(\text{CT}_n)$ and $\text{CT}' = \emptyset$. For the inductive step we must show that when the length of the class table is $n + 1$ the computed class graph remains complete. Extending the class table can be achieved by appending a new entry giving $\text{CT}_{n+1} = \text{CT}_n \cdot [C' \mapsto L]$ for some $C' \notin \text{dom}(\text{CT}_n)$. We assume that the superclass of C' is present in CT_n , otherwise the new class table would not be complete and so the conclusion would hold by default. Then given $\text{CT}' = \text{cg}(\text{CT}_{n+1}, C)$, if $C \neq C'$ then again CT' is complete by virtue of being empty. If $C = C'$ then by our assumption that the class table CT_{n+1} contains the direct superclass of C' then CT' must also be complete.

D. BASIC PROPERTIES

In this Appendix we shall show some key properties and lemmas that are necessary for the proof of our network invariance and type soundness theorem. Hereafter we often write α for U or S . We also adopt the convention that $\Gamma; \emptyset$ can be written as simply Γ .

D.1 Judgements

Lemma D.1 lists some useful properties about judgements. We write J to stand for any one of the following judgements:

$$J ::= \text{Env} \mid \sigma : \text{ok} \mid e : \alpha \mid P : \text{thread} \mid F : \text{conf} \mid N : \text{net}$$

Lemma D.1linearity has the useful property of ensuring that any channels appearing in the channel environment Δ and not in the judgement J must have the linear type **chan**.

LEMMA D.1. (Judgements)

(Permutation of environments)

(1) $\Gamma; \Delta, c : \tau, c' : \tau', \Delta' \vdash J \implies \Gamma; \Delta, c' : \tau', c : \tau, \Delta' \vdash J$.

(2) $\Gamma, u : T, u' : T', \Gamma'; \Delta \vdash J \implies \Gamma, u' : T', u : T, \Gamma'; \Delta \vdash J$. Similarly for **this**.

(Linearity of channels)

(3) $\Gamma; \Delta, c : \tau, \Delta' \vdash J \wedge c \notin \text{fn}(J) \implies \tau = \text{chan}$.

(Weakening)

(4) $\Gamma; \Delta \vdash J \wedge c \notin \text{dom}(\Delta) \implies \Gamma; \Delta, c : \text{chan} \vdash J$.

(5) $\Gamma; \Delta \vdash J \wedge \vdash T : \text{tp} \wedge x \notin \text{dom}(\Gamma) \implies \Gamma, x : T; \Delta \vdash J$.

(6) $\Gamma; \Delta \vdash J \wedge \vdash C : \text{tp} \wedge \text{this} \notin \text{dom}(\Gamma) \implies \Gamma, \text{this} : C; \Delta \vdash J$.

(Strengthening)

- (7) $\Gamma; \Delta, c : \tau \vdash J \wedge c \notin \text{fn}(J) \implies \Gamma; \Delta \vdash J$.
 (8) $\Gamma, u : T; \Delta \vdash J \wedge u \notin \text{fnv}(J) \implies \Gamma; \Delta \vdash J$.

(Implied judgements)

- (9) $\Gamma, \Gamma'; \Delta, \Delta' \vdash J \implies \Gamma; \Delta \vdash \text{Env}$.

PROOF. By induction on the size of the judgement J . All cases are straightforward. We only list the proof for weakening with the case $J = P_1 \mid P_2 : \text{thread}$. After applying rule TT-PAR we have two cases; we can apply the inductive hypothesis to either the left branch or the right branch of the parallel composition. For example, choose the left branch. Therefore $\Gamma; \Delta_1, c : \text{chan} \vdash P_1 : \text{thread}$ and $\Delta_1, c : \text{chan} \asymp \Delta_2$ as $c \notin \text{dom}(\Delta_2)$. Apply TT-PAR to yield $\Gamma; \Delta_1, c : \text{chan} \odot \Delta_2 \vdash P_1 \mid P_2 : \text{thread}$. Then $\Gamma; \Delta_1 \odot \Delta_2, c : \text{chan} \vdash P_1 \mid P_2 : \text{thread}$ by definition of $\odot$. The other case proceeds similarly. $\square$

D.2 Stores

Lemma D.2 states properties about the type-safety of store access. Store access are defined as adding new variable and object identifier mappings, updating the fields of objects and the value held by a variable, and also retrieving information from variables and object fields. Lemma D.2 allows the concatenation of disjoint stores and is useful in typing the `deserialize(e)` operation.

LEMMA D.2. (Stores) *Assuming that $\Gamma; \Delta \vdash \sigma : \text{ok}$. Then:*

- (1) *If $\Gamma \vdash v : T'$ with $x \notin \text{dom}(\Gamma)$ and $T' <: T$ then $\Gamma, x : T; \Delta \vdash \sigma \cdot [x \mapsto v] : \text{ok}$.*
- (2) *Assume $\Gamma \vdash x : T$ and $\Gamma \vdash v : T'$ with $T' <: T$. Then $\Gamma; \Delta \vdash \sigma[x \mapsto v] : \text{ok}$.*
- (3) *$\Gamma \vdash x : T$ implies $\Gamma \vdash \sigma(x) : T'$ with $T' <: T$.*
- (4) *If $\Gamma; \Delta \vdash (C, \vec{f} : \vec{v}, n, \{\vec{c}\}) : \text{ok}$ and $o \notin \text{dom}(\Gamma)$ then we have:
 $\Gamma, o : C; \Delta \vdash \sigma \cdot [o \mapsto (C, \vec{f} : \vec{v}, n, \{\vec{c}\})] : \text{ok}$.*
- (5) *If $\Gamma \vdash o : C$ and $\Gamma \vdash v : T'_i$ with $\text{fields}(C) = \vec{T}\vec{f}$ and $T'_i <: T_i$, then we have
 $\Gamma; \Delta \vdash \sigma[o \mapsto \sigma(o)[f_i \mapsto v]] : \text{ok}$.*
- (6) *Assume $\Gamma \vdash o.f_i : T_i$ with $\sigma(o) = (C, \vec{f} : \vec{v})$. Then $\Gamma \vdash v_i : T'_i$ where $T'_i <: T_i$.*
- (7) *Suppose $\Gamma, \Gamma'; \Delta, \Delta' \vdash \sigma' : \text{ok}$ with $\text{dom}(\sigma) \cap \text{dom}(\sigma') = \emptyset$, then $\Gamma, \Gamma'; \Delta, \Delta' \vdash \sigma \cup \sigma' : \text{ok}$.*
- (8) *$\Gamma; \Delta \vdash \sigma : \text{ok}$ and $\sigma' \subseteq \sigma$ implies $\Gamma; \Delta \vdash \sigma' : \text{ok}$.*
- (9) *Suppose $\Gamma; \Delta \vdash \sigma : \text{ok}$ and $\sigma' = \text{block}(\sigma, o, c)$ with $c \notin \text{dom}(\Delta)$, $o \in \text{dom}(\sigma)$. Then we have that $\Gamma; \Delta, c : \text{chan0}(\text{void}) \vdash \sigma' : \text{ok}$.*
- (10) *Suppose $\Gamma; \Delta, c : \text{chan0}(\text{void}) \vdash \sigma : \text{ok}$ and $\sigma' = \text{unblock}(\sigma, o, c)$ with $o \in \text{dom}(\sigma)$ and $c \notin (\text{fn}(\sigma) \setminus \text{fn}(\sigma(o)))$. Then we have that $\Gamma; \Delta \vdash \sigma' : \text{ok}$.*

PROOF. All are mechanical. $\square$

We list the standard lemma for the typability of the method body. The proof is routine.

LEMMA D.3. (Method body) *Suppose $\text{mbody}(m, C, \text{CT}) = (x, e)$ and $\text{mtype}(m, C) = D \rightarrow U$ with $\vdash \text{CT} : \text{ok}$. Then for some C' where $C <: C'$ and some U' where $U' <: U$ then we have $x : D, \text{this} : C' \vdash e : \text{ret}(U')$.*

D.3 Structural equivalence: Proof of Lemma 6.2

An important property to be shown is that the application of the structural equality rules given in Figure 16 preserves the typing of a term. In order to prove this property, the next lemma is important: it yields natural properties for the composability of environments and is used in many of the later proofs.

LEMMA D.4. (*Commutativity of composition and composability*)

- (1) $\Delta_1 \asymp \Delta_2$ and $(\Delta_1 \odot \Delta_2) \asymp \Delta_3 \iff \Delta_2 \asymp \Delta_3$ and $\Delta_1 \asymp (\Delta_2 \odot \Delta_3)$.
 (2) $\Delta_1 \asymp \Delta_2$ and $(\Delta_1 \odot \Delta_2) \asymp \Delta_3 \implies (\Delta_1 \odot \Delta_2) \odot \Delta_3 = \Delta_1 \odot (\Delta_2 \odot \Delta_3)$.

PROOF. In both proofs, without loss of generality we consider singleton environments such that $\Delta_1 = \{c : \text{chanI}(U)\}$ and $\Delta_2 = \{c : \text{chanO}(U)\}$ with $\Delta_1 \odot \Delta_2 = \{c : \text{chan}\}$. For (1), we show only the left-to-right direction, the opposite direction is similar. The only interesting case is that Δ_1 and Δ_2 share the same channels. By the definition of $\asymp$, we know $c \notin \text{dom}(\Delta_3)$. Since $\Delta_2 \odot \Delta_3 = \{c : \text{chanO}(U)\} \cup \Delta_3$, we have that $\Delta_2 \asymp \Delta_3$ as required. We can also easily check $\Delta_1 \odot (\Delta_2 \odot \Delta_3)$ is defined, thus by definition of $\asymp$, we have $\Delta_1 \asymp (\Delta_2 \odot \Delta_3)$, as desired. (2) proceeds in a similar manner to (1), adopting the same singleton environments. $\square$

E. PROOF OF INVARIANT PROPERTIES

E.1 Proofs of Lemma 6.3

(1) Straightforward by examining the reduction rules that modify class tables: DEFROST and DOWNLOAD.

(2) Straightforward by examining the reduction rules, starting from the initial property.

(3) Assume $\text{reachable}(\sigma_{im+1}, P_{im+1}, o)$ with $\sigma_{im+1}(o) = (C, \dots)$.

Then there are four cases.

(a) Suppose $\text{reachable}(\sigma_{im}, P_{im}, o)$ with $\sigma_{im}(o) = (C, \dots)$.

Then by the inductive hypothesis, we have two possible situations:

- i. $\text{comp}(C, \text{CT}_{im})$, hence by Lemma 6.3 (1) we have $\text{comp}(C, \text{CT}_{im+1})$.
- ii. $P_{im} \equiv E[\text{download } \vec{C} \text{ from } l_j \text{ in } e] \mid Q_{im}$ or $P_{im} \equiv E[\text{resolve } \vec{C} \text{ from } l_j \text{ in } e] \mid Q_{im}$, with a superclass of C in $\vec{C}$. Examining the reduction rules, we see that if the last rule applied was DNOTHING, then by definition $\text{comp}(C, \text{CT}_{im+1})$. If the last reduction rule applied was RESOLVE, then we see that $P_{im+1} \equiv E[\text{resolve } \vec{C}' \text{ from } l_j \text{ in } e] \mid Q_{im}$ with a superclass of C in $\vec{C}'$.

(b) Suppose $\neg \text{reachable}(\sigma_{im}, P_{im}, o)$ with $\sigma_{im}(o) = (C, \dots)$.

Then the last reduction rule applied must have been RETURN or LEAVE. We shall consider the case of the latter; the former is similar. Then we have that $P_{im+1} \equiv o.m(\text{deserialize}(v))$ with $c \mid Q_{im+1}$. Since o moved from another location, we can conclude $\text{remote}(C)$, hence it must have been created at location l_i by NEW, or was there in the initial network (recall that remote object identifiers cannot leak to other locations, nor can their store entry be carried in a frozen value).

In the case of NEW, we had that for some step k (with $N_0 \rightarrow N_k \rightarrow N_m$) when o was created, $\text{comp}(C, \text{CT}_{ik})$. Then by Lemma 6.3 (1), $\text{comp}(C, \text{CT}_{im+1})$. If o was

present in the initial network, then by initial conditions and application of this Lemma again, $\text{comp}(C, \text{CT}_{im+1})$.

When o becomes reachable, all objects that it points to (both directly and indirectly) now become reachable. Suppose without loss of generality that $[o \mapsto (C, f : o')] \in \sigma_{im}$, i.e. we only have one field, that points to a newly reachable object o' . If $o' \notin \text{dom}(\sigma_{im})$, then it must be remote, and so the property holds immediately.

If $o' \in \text{dom}(\sigma_{im})$, we must investigate further. Suppose there exists some step k such that $P_{ik} \equiv E[\text{new } C(o')] \mid Q_{ik} \longrightarrow_{l_i} E[o] \mid Q_{ik+1}$. Since o' is a non-remote identifier, we have that it was either instantiated by NEW (hence the class it belongs to is complete), or it is the result of a defrost operation. By examining the rule DEFROST, before this assignment could take place, we must have downloaded all superclasses of o' 's class, hence it must be complete. Similarly if $P_{ik} \equiv E[o.f = o'] \mid Q_{ik}$.

If o was added to the current location by DEFROST, then by the completeness of the object graph of the frozen object, we had that o' was also added in the same step. Hence after downloading all superclasses of o , we must have also downloaded all superclasses of o' .

(c) Suppose $\text{reachable}(\sigma_{im}, P_{im}, o)$ with $o \notin \text{dom}(\sigma_{im})$.

For this situation to arise, we have o of some class C such that $\text{remote}(C)$. Since remote object identifiers cannot move, this case is complete by contradiction: no reduction to N_{m+1} can occur.

(d) Suppose $\neg \text{reachable}(\sigma_{im}, P_{im}, o)$ with $o \notin \text{dom}(\sigma_{im})$.

The last reduction rule applied was DOWNLOAD or NEW. In the case of the former, we see that $P_{im} \equiv E[\text{defrost}(v; \dots)] \mid Q_{im}$ reduces to

$P_{im+1} \equiv E[\text{download } \vec{C} \text{ from } l_j \text{ in } e] \mid Q_{im+1}$ as required. In the case of the latter, we had $\text{comp}(C, \text{CT}_{im})$ and so by Lemma 6.3 (1) $\text{comp}(C, \text{CT}_{im+1})$.

(4) Obvious by repeating DOWNLOAD and RESOLVE until we reach resolution of C' . Note that these reductions terminate as the inheritance relations in a well-formed class table is acyclic.

E.2 Proofs of Lemma 6.4

Induction on k . Below we assume $c \in \text{blocked}(o, \sigma_{im+1})$ in $\text{waiting}(c) \dots$

(1) (a) The base case is when $k = 1$. To generate $E_1[\text{insync } o \{e\}], \sigma_{i1}$, it must have been the case that:

$$E_1[\text{sync } (o) \{e\}], \sigma_{i0} \longrightarrow_l E_1[\text{insync } o \{e\}], \sigma_{i1}$$

By the initial conditions, $\text{getLock}(\sigma_{i0}, o) = 0$ and so by application of SYNC in we have that $\text{getLock}(\sigma_{i1}, o) = 1$ as required.

For the inductive case, we assume the hypothesis for N_m and show for N_{m+1} . Suppose we have a configuration of the form:

$$E_1[\dots E_p[\text{insync } o \{e_p\}] \dots], \sigma_{im+1}, \\ \text{with } e_p \neq E'[\text{insync } o \{e'\}], E'[\text{waiting}(c) \dots], E'[\text{ready } o \dots]$$

Then the last reduction was either $e \longrightarrow_l e_p$ for some e or $E_p[\text{sync } (o) \{e_p\}] \longrightarrow_l$

$E_p[\text{insync } o \{e_p\}]$.

By our assumption that $e_p \neq E'[\text{insync } o \{e'\}]$, in the first case the only interesting reduction rule to consider is the application of LEAVECRITICAL. Suppose $e \equiv E[\text{insync } o \{v\}]$, then there are $p + 1$ nested acquisitions of the monitor o . By the inductive hypothesis we have that $\text{getLock}(\sigma_{im}, o) = p + 1$. By premise of LEAVECRITICAL $\sigma_{im+1} = \text{setLock}(\sigma_{im}, o, p)$, and so $\text{getLock}(\sigma_{im+1}, o) = p$ as required. The case for $e \equiv E[\text{insync } o \{\text{return}(c) v\}]$ is similar.

For the second case, the last reduction rule applied was SYNC. Before application there are $p - 1$ levels of nested monitors. By the inductive hypothesis it must be the case that $\text{getLock}(\sigma_{im}, o) = p - 1$. Then by the premise of SYNC we see that $\sigma_{im+1} = \text{setLock}(\sigma_{im}, o, p)$, and so $\text{getLock}(\sigma_{im+1}, o) = p$ as required.

(b) Straightforward using (a) and inspecting NOTIFY.

(c) Establishing that $p = n'$ is straightforward using (a).

(2) Base case, $k = 1$. Now suppose: $P_{i0}, \sigma_{i0}, \text{CT}_{i0} \longrightarrow_l P_{i1}, \sigma_{i1}, \text{CT}_{i1}$. By the initial conditions $\text{getLock}(\sigma_{i0}, o) = 0$, and by assumption $\text{getLock}(\sigma_{i1}, o) = n$ with $n > 0$. As no run-time syntax can exist in the network initially, the reduction rule applied was SYNC. This means that $P_{i0} \equiv E[\text{sync } (o) \{e\}] \mid Q_{i0}$, and examining the conclusion of the rule $P_{i1} \equiv E[\text{insync } o \{e\}] \mid Q_{i1}$. Again by the initial conditions, e cannot contain $\text{insync } o \{\dots\}$ as a sub-term, completing this case.

For the inductive step, suppose $P_{im}, \sigma_{im}, \text{CT}_{im} \longrightarrow_l P_{im+1}, \sigma_{im+1}, \text{CT}_{im+1}$. By assumption $\text{getLock}(\sigma_{im+1}, o) = p$ and $p > 0$. There are four distinct cases:

- (1) $\text{getLock}(\sigma_{im}, o) = 0$. The last rule applied to derive P_{im} could be either SYNC or READY. To apply the former it must be the case that $P_{im} \equiv E_1[\text{sync } (o) \{e\}] \mid Q_{im}$. By the conclusion of this rule $P_{im+1} \equiv E_1[\text{insync } o \{e\}] \mid Q_{im+1}$ as required. For application of READY, we have that

$$P_{im} \equiv E_1[\text{insync } o \{\dots E_x[\text{insync } o \{E[\text{ready } o x]\}] \dots\}] \mid Q_{im}$$

with $\text{insync } o \{\dots\}$ not a sub-term of E by 1.(a). Then it remains to show that $x = p$, however this is immediate by inspection of the rule READY.

- (2) $\text{getLock}(\sigma_{im}, o) = p - 1$. The only rule applicable in this situation is SYNC. Therefore by the inductive hypothesis:

$$P_{im} \equiv E_1[\text{insync } o \{\dots E_{p-1}[\text{insync } o \{E_p[\text{sync } (o) \{e\}]\}] \dots\}] \mid Q_{im}$$

with $\text{insync } o \{\dots\}$ not a sub-term of E_p . Examining SYNC, we have that:

$$P_{im+1} \equiv E_1[\text{insync } o \{\dots E_{p-1}[\text{insync } o \{E_p[\text{insync } o \{e\}]\}] \dots\}] \mid Q_{im+1}$$

By the initial conditions $\text{insync } o \{\dots\}$ cannot be a sub-term of e , so this completes the case.

- (3) $\text{getLock}(\sigma_{im}, o) = p$. Straightforward.

- (4) $\text{getLock}(\sigma_{im}, o) = p + 1$. Only one rule is applicable in this situation: LEAVECRITICAL.

To apply this, P_{im} must be $E_1[\text{insync } o \{\dots E_{p+1}[\text{insync } o \{e\}] \dots\}] \mid Q_{im}$ with $e = v$ or $e = \text{return}(c) v$. We consider the case for v , the case for a return is similar. Examining LEAVECRITICAL, P_{im+1} must be $E_1[\text{insync } o \{\dots E_{p+1}[v] \dots\}] \mid Q_{im+1}$ with $\text{insync } o \{\dots\}$ not a sub-term of E_{p+1} by our earlier assumption, completing this case. $\square$

E.3 Proofs of the Invariant properties

Inv(1) By Lemma 6.3 (1).

Inv(2) Suppose $P_{im} \not\equiv E[\mathbf{new} C(\vec{v})] \mid Q_{im} \longrightarrow_{l_i} P_{im+1} \equiv E[\mathbf{new} C(\vec{v})] \mid Q_{im}$, then one of two possible reduction rules was applied:

- (a) We applied CONG. Then $P_{im} \equiv P'_{im} \mid Q_{im}$ with $P'_{im} \longrightarrow_{l_i} E[\mathbf{new} C(\vec{v})]$.
Then if $m = 0$, by the initial condition Inv(2)' and Lemma 6.3 (1) we have $\mathbf{comp}(C, \mathbf{CT}_{im+1})$, irrespective of the reduction rule applied.
Suppose $m > 0$. Therefore $P'_{im} \equiv E'[\mathbf{new} C(\vec{v})]$ for some context E' . By the inductive assumption we have $\mathbf{comp}(C, \mathbf{CT}_{im})$ and again by Lemma 6.3 (1) it is the case that $\mathbf{comp}(C, \mathbf{CT}_{im+1})$.
- (b) We applied DNOTHING. Then $P_{im} \equiv E[\mathbf{download} \vec{C} \text{ from } l_j \text{ in } \mathbf{new} C(\vec{v}) \mid Q_{im}$ with $\vec{C} \in \mathbf{dom}(\mathbf{CT}_i)$. Then it must be the case that $m > 0$ since the download expression is not permissible runtime syntax in an initial network. In order to download nothing, it must have been the case that $P_{ik} \equiv E[\mathbf{download} \vec{D} \text{ from } l_j \text{ in } e]$ with $C \in \vec{D}$ and $k < m$ (i.e. class C was downloaded at some point in the past). Then examining the rules DOWNLOAD and RESOLVE we can straightforwardly observe that they iterate until all superclasses of C are downloaded. Therefore using Lemma 6.3 (1) we have trivially that $\mathbf{comp}(C, \mathbf{CT}_{im+1})$.

Inv(3) There are two interesting sub-cases:

- (a) The last applied reduction rule was DOWNLOAD. Then

$$P_{im} \equiv E[\mathbf{download} C \text{ from } l_j \text{ in } e] \mid Q_{im}$$

$$\text{and } P_{im} \longrightarrow_{l_i} E[\mathbf{resolve} C \text{ from } l_j \text{ in } e]$$

Since downloading did not fail (the assumption that $N_{m+1} \not\models \mathbf{Err}$), there must exist a location l_j with $C \in \mathbf{dom}(\mathbf{CT}_{jm})$. By the premise of DOWNLOAD, $C \in \mathbf{dom}(\mathbf{CT}_{im+1})$ and $\mathbf{CT}_{im+1}(C) = \mathbf{CT}_{jm}$ modulo class labelling as required.

- (b) The last applied reduction rule was DEFROST. Then

$$P_{im} \equiv E[\mathbf{defrost}(v; \lambda(T x).(\nu \vec{u})(l_j, e, \sigma, \mathbf{CT}))]$$

with $\mathbf{CT} \subseteq \mathbf{CT}_{jk}$ where $k < m$. Straightforwardly, by premise of DEFROST we have that $\mathbf{CT}_{im+1} = \mathbf{CT}_{im} \cup \mathbf{CT}[\vec{C}^m / \vec{C}]$ for some classes $\vec{C}$, and so making a similar argument to the previous sub-case, any duplicate classes *must* have the same definition, modulo class labelling.

Inv(4) The only interesting cases are those where the set of free variables of a term changes with reduction. There are three such cases. Without loss of generality, we consider only a single thread containing no free variables for a single location with an empty store:

- (a) The last applied reduction rule was DEC. Then suppose

$$l_i[E[T x = v], \emptyset, \mathbf{CT}_i] \longrightarrow \equiv (\nu x)(l_i[E[v], [x \mapsto v], \mathbf{CT}_i])$$

By Inv(12), we have $\mathbf{fv}(v) = \emptyset$. Before reduction we have $\mathbf{fv}(E[T x = v]) = \emptyset$, after reduction we potentially have that $\mathbf{fv}(E[v]) = \{x\}$, however we see that $\mathbf{dom}([x \mapsto v]) = \{x\}$ and the new identifier is restricted at the network level. Therefore this case is complete.

(b) The last applied reduction rule was METHINVOKE. Then suppose

$$l_i[o.m(v) \text{ with } c, \emptyset, \text{CT}_i] \longrightarrow \equiv (\nu x)(l_i[e[o/\text{this}][\text{return}(c)/\text{return}], [x \mapsto v], \text{CT}_i])$$

where $\text{mbody}(m, C, \text{CT}_i) = (x, e)$, and again by $\text{Inv}(12)$, $\text{fv}(v) = \emptyset$. We must show that $\text{fv}(e) \subseteq [x \mapsto v] \subseteq \{x\}$. However by definition of substitution, we know that $\text{fv}(e[o/\text{this}][\text{return}(c)/\text{return}]) = \text{fv}(e)$. Given that the network configuration is well-typed, it must be the case that $x : D, \text{this} : C \vdash e : \text{ret}(U)$, i.e. $\text{fv}(e) \subseteq \{x\}$. This concludes the case.

(c) The last applied reduction rule was DEFROST. Suppose

$$\begin{aligned} & l_i[E[\text{defrost}(v; \lambda(T \ x).(\nu \vec{u})(l_j, e, \sigma, \text{CT}))], \emptyset, \text{CT}_i] \\ & \longrightarrow \equiv (\nu \vec{u}x)(l_i[E[\text{download } \vec{F} \text{ from } l_j \text{ in sandbox } \{e\}], \sigma \cdot [x \mapsto v], \text{CT} \cup \text{CT}_i[\vec{C}^{l_j}/\vec{C}]] \end{aligned}$$

where $\text{fv}(v) = \text{fv}(\lambda(T \ x).(\nu \vec{u})(l_j, e, \sigma, \text{CT})) = \emptyset$ by $\text{Inv}(12)$. Straightforwardly we have that $\text{fv}(e) \subseteq \text{dom}(\sigma \cdot [x \mapsto v]) \subseteq \{\vec{u}x\}$ to complete this case.

Inv(5) Here we only need to examine the case when the last applied rule was DEFROST. This case is straightforward: all new identifiers added to the store during defrosting are restricted with fresh names, therefore there can be no overlap of store entries between producer and consumer locations. Moreover, by $\text{Inv}(15)$, frozen values contain only remote identifiers, hence the rules METHREMOTE and LEAVE cannot move shared store entries across the network.

Inv(6) For the first case suppose that:

$$o \in \text{fn}(F_{im+1}) \cap \text{fn}(F_{jm+1}), \quad o \in \text{fn}(F_{im}) \cap \text{fn}(F_{jm}), \quad \exists !k. \sigma_{km}(o) = (C, \dots) \text{ with } \text{remote}(C)$$

Then by $\text{Inv}(5)$ we have that o cannot exist as an entry in more than one store. By $\text{Inv}(8)$ we observe that since $o \in \text{fn}(F_{im+1})$, the entry for o cannot have been garbage collected. Hence $\exists !k. \sigma_{km+1}(o) = (C, \dots)$, and since there are no operations to change the “remoteness” of a class, $\text{remote}(C)$ holds, completing the case.

For the second case suppose that:

$$o \in \text{fn}(F_{im+1}) \cap \text{fn}(F_{jm+1}), \quad o \notin \text{fn}(F_{im}) \cap \text{fn}(F_{jm}) \quad (8)$$

This indicates that a free name moved between two locations. This can happen in two ways: by application of METHREMOTE or RETURN. We show the case of the latter, since the former is proved by the same argument. Assume

$$l_j[\text{go } v \text{ to } c] \mid l_i[Q_{im}] \longrightarrow l_j[\dots] \mid l_i[\text{return}(c) \text{ deserialize}(v) \mid Q_{im}]$$

By typability of F_{im+1} we have that $\Gamma; \Delta \vdash v : \text{unit} \rightarrow D$ for some class D , i.e. v is a *frozen* value. Assume there exists o such that $o \in \text{fn}(v)$ and (8) holds. By $\text{Inv}(8)$, there exists a store entry for o “somewhere” and by $\text{Inv}(5)$ this entry must be unique. Hence there exists unique k such that $\sigma_{km+1}(o) = (C, \dots)$. Now by $\text{Inv}(15)$, the only free identifiers in a frozen value must be of remote classes, hence $\text{remote}(C)$. This completes the case.

Inv(7) There are two sub-cases. First assume that

$$o \in \text{fn}(F_{im+1}) \wedge \exists k. \sigma_{km+1}(o) = (C, \dots) \wedge \text{local}(C)$$

and we shall prove that $k = i$.

- (a) Suppose $o \in \text{fn}(F_{im}) \wedge \exists k'. \sigma_{k'm}(o) = (C, \dots) \wedge \text{local}(C) \wedge k' = i$. Then by the inductive assumption we can derive that $o \in \text{dom}(\sigma_{im})$. By **Inv(8)** we have that $o \in \text{dom}(\sigma_{im+1})$ which implies $i = k = k'$ as required.
- (b) Suppose that $o \notin \text{fn}(F_{im})$, then the last applied reduction step applied must have somehow created a new object identifier o . Examining the reduction rules we have four cases (although two are trivial):
 - i. The last applied rule was **NEW**.

$$E[\text{new } C(\vec{v})] \mid Q_{im}, \sigma_{im} \longrightarrow_{l_i} (\nu o)(E[o] \mid Q_{im}, \sigma_{im} \cdot [o \mapsto \dots])$$

Then straightforwardly $k = i$ because $o \in \text{dom}(\sigma)_{im+1}$.

- ii. The last applied rule was **DEFROST**.

$$E[\text{defrost}(v; \lambda(T \ x).(\nu \vec{u})(l_j, e, \sigma, \text{CT}))] \mid Q_{im}, \sigma_{im} \longrightarrow_{l_i} (\nu \vec{u}x)(E[\dots] \mid Q_{im}, \sigma_{im} \cup \sigma \cdot [x \mapsto v])$$

Then by **Inv(15)**, the frozen value can contain no free local object identifiers and so $o \in \vec{u}$ which entails that $o \in \text{dom}(\sigma)$ and $k = i$.

- iii. The last applied rule was **LEAVE** or **METHREMOTE**. Then (omitting stores and class tables for conciseness)

$$l_j[\text{go } o.m(v) \text{ with } c] \mid l_i[Q_{im}] \longrightarrow l_j[\dots] \mid l_i[o.m(\text{deserialize}(v)) \text{ with } \mid Q_{im}]$$

By assumption, $\Gamma; \Delta \vdash N : \text{net}$ and therefore $\Gamma \vdash v : \text{unit} \rightarrow D$ for some class D . Then by **Inv(15)**, we know that any free names appearing inside v must be *remote* object identifiers. Therefore this case holds vacuously.

Inv(8) Assume: $o \in \text{fn}(F_{im}) \wedge \exists k \ 1 \leq k \leq n. o \in \text{dom}(\sigma_{km})$ and $o \in \text{fn}(F_{im+1})$. Then, by examination of the rules for structural equivalence in Figure 16, we see that the only way to remove an object identifier is when it does not exist in the free names of the remainder of the network. Since $o \in \text{fn}(F_{im+1})$ by assumption, it must be the case that $\exists k \ 1 \leq k \leq n. o \in \text{dom}(\sigma_{km+1})$ as required.

Inv(9) Only the cases for $R_i \equiv o.m(e) \text{ with } c$ and $R_i \equiv E[o.f = e]$ are shown, as the others use the same basic method.

- (a) Suppose $P_{im+1} \equiv o.m(e) \text{ with } c \mid Q_{im+1}$. Examining the structure of this thread, we see that there were two possible rules applied in the last reduction step: **LEAVESANDBOX** or **METHLOCAL**.

- i. Let $P_{im} \equiv E[o.m(v)] \mid Q_{im}$. Then this is a local method call, and by the premises of **METHLOCAL**, we had $o \in \text{dom}(\sigma_{im})$, hence $\sigma_{im} = (C, \dots)$. Then by Lemma 6.3 (3) we have $\text{comp}(C, \text{CT}_{im})$. By monotonicity of stores and class tables, $\sigma_{im+1}(o) = (C, \dots)$ and $\text{comp}(C, \text{CT}_{im+1})$ as required.

- ii. Let $P_{im} \equiv Q_{im}$. Then this is a remote method call, and by the premises of **LEAVE**, we have $o \in \text{dom}(\sigma_{im})$. Then proof proceeds as in the previous case.

- (b) Assume $P_{im+1} \equiv E[o.f = e] \mid Q_{im+1}$. Then we have two main cases. If $P_{im} \equiv E[o.f = e'] \mid Q_{im}$ then by the inductive hypothesis, this case is complete. However suppose $P_{im} \equiv E[e'.f = e] \mid Q_{im}$, then we must perform a case analysis on the step $e' \longrightarrow o$.

- i. Suppose the last rule applied was **NEW**. Then by **Inv(2)** we have $\text{comp}(C, \text{CT}_{im})$, hence by Lemma 6.3 (1) $\text{comp}(C, \text{CT}_{im+1})$.
- ii. Suppose the last applied rule was **FLD**. Then $P_{im} \equiv E[(o'.f').f = e] \mid Q_{im}$. Then by typability of N_m we have that $\Gamma, \vec{u} : \vec{T} \vdash (o'.f').f : C$ and $\text{local}(C)$. Then by **Inv(8)** and **Inv(7)** we have that $\sigma_{im}(o) = (C, \dots)$. Then by Lemma 6.3 (3),

- $\text{comp}(C, \text{CT}_{im})$. By monotonicity of class tables and stores, $\sigma_{im+1}(o) = (C, \dots)$ and $\text{comp}(C, \text{CT}_{im+1})$ as required.
- iii. Suppose the last rule applied was VAR. Then $P_{im} \equiv E[x.f = e] \mid Q_{im}$. By typability, $\Gamma, \vec{u} : \vec{T} \vdash x.f : C$ with $\text{local}(C)$. By reduction, we must have that $[x \mapsto o] \in \sigma_{im}$ and so o is reachable from thread P_{im} . Again by Inv(8) and Inv(7) we have that $\sigma_{im}(o) = (C, \dots)$. Then by Lemma 6.3 (3), $\text{comp}(C, \text{CT}_{im})$. By monotonicity of class tables and stores, $\sigma_{im+1}(o) = (C, \dots)$ and $\text{comp}(C, \text{CT}_{im+1})$ as required.
 - iv. Suppose the last rule applied was ASS. Then $P_{im} \equiv E[(x = o).f = e] \mid Q_{im}$. Then by typability, $\Gamma, \vec{u} : \vec{T} \vdash x = o : C$ for some C such that $\text{local}(C)$. Hence by typability $\Gamma, \vec{u} : \vec{T} \vdash o : C'$ such that $C' <: C$. By our convention, it must be the case that $\text{local}(C')$ also. Then this case straightforward as above, by establishing that o must be in the store due to the locality of its class.
 - v. Suppose the last rule applied was FLDASS. Similar to the case for ASS.
 - vi. The last rule applied was LEAVESANDBOX. Then $P_{im} \equiv E[\text{sandbox } \{o\}.f = e] \mid Q_{im}$. Then this is also straightforward, noting that $\Gamma, \vec{u} : \vec{T} \vdash o : C$ with $\text{local}(C)$.
- Inv(10) Straightforward by the definition of $\Delta_1 \asymp \Delta_2$.
- Inv(11) Straightforward by the definition of $\Delta_1 \asymp \Delta_2$.
- Inv(12) We investigate the cases where a value comes into a redex position. Assume $P_{im+1} \equiv E[v] \mid Q_{im+1}$, then we perform a case analysis as follows.
- (a) The last rule was NEW. Then this case is trivial.
 - (b) The last rule was VAR. Then $P_{im} \equiv E[x] \mid Q_{im}$, and this case is straightforward by Inv(13).
 - (c) The last rule was FLD. Then $P_{im} \equiv E[o.f] \mid Q_{im}$, and this case is straightforward by Inv(14).
 - (d) The last rule was FREEZE. Then $P_{im} \equiv E[\text{freeze}[t](T x)\{e\}] \mid Q_{im}$. By reduction, $v = \lambda(T x).(\nu \vec{u})(l_i, e, \sigma, \text{CT})$. Examining the definition of fv , we see that $\text{fv}(v) = (\text{fv}(e) \setminus \{x\}) \cup \text{fv}(\sigma) \cup \text{fv}(\text{CT})$. By premises of FREEZE, we see that σ is derived from σ_{im} . Hence by Inv(13) and Inv(14), all values must be closed, so $\text{fv}(\sigma) = \emptyset$. Similarly, CT is derived from CT_{im} and so since $\vdash \text{CT}_{im} : \text{ok}$, we have that $\text{fv}(\text{CT}) = \emptyset$. By typability of P_{im} , we see that $\text{fv}(e) \subseteq \{x\}$, so we can conclude that $\text{fv}(v) = \emptyset$ as required.
 - (e) The last rule was DNOTHING. Then $P_{im} \equiv E[\text{download } \vec{C} \text{ from } l_j \text{ in sandbox } \{v\}] \mid Q_{im}$. However this situation only arises after a frozen value has been defrosted. Therefore by the inductive hypothesis, we know that v can contain no free variables.
- Inv(13) For this invariant, we check the cases where new variable mappings are added to the store, or when an existing mapping is changed. Assume $\sigma_{im+1}(x) = v$. Then
- (a) Suppose the last reduction rule applied was DEC. Then $P_{im} \equiv E[T x = v] \mid Q_{im}$. Then by Inv(12) we have that $\text{fv}(v) = \emptyset$, so this case is straightforward.
 - (b) The last rule applied was DEFROST. Then $P_{im} \equiv E[\text{defrost}(v; \lambda(T x).(\nu \vec{u})(l_j, e, \sigma, \text{CT}))] \mid Q_{im}$. However by Inv(12), we

have that the frozen value contains no free variables, hence $\text{fv}(\sigma) = \emptyset$ and so any mappings added to σ_{im} are closed.

(c) The last rule was ASS. Again straightforward by Inv(12).

Inv(14) For this invariant, check the cases where new object mappings are added to the store, or when an existing mapping is changed. Assuming $\sigma_{im+1}(o) = (C, \vec{f} : \vec{v})$, we investigate when the last rule was NEW, DEFROST or FLDASS. All are straightforward by application of Inv(12).

Inv(15) The only interesting case is when the last applied reduction rule was FREEZE. We show only the case when $t = \mathbf{eager}$, as the others are similar. Suppose:

$$P_{im} \equiv E[\mathbf{freeze}[\mathbf{eager}](T\ x)\{e\}] \mid Q_{im} \longrightarrow_{l_i} P_{im+1} \equiv E[\lambda(T\ x).(\nu \vec{u})(l, e, \sigma, \mathbf{CT})] \mid Q_{im}$$

Then by the premises of FREEZE we have:

$$\begin{aligned} \{\vec{y}\} &= \text{fv}(e) \setminus \{x\} \text{ with } \sigma_y = \bigcup \sigma_{im}(y_i), \quad \mathbf{CT}' = \mathbf{cg}(\mathbf{CT}, \text{fcl}(e) \cup \text{fcl}(\sigma)) \\ \sigma &= \mathbf{og}(\sigma_{im}, \text{fn}(e) \cup \text{fn}(\sigma_y)) \cup \sigma_y \text{ with } \{\vec{u}\} = \text{dom}(\sigma) \end{aligned}$$

As a preliminary note, by typability of N_m we have that $\vdash \mathbf{CT} : \mathbf{ok}$ and so $\text{fn}(\mathbf{CT}) = \emptyset$. Therefore: $\text{fn}(\lambda(T\ x).(\nu \vec{u})(l, e, \sigma, \mathbf{CT})) = (\text{fn}(e) \cup \text{fn}(\sigma)) \setminus \{\vec{u}\}$. By Lemma 1 (1), and that $\{\vec{u}\} = \text{dom}(\sigma)$:

$$\forall u \in (\text{fn}(\mathbf{og}(\sigma_{im}, \text{fn}(e) \cup \text{fn}(\sigma_y))) \setminus \{\vec{u}\}). \Gamma \vdash u : C \wedge \mathbf{remote}(C)$$

By Lemma 2 (2), all *local* object identifiers reachable from e or σ_y must be included in the computed object graph and hence in $\vec{u}$. Therefore all remaining free names in e or σ_y must point to remote objects. So:

$$\begin{aligned} &\forall u \in (\text{fn}(e) \cup \text{fn}(\sigma_y)) \setminus \{\vec{u}\}. \Gamma \vdash u : C \wedge \mathbf{remote}(C) \wedge (9) \\ \implies &\forall u \in (\text{fn}(e) \cup \text{fn}(\sigma)) \setminus \{\vec{u}\}. \Gamma \vdash u : C \wedge \mathbf{remote}(C) \\ \implies &\forall u \in \text{fn}(\lambda(T\ x).(\nu \vec{u})(l, e, \sigma, \mathbf{CT})). \Gamma \vdash u : C \wedge \mathbf{remote}(C) \end{aligned}$$

Inv(16) Suppose $P_{im+1} \equiv E[\mathbf{ready}\ o\ n] \mid Q_{im+1}$. There are only two interesting cases: the last reduction rule applied was NOTIFY, or it was NOTIFYALL.

(a) Suppose $P_{im} \equiv E'[o.\mathbf{notify}] \mid E[\mathbf{waiting}(c)\ n] \mid Q_{im}$. By typability of this term we have that $n > 0$, and by the premises of NOTIFY we see that $c \in \mathbf{blocked}(\sigma_{im}, o)$. Then by Lemma 6.4 we have that $\mathbf{insync}(o, E)$ with n levels of nesting.

(b) Suppose $P_{im} \equiv E'[o.\mathbf{notifyAll}] \mid E[\mathbf{waiting}(c)\ n] \mid Q_{im}$. Then this case follows in the same way as the previous.

Inv(17) Assume $P_{im+1} \equiv E[\mathbf{waiting}(c)\ n] \mid Q_{im+1}$. Then there is only one interesting case to consider, when $P_{im} \equiv E[o.\mathbf{wait}] \mid Q_{im}$. By premise of WAIT we have that $\mathbf{insync}(o, E)$, and consequently by Lemma 6.4, $n > 0$. Since channel c is created fresh, we know that it is stored in the blocked queue of at most one object, and again by the premise of WAIT we see that $\sigma_{im+1} = \mathbf{block}(\sigma_{im}, o, c)$, therefore it exists in exactly one place: the blocked queue of o . This completes the case.

E.4 Proofs for progress with synchronisation primitives

This subsection proves Lemma 6.1.

By induction on the number of threads synchronised on the object o , written n . The base case is straightforward; take $n = 1$ then $P \equiv E_1[\text{insync } o \{e_1\}] \mid Q$. e_1 can be of any form and still satisfy the condition that at most one thread can execute in its critical section.

For the inductive step, we assume that the property holds for $n - 1$ threads in parallel. Now we write P as follows:

$$P \equiv E_1[\text{insync } o \{e_1\}] \mid \dots \mid E_{n-1}[\text{insync } o \{e_{n-1}\}] \mid E_n[\text{insync } o \{e_n\}] \mid Q \quad (9)$$

We shall show that either:

$$\forall j. 1 \leq j \leq n. (e_j = E'_j[\text{waiting}(c) \ o \dots] \vee e_j = E'_j[\text{ready } o \dots]) \quad (10)$$

$$\text{or } \exists! j. 1 \leq j \leq n. (e_j \neq E'_j[\text{waiting}(c) \ o \dots] \wedge e_j \neq E'_j[\text{ready } o \dots]) \quad (11)$$

with $c \in \text{blocked}(o, \sigma)$. By the inductive assumption, we have that:

$$\forall j. 1 \leq j \leq n-1. (e_j = E'_j[\text{waiting}(c) \ o \dots] \vee e_j = E'_j[\text{ready } o \dots]) \quad (12)$$

$$\text{or } \exists! j. 1 \leq j \leq n-1. (e_j \neq E'_j[\text{waiting}(c) \ o \dots] \wedge e_j \neq E'_j[\text{ready } o \dots]) \quad (13)$$

For (12) if $e_n = \text{waiting}(c) \ o \dots$ or $e_n = \text{ready } o \dots$ then we can immediately conclude (10) as required. Similarly, if e_n is not of this form then we can safely conclude (11).

However, if we have the situation (13) then the nature of the new thread is important—it cannot be executing inside its critical section. If e_n is waiting or ready, then (11) is preserved. However consider that e_n is executing within its critical section. We shall show that this situation cannot arise by showing a contradiction.

Without loss of generality, consider only two threads executing in their critical section simultaneously:

$$E_1[\text{insync } o \{e_1\}] \mid E_2[\text{insync } o \{e_2\}]$$

Assume neither e_1 nor e_2 are of the form $E''[\text{waiting}(c) \ o \dots]$ or $E''[\text{ready } o \dots]$. In order to reach such a situation, one thread must have entered its critical section while the other was still active in its. Therefore:

$$E_1[\text{insync } o \{e_1\}] \mid E_2[e'_2], \sigma, \text{CT} \longrightarrow_l E_1[\text{insync } o \{e_1\}] \mid E_2[\text{insync } o \{e_2\}], \sigma', \text{CT}$$

e'_2 can take two shapes: $e'_2 = \text{ready } o \dots$ or $e'_2 = \text{sync } (o) \{e_2\}$. In the first case, the only reduction rule applicable is **READY**, which has $\text{getLock}(\sigma, o) = 0$ as a premise. However by Lemma 6.4 we can conclude that $\text{getLock}(\sigma, o) > 0$, giving rise to an immediate contradiction. The same argument may be made for the second form of e'_2 , where **SYNC** is used.

Using this, it is possible to conclude that e_n must be of the form $E'_n[\text{waiting}(c) \ o \dots]$ or $E'_n[\text{ready } o \dots]$, which establishes (11) as required. $\square$

F. PROOFS FOR TYPE SOUNDNESS

F.1 Proofs for Substitution Lemma 6.5

By induction on the structure of the expression e using Lemma 6.2. The proof is standard, so we only list one case. Suppose $\Gamma, x : T; \Delta, c : \text{chanI}(U) \vdash E[\text{await } c]^U : \text{thread}$. Then this is derived from TT-AWAIT with the premise: $\Gamma, x : T; \Delta \vdash E[\]^U : \text{thread}$ with $c \notin \text{dom}(\Delta)$. We apply the inductive hypothesis obtaining $\Gamma; \Delta \vdash E[\]^U[v/x] : \text{thread}$. Since Δ is unchanged, the side condition $c \notin \text{dom}(\Delta)$ still holds, and we can apply rule TT-AWAIT to yield $\Gamma; \Delta, c : \text{chanI}(U) \vdash E[\text{await } c]^U[v/x] : \text{thread}$, as required. $\square$

F.2 Proof of Theorem 6.1

(1) The proof proceeds by induction on the length of reduction sequence with a case analysis on the final reduction rule applied. When $\sigma = \sigma'$ or $\text{CT} = \text{CT}'$ we shall omit them.

Case VAR. Use Lemma D.2 (3).

Case COND. Straightforward by the inductive hypothesis.

Case WHILE. Standard.

Case FLD. Straightforward by Lemma D.2 (6).

Case SEQ. Assume $e_1; e_2, \sigma, \text{CT} \longrightarrow_l (\nu \vec{u})(e_2, \sigma', \text{CT}')$ with $\vec{u} \notin \text{fnv}(e_2)$ and $\Gamma \vdash e_1; e_2 : S$. To derive this, TE-SEQ was applied with the premises that $\Gamma \vdash e_1 : U$ and $\Gamma \vdash e_2 : S$. By premises of SEQ we have $e_1, \sigma, \text{CT} \longrightarrow_l v, \sigma', \text{CT}'$, and so by the inductive hypothesis we have $\Gamma, \vec{u} : \vec{T} \vdash v : U'$ with $U' < U$ and $\Gamma, \vec{u} : \vec{T}; \Delta \vdash \sigma' : \text{ok}$, $\vdash \text{CT}' : \text{ok}$. Then to complete the case we apply weakening to derive $\Gamma, \vec{u} : \vec{T} \vdash e_2 : S$.

Case ASS. Straightforward using Lemma D.2 (2).

Case FLDASS. Use Lemma D.2 (5).

Case NEW. Assume $\text{new } C(\vec{v}), \sigma \longrightarrow_l (\nu o)(o, \sigma \cdot [o \mapsto (C, \vec{f} : \vec{v}, 0, \emptyset)])$ and $\Gamma \vdash \text{new } C(\vec{e}) : C$. To derive this, TE-NEW must have been applied with premises $\text{fields}(C) = \vec{T}\vec{f}, T_i' < T_i, \Gamma \vdash e_i : T_i'$ and $\vdash C : \text{tp}$. Using this we can derive that $\Gamma \vdash (C, \vec{f} : \vec{v}, 0, \emptyset) : \text{ok}$. Since o is fresh and therefore not in Γ , we can apply Lemma D.2 (4) to complete the case.

Case NEWR. Similar to the case for NEW. Assume $\Gamma \vdash \text{new } C^l(\vec{e}) : C$, we see TE-NEW was applied with the important premise that $\vdash C : \text{tp}$. Then we can apply TE-CLASSLOAD to derive $\Gamma \vdash \text{download } C \text{ from } l \text{ in } \text{new } C^l(\vec{e}) : C$ as required.

Case DEC. Straightforward by Lemma D.2 (1).

Case CONG. Use Lemma 6.5(3).

Case RESOLVE. By TE-CLASSLOAD.

Case DNOTHING. Assume By TE-CLASSLOAD.

Case FREEZE. We shall assume the **eager**-mode of operation, others are similar. Assume $\text{freeze}[\text{eager}](T \ x)\{e\}, \sigma, \text{CT} \longrightarrow_l \lambda(T \ x).(\nu \vec{u})(l, e, \sigma', \text{CT}'), \sigma, \text{CT}$ and $\Gamma \vdash$

freeze[eager]($T\ x\}\{e\} : T \rightarrow U$. To infer this, TE-FREEZE was used with premise $\Gamma, x : T \vdash e : U$. The premises of FREEZE are

$$\sigma' = \text{og}(\sigma, \text{fn}(e) \cup \text{fn}(\sigma_y)) \cup \sigma_y, \quad \text{with } \{\vec{y}\} = \text{fv}(e) \setminus \{x\}, \sigma_y = \bigcup \sigma(y_i) \quad (\text{a})$$

$$\{\vec{u}\} = \text{dom}(\sigma') \quad (\text{b})$$

$$\text{CT}' = \text{cg}(\text{CT}, \text{fcl}(e) \cup \text{fcl}(\sigma')) \quad (\text{c})$$

From (a) we can apply Lemma 4.1 (1) to obtain

$$\Gamma; \emptyset \vdash \text{og}(\sigma, \text{fn}(e) \cup \text{fn}(\sigma_y)) : \text{ok}$$

By definition, we have $\sigma_y \subseteq \sigma$. Then by Lemma D.2 (8), $\Gamma; \Delta \vdash \sigma_y : \text{ok}$. However σ_y , by construction, only contains mappings from variables, i.e. there are no store objects (and hence no channels) in its co-domain, therefore $\Gamma; \emptyset \vdash \sigma_y : \text{ok}$ by Lemma D.1 (strengthening). From this knowledge, we can apply Lemma D.2 (7) to obtain: $\Gamma; \emptyset \vdash \sigma' : \text{ok}$.

Now considering CT' , by (c) we see that CT' is constructed using the class dependency algorithm. Trivially if $C \in \text{fcl}(e) \cup \text{fcl}(\sigma')$ then $C \in \text{dom}(\text{CSig})$. So we apply Lemma 4.1 (3) to obtain $\vdash \text{CT}' : \text{ok}$.

Applying TV-FROZEN to e , σ' and CT' , we derive $\Gamma' \vdash \lambda(T\ x).(\nu \vec{u})(l, e, \sigma', \text{CT}') : T \rightarrow U$ where Γ' is a subset of Γ such that $u_i \notin \text{dom}(\Gamma')$. Then we apply Lemma D.1 (weakening) to obtain $\Gamma \vdash \lambda(T\ x).(\nu \vec{u})(l, e, \sigma', \text{CT}') : T \rightarrow U$ as required. Since σ and CT are unchanged, this completes the case.

Case DEFROST. Assume

$$\text{defrost}(v; \lambda(T\ x).(\nu \vec{u})(m, e, \sigma', \text{CT}')), \sigma, \text{CT} \quad (\text{b})$$

$$\longrightarrow_l (\nu x \vec{u})(\text{download } \vec{F} \text{ from } m \text{ in sandbox } \{e[\vec{C}^m/\vec{C}]\}, \sigma \cup \sigma' \cdot [x \mapsto v], \text{CT} \cup \text{CT}')$$

and $\Gamma \vdash \text{defrost}(v; \lambda(T\ x).(\nu \vec{u})(m, e, \sigma', \text{CT}')) : U$. To derive this, the last typing rule applied was TE-DEFROST with premises

$$\Gamma \vdash v : T' \text{ with } T' <: T \text{ and } \Gamma \vdash \lambda(T\ x).(\nu \vec{u})(m, e, \sigma', \text{CT}') : T \rightarrow U \quad (\text{a})$$

We shall prove

$$\Gamma, x : T, \vec{u} : \vec{T} \vdash \text{download } \vec{F} \text{ from } m \text{ in sandbox } \{e[\vec{C}^m/\vec{C}]\} : U \quad (\text{b})$$

$$\Gamma, x : T, \vec{u} : \Delta \vdash \sigma \cup \sigma' \cdot [x \mapsto v] : \text{ok} \quad (\text{c})$$

$$\vdash \text{CT} \cup \text{CT}' : \text{ok} \quad (\text{d})$$

To infer (a), TV-FROZEN was applied with the premises

$$\Gamma, x : T, \vec{u} : \vec{T} \vdash e : U \quad \Gamma, \vec{u} : \vec{T} \vdash \sigma' : \text{ok} \quad \vdash \text{CT}' : \text{ok}$$

Since C^m and C are typed by the same rule, we infer that

$$\Gamma, x : T, \vec{u} : \vec{T} \vdash e[\vec{C}^m/\vec{C}] : U$$

By application of TE-SANDBOX we have $\Gamma, x : T, \vec{u} : \vec{T} \vdash \text{sandbox } \{e[\vec{C}^m/\vec{C}]\} : U$.

By the premise of DEFROST, $\{\vec{F}\} = \text{fcl}(\sigma') \setminus \text{dom}(\text{CT}')$, and in order to judge σ' , it must be the case that for all F_i , there exists some mapping $o : F_i$ in $\Gamma, \vec{u} : \vec{T}$. Then by Lemma D.1 we have that $\Gamma, \vec{u} : \vec{T} \vdash \text{Env}$. To infer this, we must have used

the rules for well-formedness of environments (Figure 13), and so can deduce that $\vdash F_i : \mathbf{tp}$ for all F_i . Taking this fact, we can apply TE-CLASSLOAD to obtain (b). Now to show that the two stores are compatible, we apply Lemma D.2 (7) to derive $\Gamma, \vec{u} : \vec{T} \vdash \sigma \cup \sigma' : \mathbf{ok}$. Then by application of Lemma D.2 (1) we have (c). (d) then follows from $\mathbf{Inv}(3)$ to complete this case.

(2) Assume $\Gamma; \Delta \vdash F : \mathbf{conf}$, $F \longrightarrow_l F'$ and $F' \not\models \mathbf{Err}$. Then we have $\Gamma; \Delta \vdash F' : \mathbf{conf}$. To type a configuration, we apply TC-CONF which has the premises $\Gamma; \Delta_1 \vdash P : \mathbf{thread}$, $\Gamma; \Delta_2 \vdash \sigma : \mathbf{ok}$, $\vdash \mathbf{CT} : \mathbf{ok}$, $\mathbf{FCT} \subseteq \mathbf{CT}$, $\Delta_1 \asymp \Delta_2$ where $\Delta = \Delta_1 \odot \Delta_2$. Proofs proceed from this point, and we omit the store σ and class table $\mathbf{CT}$ when they do not change.

Case LEAVESANDBOX. Straightforward by the inductive hypothesis.

Case METHLOCAL. Assume $E[o.m(v)] \mid P \longrightarrow_l (\nu c)(E[\mathbf{await } c] \mid o.m(v) \mathbf{with } c \mid P)$ and $\Gamma; \Delta_1 \vdash E[o.m(v)] \mid P : \mathbf{thread}$. To type this, rule TT-PAR was applied with the premises

$$\Gamma; \Delta_{11} \vdash E[o.m(v)] : \mathbf{thread} \quad \Gamma; \Delta_{12} \vdash P : \mathbf{thread} \quad \text{with } \Delta_{11} \asymp \Delta_{12} \quad (\text{a})$$

To type (a), we applied Lemma 6.5(3) with premises

$$\Gamma \vdash o.m(v) : U \quad \Gamma; \Delta_{11} \vdash E[\]^U : \mathbf{thread} \quad (\text{b})$$

Pick a fresh channel c , then apply TT-METHWITH to (a) to obtain

$$\Gamma; \mathbf{chan0}(U) \vdash o.m(v) \mathbf{with } c : \mathbf{thread}$$

With the same fresh channel, apply TT-AWAIT to (b) to obtain

$$\Gamma; \Delta_{11}, c : \mathbf{chanI}(U) \vdash E[\mathbf{await } c] : \mathbf{thread}$$

By Definition 5.1, we have $\Delta_{11}, c : \mathbf{chanI}(U) \asymp c : \mathbf{chan0}(U)$ with

$\Delta_{11}, c : \mathbf{chanI}(U) \odot c : \mathbf{chan0}(U) = \Delta_{11}, c : \mathbf{chan}$. Since c was chosen fresh, $c \notin \mathbf{dom}(\Delta_{12})$ therefore $\Delta_{11}, c : \mathbf{chan} \asymp \Delta_{12}$ and we can apply TT-PAR twice to get

$$\Gamma; \Delta_{11}, c : \mathbf{chan} \odot \Delta_{12} \vdash E[\mathbf{await } c] \mid o.m(v) \mathbf{with } c \mid P : \mathbf{thread}$$

Then by permutation of the environment (Lemma D.1) we have $\Delta_{11}, c : \mathbf{chan} \odot \Delta_{12} = \Delta_1, c : \mathbf{chan}$. Applying TT-RES yields:

$$\Gamma; \Delta_1 \vdash (\nu c)(E[\mathbf{await } c] \mid o.m(v) \mathbf{with } c \mid P) : \mathbf{thread}$$

as required.

Case METHREMOTE. Assume

$$E[o.m(v)] \mid P \longrightarrow_l (\nu c)(E[\mathbf{await } c] \mid \mathbf{go } o.m(\mathbf{serialize}(v)) \mathbf{with } c \mid P)$$

with $o \notin \mathbf{dom}(\sigma)$ and $\Gamma; \Delta_1 \vdash E[o.m(v)] \mid P : \mathbf{thread}$. This case is similar to the previous case up to (b). We shall proceed from this point. To type (b), we must

have applied TE-METH, with the premise that $\Gamma \vdash o : C$. By the side condition that $o \notin \text{dom}(\sigma)$, $\text{Inv}(8)$ and $\text{Inv}(6)$, we have $\text{remote}(C)$. Therefore we can apply TT-GOSER to derive:

$$\Gamma; c : \text{chan0}(U) \vdash \text{go } o.m(\text{serialize}(v)) \text{ with } c : \text{thread}$$

Proof then proceeds from this point as in the case for METHLOCAL to derive $\Gamma; \Delta_1 \vdash (\nu c)(E[\text{await } c] \mid \text{go } o.m(\text{serialize}(v)) \text{ with } c \mid P) : \text{thread}$, as required.

Case METHINVOKE. Assume

$$o.m(v) \text{ with } c, \sigma \longrightarrow_l (\nu x)(e[o/\text{this}][\text{return}(c)/\text{return}], \sigma \cdot [x \mapsto v])$$

and $\Gamma; \Delta_1 \odot \Delta_2 \vdash o.m(v) \text{ with } c, \sigma : \text{conf}$. To infer this, we applied TC-CONF with premises

$$\Gamma; \Delta_1 \vdash o.m(v) \text{ with } c : \text{thread} \quad \Gamma; \Delta_2 \vdash \sigma : \text{ok} \quad \Delta_1 \preceq \Delta_2$$

By premises of TT-METHWITH, we have $\Gamma \vdash o.m(v) : U$ and $\Delta_1 = c : \text{chan0}(U)$. By application of METHINVOKE in the reduction step, and the fact that to infer the above, we had to apply TE-METH we have

$$\begin{aligned} \sigma(o) &= (C, \dots) & \text{mbody}(m, C, \text{CT}) &= (x, e) & \text{mtype}(m, C) &= D \rightarrow U \\ \Gamma \vdash o : C & & \Gamma \vdash v : D' & & D' <: D \end{aligned} \quad (\text{a})$$

By assumption that $\vdash \text{CT} : \text{ok}$, we can apply Lemma D.3 to (a) to obtain that $x : D, \text{this} : C \vdash e : \text{ret}(U')$ with $U' <: U$ for a freshly chosen x . By application of Lemma 6.5 (substitution) followed by Lemma D.1 (strengthening) we have $\Gamma, x : D \vdash e[o/\text{this}] : \text{ret}(U')$. Then by applying TT-RETURN we have

$$\Gamma, x : D; c : \text{chan0}(U) \vdash e[o/\text{this}][\text{return}(c)/\text{return}] : \text{thread}$$

By application of Lemma D.2 (1) we then have that $\Gamma, x : D \vdash \sigma \cdot [x \mapsto v] : \text{ok}$. To complete the case we then apply TC-CONF followed by TC-RESID giving

$$\Gamma; \Delta_1 \odot \Delta_2 \vdash (\nu x)(e[o/\text{this}][\text{return}(c)/\text{return}], \sigma \cdot [x \mapsto v]) : \text{conf}$$

as required.

Case AWAIT. Assume $E[\text{await } c] \mid \text{return}(c) v \longrightarrow_l E[v]$ and $\Gamma; \Delta_1 \vdash E[\text{await } c] \mid \text{return}(c) v : \text{thread}$. To type this, we applied TT-PAR with premises:

$$\Gamma; \Delta_{11} \vdash E[\text{await } c] : \text{thread} \quad \Gamma; \Delta_{12} \vdash \text{return}(c) v : \text{thread} \quad \Delta_{11} \preceq \Delta_{12} \quad (\text{a})$$

To type the first conjunct of (a), we must have applied TT-AWAIT. To type the second conjunct, we applied TT-RETURN. These give us that

$$\begin{aligned} \Gamma; \Delta'_{11} \vdash E[\]^U : \text{thread} & \quad \text{with } \Delta_{11} = \Delta'_{11}, c : \text{chanI}(U) \\ \Gamma \vdash \text{return } v : \text{ret}(U') & \quad \text{with } U' <: U \end{aligned} \quad (\text{b})$$

To derive (b), we applied TE-RETURN with the premise that $\Gamma \vdash v : U'$. By Lemma 6.5(3) we obtain $\Gamma; \Delta'_{11} \vdash E[v] : \mathbf{thread}$. To complete the case we apply Lemma D.1 (weakening) to the environment Δ'_{11} to give $\Gamma; \Delta_1 \vdash E[v] : \mathbf{thread}$.

Case FORK. Similar to the above case, using Lemma 6.5(3).

Case THREADDEATH. Trivial.

Case SYNC. Assume $E[\mathbf{sync}(o) \{e\}], \sigma \longrightarrow_l E[\mathbf{insync} o \{e\}], \sigma'$ with $\Gamma; \Delta_1 \vdash E[\mathbf{sync}(o) \{e\}] : \mathbf{thread}$ and $\Gamma; \Delta_2 \vdash \sigma : \mathbf{ok}$.

To derive this, we applied Lemma 6.5(3) with premises

$$\Gamma; \Delta_1 \vdash E[\]^S \quad \Gamma \vdash \mathbf{sync}(o) \{e\} : S \quad (\text{a})$$

To derive (a), we applied TE-SYNC with premises $\Gamma \vdash o : C$ and $\Gamma \vdash e : S$. We can then apply TE-INSYNC to derive $\Gamma \vdash \mathbf{insync} o \{e\} : S$. Showing the resulting thread is well-typed, $\Gamma; \Delta_1 \vdash E[\mathbf{insync} o \{e\}] : \mathbf{thread}$, is straightforward from this point, so all that remains is to show $\Gamma; \Delta_2 \vdash \sigma' : \mathbf{ok}$. However this is straightforward, since in the reduction step only the lock count of a store entry is changed. As σ holds, we see that it must be the case that lock counts cannot be set to a negative number, and so we can conclude the case.

Case WAIT. Assume $E[o.\mathbf{wait}] \mid P, \sigma \longrightarrow_l (\nu c)(E[\mathbf{waiting}(c) n] \mid P, \sigma')$ and $\Gamma; \Delta_1 \vdash E[o.\mathbf{wait}] \mid P : \mathbf{thread}$, $\Delta_1 \asymp \Delta_2$ and $\Gamma; \Delta_2 \vdash \sigma : \mathbf{ok}$. To derive this, TT-PAR was applied with premises:

$$\Gamma; \Delta_{11} \vdash E[o.\mathbf{wait}] : \mathbf{thread} \quad \Gamma; \Delta_{12} \vdash P : \mathbf{thread} \quad \text{and} \quad \Delta_{11} \asymp \Delta_{12} \quad (\text{a})$$

To derive the left conjunct of (a), we applied Lemma 6.5(3) with the premise:

$$\Gamma; \Delta_{11} \vdash E[\]^{\mathbf{void}} : \mathbf{thread} \quad \text{and} \quad \Gamma \vdash o.\mathbf{wait} : \mathbf{void}$$

By the premise of the reduction rule WAIT, we have

$$\mathbf{insync}(o, E) \quad \mathbf{getLock}(\sigma, o) = n \quad \mathbf{setLock}(\sigma, o, 0) = \sigma'' \quad \mathbf{block}(\sigma'', o, c) = \sigma'$$

By Lemma 6.4, we have that $n > 0$, and since c is fresh we can apply TT-WAITING to obtain

$$\Gamma; \Delta_{11}, c : \mathbf{chanI}(\mathbf{void}) \vdash E[\mathbf{waiting}(c) n]^{\mathbf{void}} : \mathbf{thread} \quad (\text{h})$$

Since c is fresh, then by Definition 5.1, $\Delta_{11}, c : \mathbf{chanI}(\mathbf{void}) \asymp \Delta_{12}$. We then apply TT-PAR to yield $\Gamma; \Delta_1, c : \mathbf{chanI}(\mathbf{void}) \vdash E[\mathbf{waiting}(c) n] \mid P : \mathbf{thread}$. By Lemma D.2 (9), $\Gamma; \Delta_2, c : \mathbf{chan0}(\mathbf{void}) \vdash \sigma' : \mathbf{ok}$. By Definition 5.1 followed by TC-RESC we can derive $\Gamma; \Delta_1 \odot \Delta_2 \vdash (\nu c)(E[\mathbf{waiting}(c) n] \mid P, \sigma') : \mathbf{conf}$, as required.

Case NOTIFY. Assume

$$\begin{aligned} & E[o.\mathbf{notify}] \mid E_1[\mathbf{waiting}(c) n], \sigma \longrightarrow_l E[\epsilon] \mid E_1[\mathbf{ready} o n], \sigma' \\ & \Gamma; \Delta_1 \odot \Delta_2 \vdash E[o.\mathbf{notify}] \mid E_1[\mathbf{waiting}(c) n], \sigma : \mathbf{conf} \quad \Delta_1 \asymp \Delta_2 \end{aligned}$$

To derive the above, we applied TC-CONF with premises

$$\begin{aligned} \Gamma; \Delta_{11} \odot \Delta_{12} \vdash E[o.\text{notify}] \mid E_1[\text{waiting}(c) \ n] : \text{thread} \quad \Delta_1 = \Delta_{11} \odot \Delta_{12} \quad (a) \\ \Gamma; \Delta_2 \vdash \sigma : \text{ok} \quad (b) \end{aligned}$$

To derive (a) we applied TT-PAR with premises

$$(i) \Gamma; \Delta_{11} \vdash E[o.\text{notify}] : \text{thread} \quad (ii) \Gamma; \Delta_{12} \vdash E_1[\text{waiting}(c) \ n] : \text{thread} \quad (c)$$

To derive (c-i) there were two possible rules applied. Either E contains a return statement, or E is a forked thread. We show the case of the former (the latter is similar). Here the typing rule applied was TT-RETURN and we have that $\Delta_{11} = d : \text{chan0}(U)$ with the premise that $\Gamma \vdash E'[o.\text{notify}] : \text{ret}(U')$ with $U' <: U$ (E' is the context prior to the substitution of return statements). Then we see that we applied Lemma 6.5(3) with the premise that $\Gamma \vdash o.\text{notify} : \text{void}$ and $\Gamma \vdash E'[\]^{\text{void}} : \text{ret}(U')$. Then, to derive $\Gamma \vdash o.\text{notify} : C$ we applied TE-MONITOR with premise $\Gamma \vdash o : C$. To derive (c-ii) we applied TT-WAITING with premises:

$$\Gamma; \Delta'_{12} \vdash E_1[\]^{\text{void}} : \text{thread}, \quad c \notin \text{dom}(\Delta'_{12}) \quad n > 0 \quad \Delta_{12} = \Delta'_{12}, c : \text{chanI}(\text{void}) \quad (d)$$

Since $\Gamma \vdash \epsilon : \text{void}$, we can safely conclude that $\Gamma; \Delta_{11} \vdash E[\epsilon] : \text{ret}(U')$. Then since $\Gamma \vdash o : C$ and $n > 0$ we can apply TE-READY to deduce $\Gamma \vdash \text{ready } o \ n : \text{void}$. Then taking this fact and (d), we can fill the whole in context E_1 to obtain

$$\Gamma; \Delta'_{12} \vdash E_1[\text{ready } o \ n] : \text{thread}$$

Now by the premise of the reduction rule NOTIFY, we have that $c \in \text{blocked}(\sigma, o)$, therefore by typability of (b), we have that $c : \text{chan0}(\text{void}) \in \Delta_2$. Since $\Delta_1 \asymp \Delta_2$ we cannot have another output on channel c in Δ_{11} , therefore we can safely say that $\Delta_{11} \asymp \Delta'_{12}$, and then apply TT-PAR to obtain

$$\Gamma; \Delta_{11} \odot \Delta'_{12} \vdash E[\epsilon] \mid E_1[\text{ready } o \ n] : \text{thread}$$

Now we must show that the new store, σ' is safe. Taking $\Delta_2 = \Delta'_2, c : \text{chan0}(\text{void})$ we have by Lemma D.1 that $\Gamma; \Delta_2 \vdash \text{Env}$, and so $c \notin \text{dom}(\Delta'_2)$. By premise of the reduction rule, we have that $\sigma' = \text{unblock}(\sigma, o, c)$, and so by applying Lemma D.2 (10) it must be the case that $\Gamma; \Delta'_2 \vdash \sigma' : \text{ok}$. Trivially we have $(\Delta_{11} \odot \Delta'_{12}) \asymp \Delta'_2$, and can apply TC-CONF to yield (where $\Delta'_1 = \Delta'_{11} \odot \Delta'_{12}$)

$$\Gamma; \Delta'_1 \odot \Delta'_2 \vdash E[\epsilon] \mid E_1[\text{ready } o \ n], \sigma' : \text{conf}$$

Finally we apply TC-WEAK to add the channel c that was removed to obtain $\Gamma; \Delta_1 \odot \Delta_2 \vdash E[\epsilon] \mid E_1[\text{ready } o \ n], \sigma' : \text{conf}$.

Case NOTIFYALL. Similar to the case for NOTIFY.

Case NOTIFYNONE. Straightforward.

Case READY. Assume [this should be in expressions](#) $\text{ready } o \ n, \sigma \longrightarrow_l \epsilon, \sigma'$ with $\Gamma \vdash \text{ready } o \ n : \text{void}$ and $\Gamma; \Delta_2 \vdash \sigma : \text{ok}$. Trivially, $\Gamma \vdash \epsilon : \text{ok}$, therefore it remains to show that $\Gamma; \Delta_2 \vdash \sigma' : \text{ok}$. However since to derive $\text{ready } o \ n$, we had to apply TE-READY which has the premise that $n > 0$, then we can trivially conclude that $\Gamma; \Delta_2 \vdash \sigma' : \text{ok}$.

Case LEAVECRITICAL. We shall consider the case for return, as the other case is similar. Assume $\text{insync } o \ \{\text{return}(c) \ v\}, \sigma \longrightarrow_l \text{return}(c) \ v, \sigma'$ with $\Gamma; c, \text{chan0}(U) \vdash \text{insync } o \ \{\text{return}(c) \ v\} : \text{thread}$ and $\Gamma; \Delta_2 \vdash \sigma : \text{ok}$. To derive the first judgement, we applied TT-RETURN with the premise that $\Gamma \vdash \text{insync } o \ \{\text{return } v\} : \text{ret}(U')$ with $U' <: U$. Therefore we conclude TE-INSYNC was applied with premises $\Gamma \vdash o : C$ and $\Gamma \vdash \text{return } v : \text{ret}(U')$. Then by applying TT-RETURN we deduce $\Gamma; \text{chan0}(U) \vdash \text{return}(c) \ v : \text{thread}$, as required. By the premise of LEAVECRITICAL we have that $\text{getLock}(\sigma, o) = n$ and $\text{setLock}(\sigma, o, n - 1) = \sigma'$, and by the assumption of σ , we have that $n \geq 0$. By Lemma 6.4 we have that $n \neq 0$ i.e. $n > 0$. When creating σ' , we know that the new lock count can not be negative, therefore we have $\Gamma; \Delta_2 \vdash \sigma' : \text{ok}$ as required.

Case RC-PAR. Assume $P_1 \mid P_2, \sigma, \text{CT} \longrightarrow_l (\nu \vec{u})(P'_1 \mid P_2, \sigma', \text{CT}')$ with $\vec{u} \notin \text{fnv}(P_2)$ and $\Gamma; \Delta_1 \vdash P_1 \mid P_2 : \text{thread}$ with $\Delta_1 \asymp \Delta_2$, $\Gamma; \Delta_2 \vdash \sigma : \text{ok}$ and $\vdash \text{CT} : \text{ok}$. To derive $P_1 \mid P_2$, TT-PAR was applied with premises

$$\Gamma; \Delta_{11} \vdash P_1 : \text{thread}, \Gamma; \Delta_{12} \vdash P_2 : \text{thread}, \Delta_{11} \asymp \Delta_{12} \Delta_{11} \odot \Delta_{12} = \Delta_1 \quad (\text{a})$$

By the premise of RC-PAR we have that $P_1, \sigma, \text{CT} \longrightarrow_l (\nu \vec{u})(P'_1, \sigma', \text{CT}')$. Since $\Delta_{11} \asymp \Delta_{12}$, $(\Delta_{11} \odot \Delta_{12}) \asymp \Delta_2$ then by Lemma D.4, $\Delta_{11} \asymp \Delta_2$. So applying the inductive hypothesis to P_1 we obtain $\Gamma; \Delta_{11} \odot \Delta_2 \vdash (\nu \vec{u})(P'_1, \sigma', \text{CT}') : \text{conf}$. By virtue of the fact that $\vec{u} \notin \text{fnv}(P_2)$, we can apply weakening to (a), ensuring that thread P_2 is well typed in the new environment. This allows us to conclude $\Gamma; \Delta_1 \odot \Delta_2 \vdash (\nu \vec{u})(P'_1 \mid P_2, \sigma', \text{CT}') : \text{conf}$, as required.

Case RC-STR. Straightforward by Lemma 6.2.

Case RC-RES. We shall prove the case when u is a channel name. The others are similar. Assume $(\nu c\vec{u})(P, \sigma, \text{CT}) \longrightarrow_l (\nu c\vec{u}')(P', \sigma', \text{CT}')$ and $\Gamma; \Delta \vdash (\nu c\vec{u})(P, \sigma, \text{CT}) : \text{conf}$. There are two cases: the last applied typing rule was TC-WEAK, or it was TC-RESC. In the case of the former, this rule has the premises

$\Gamma; \Delta' \vdash (\nu \vec{u})(P, \sigma, \text{CT}) : \text{conf}$ with $\Delta = \Delta', c : \text{chan}$. Then by the inductive hypothesis, $\Gamma; \Delta' \vdash (\nu \vec{u}')(P', \sigma', \text{CT}') : \text{conf}$. Applying TC-WEAK we have $\Gamma; \Delta \vdash (\nu c\vec{u}')(P', \sigma', \text{CT}') : \text{conf}$ as required. When the last reduction rule was TC-RESC, we have the premises: $\Gamma; \Delta, c : \text{chan} \vdash (\nu \vec{u})(P, \sigma, \text{CT}) : \text{conf}$. Then again by the inductive hypothesis, $\Gamma; \Delta, c : \text{chan} \vdash (\nu \vec{u})(P, \sigma, \text{CT}) : \text{conf}$ and we can apply TC-RESC to conclude the required result.

(2)

Case RN-CONF. By the premises of RN-CONF, $F \longrightarrow_l F'$. From the structure of N , we see that the last typing rule applied must have been TN-CONF with premise $\Gamma; \Delta \vdash F : \text{conf}$. Given this and the assumption that $F \longrightarrow_l F'$ we can apply Theorem 6.1 (1) to obtain $\Gamma; \Delta \vdash F' : \text{conf}$. We can then re-apply TN-CONF

to deduce $\Gamma; \Delta \vdash l[F'] : \mathbf{net}$ as required.

Case DOWNLOAD. Assume $l_1[E[\mathbf{download} \vec{C} \text{ from } l_2 \text{ in } e] \mid P, \sigma_1, \mathbf{CT}_1] \mid l_2[P_2, \sigma_2, \mathbf{CT}_2] \longrightarrow l_1[E[\mathbf{resolve} \vec{D} \text{ from } l_2 \text{ in } e] \mid P, \sigma_1, \mathbf{CT}_1 \cup \mathbf{CT}'] \mid l_2[P_2, \sigma_2, \mathbf{CT}_2]$ and $\Gamma; \Delta \vdash l_1[E[\mathbf{download} \vec{C} \text{ from } l_2 \text{ in } e] \mid P, \sigma_1, \mathbf{CT}_1] \mid l_2[P_2, \sigma_2, \mathbf{CT}_2] : \mathbf{net}$. To derive this, we applied TN-PAR with premises:

$$\begin{aligned} \Gamma; \Delta_1 \vdash l_1[E[\mathbf{download} \vec{C} \text{ from } l_2 \text{ in } e] \mid P, \sigma_1, \mathbf{CT}_1] : \mathbf{net} \quad \Delta_1 \asymp \Delta_2 \quad (a) \\ \Gamma; \Delta_2 \vdash l_2[P_2, \sigma_2, \mathbf{CT}_2] : \mathbf{net} \end{aligned}$$

To type the network location in (a), we had to apply TE-CLASSLOAD at some point, with the premises that $\Gamma \vdash e : U$ and $\vdash \vec{C} : \mathbf{tp}$. Then by examination of the premises of reduction rule DOWNLOAD, we have that $\vec{D} \subseteq \vec{C}$ and so trivially $\vdash \vec{D} : \mathbf{tp}$. Thus we can apply TE-CLASSLOAD again to derive that $\Gamma \vdash \mathbf{resolve} \vec{D} \text{ from } l_2 \text{ in } e : U$ as required. Now it remains to show that $\vdash \mathbf{CT}_1 \cup \mathbf{CT}' : \mathbf{ok}$. Again by inspecting the premises of DOWNLOAD, we see that $\mathbf{CT}'$ is a subset of $\mathbf{CT}_2$ with some substitutions applied. Since these do not affect well-formedness, we deduce that $\vdash \mathbf{CT}' : \mathbf{ok}$. By Inv(3) we see that if $\mathbf{dom}(\mathbf{CT}_1) \cap \mathbf{dom}(\mathbf{CT}') \neq \emptyset$, any shared classes will have the same definition. This means we can immediately derive $\vdash \mathbf{CT}_1 \cup \mathbf{CT}' : \mathbf{ok}$. After this, to complete the case there is merely the mechanical rebuilding of the derivation of the following required result:

$$\Gamma; \Delta \vdash l_1[E[\mathbf{resolve} \vec{D} \text{ from } l_2 \text{ in } e] \mid P, \sigma_1, \mathbf{CT}_1 \cup \mathbf{CT}'] \mid l_2[P_2, \sigma_2, \mathbf{CT}_2] : \mathbf{net}$$

Case LEAVE. Assume $l_1[\mathbf{go} \text{ } o.m(v) \text{ with } c \mid P_1, \sigma_1, \mathbf{CT}_1] \mid l_2[P_2, \sigma_2, \mathbf{CT}_2] \longrightarrow l_1[P_1, \sigma_1, \mathbf{CT}_1] \mid l_2[o.m(\mathbf{deserialize}(v)) \text{ with } c \mid P_2, \sigma_2, \mathbf{CT}_2]$ and $\Gamma; \Delta \vdash l_1[\mathbf{go} \text{ } o.m(v) \text{ with } c \mid P_1, \sigma_1, \mathbf{CT}_1] \mid l_2[P_2, \sigma_2, \mathbf{CT}_2] : \mathbf{net}$. In this derivation, we had to judge $\Gamma; c : \mathbf{chan0}(U) \vdash \mathbf{go} \text{ } o.m(v) \text{ with } c : \mathbf{thread}$. This is typed by TT-DESERWITH, as is $o.m(\mathbf{deserialize}(v)) \text{ with } c$. Therefore it remains to show that the channel environments can be safely composed, however this is straightforward by Lemma D.4 and noting that the operator $\odot$ and predicate $\asymp$ are commutative. Hence we obtain

$$\Gamma; \Delta \vdash l_1[P_1, \sigma_1, \mathbf{CT}_1] \mid l_2[o.m(\mathbf{deserialize}(v)) \text{ with } c \mid P_2, \sigma_2, \mathbf{CT}_2] : \mathbf{net}$$

as required.

Case RETURN. Similar to case LEAVE.

Case SERRETURN. Assume $l[\mathbf{return}(c) \text{ } v \mid P, \sigma, \mathbf{CT}] \longrightarrow l[\mathbf{go} \text{ } \mathbf{serialize}(v) \text{ to } c \mid P, \sigma, \mathbf{CT}]$ and $\Gamma; \Delta \vdash l[\mathbf{return}(c) \text{ } v \mid P, \sigma, \mathbf{CT}] : \mathbf{net}$. To type this, we applied TN-CONF with premise $\Gamma; \Delta \vdash l[\dots] : \mathbf{conf}$. To type this, we applied, TC-CONF. This has the following premises (we omit stores and class tables, since they are invariant under this reduction and therefore trivially well-typed

$$\begin{aligned} \Gamma; \Delta_{11} \vdash \mathbf{return}(c) \text{ } v : \mathbf{thread} \quad \Delta = \Delta_1 \odot \Delta_2 \quad \Delta_1 = \Delta_{11} \odot \Delta_{12} \quad (a) \\ \Gamma; \Delta_{12} \vdash P : \mathbf{thread} \quad \Delta_{11} \asymp \Delta_{12} \end{aligned}$$

To infer (a), we applied TT-RETURN with premise

$$\Gamma \vdash \mathbf{return} \text{ } v : \mathbf{ret}(U') \quad U' <: U \quad \Delta_{11} = c : \mathbf{chan0}(U) \quad (b)$$

To type (b), we applied TE-RETURN, with the premise that $\Gamma \vdash v : U'$. Then by applying TT-VALTO we have $\Gamma, c : \text{chan0}(U) \vdash \text{go serialize}(v) \text{ to } c : \text{thread}$. To complete the case we rebuild the network by applying TC-CONF and TN-PAR and obtain $\Gamma; \Delta \vdash l[\text{go serialize}(v) \text{ to } c \mid P, \sigma, \text{CT}] : \text{net}$.

Case RN-PAR. Straightforward by the inductive hypothesis.

Case RN-RES. We consider the case when the restricted name is a channel. Assume $(\nu c)N \longrightarrow (\nu c)N'$ and $\Gamma; \Delta \vdash (\nu c)N : \text{net}$. To derive this, we applied TN-RESC with the premise $\Gamma; \Delta, c : \text{chan} \vdash N : \text{net}$. By premise of RN-RES, $N \longrightarrow N'$ and so by the inductive hypothesis we have that $\Gamma; \Delta, c : \text{chan} \vdash N' : \text{net}$. Then applying TN-RESC we obtain $\Gamma; \Delta \vdash (\nu c)N' : \text{net}$ as required.

Case RN-STR. Straightforward using Lemma 6.2.

G. PROOFS OF PROPOSITION 7.1 (2)

We define the encoding from DJ with the methods with multiple parameters into those with a single parameter. The mapping forms $\llbracket \cdot \rrbracket_{\vec{x}, z}^\Gamma$ where $\vec{x}$ is multiple parameters of the source language, while z is a single parameter of the target one. Γ is an environment.

$$\begin{aligned}
\llbracket \text{CT} \cdot [C \mapsto L] \rrbracket &\stackrel{\text{def}}{=} \llbracket \text{CT} \rrbracket \cdot [C \mapsto L'] \cup \text{CT}' \text{ where } (L', \text{CT}') = \llbracket L \rrbracket \\
\llbracket \text{CSig} \cdot C : \text{extends } D [\text{remote}] \vec{T} \vec{f} \{ \mathbf{m}_i : \vec{T}_i \rightarrow U_i \} \rrbracket &\stackrel{\text{def}}{=} \llbracket \text{CSig} \rrbracket \cdot C : \text{extends } D [\text{remote}] \vec{T} \vec{f} \{ \mathbf{m}_i : C_{m_i} \rightarrow U_i \} \cdot C_{m_i} : \vec{T}_i \vec{f}_i \\
\llbracket \text{class } C \text{ extends } D \{ \vec{T} \vec{f}; K \vec{M} \} \rrbracket &\stackrel{\text{def}}{=} (\text{class } C \text{ extends } D \{ \vec{T} \vec{f}; K \vec{M}' \}, \text{CT}) \text{ where } (\vec{M}', \text{CT}) = \llbracket \vec{M} \rrbracket_{\text{this}:C} \\
\llbracket U m (\vec{T} \vec{x}) \{ e \} \rrbracket_{\text{this}:C} &\stackrel{\text{def}}{=} (U m (C_m z) \{ e' \}, [C_m \mapsto \text{class } C_m \{ \vec{T} \vec{f}; K \} \cup \text{CT}) \\
&\text{where } (e', \text{CT}) = \llbracket e \rrbracket_{\text{this}:C, \vec{x}: \vec{T}}^{\vec{x}, z} \\
\llbracket \lambda(T x).(\nu \vec{u})(l, e, \sigma, \text{CT}) \rrbracket &\stackrel{\text{def}}{=} \lambda(T x).(\nu \vec{u})(l, e', \sigma', \text{CT} \cup \text{CT}') \\
&\text{where } (e', \text{CT}_0) = \llbracket e \rrbracket_{\vec{x}: \vec{T}, \vec{u}: \vec{T}'}^{\vec{x}, z}, (\sigma', \text{CT}_1) = \llbracket \sigma \rrbracket_{\vec{x}: \vec{T}, \vec{u}: \vec{T}'}^{\vec{x}, z}, \text{CT}' = \text{CT}_0 \cup \text{CT}_1
\end{aligned}$$

The encoding of other values are identical. Next we define the main rules for the expressions. Others are just homomorphic like $e_0; e_1$ below. We let $(e'_i, \text{CT}_i) = \llbracket e_i \rrbracket_{\Gamma}^{\vec{x}, z}$ for $i \geq 0$ in the following.

$$\begin{aligned}
\llbracket y \rrbracket_{\Gamma}^{x_1, \dots, x_n, z} &\stackrel{\text{def}}{=} \begin{cases} (z.\mathbf{f}_i, \emptyset) & \text{if } y = x_i \\ (y, \emptyset) & \text{otherwise} \end{cases} & \llbracket \text{this} \rrbracket_{\Gamma}^{\vec{x}, z} &\stackrel{\text{def}}{=} (\text{this}, \emptyset) \\
\llbracket y := e_0 \rrbracket_{\Gamma}^{x_1, \dots, x_n, z} &\stackrel{\text{def}}{=} \begin{cases} (z.\mathbf{f}_i := e'_0, \text{CT}_0) & \text{if } y = x_i \\ (y := e'_0, \text{CT}_0) & \text{otherwise} \end{cases} & \llbracket e_0; e_1 \rrbracket_{\Gamma}^{\vec{x}, z} &\stackrel{\text{def}}{=} (e'_0; e'_1, \text{CT}_i) \\
\llbracket e_0.m(e_1, \dots, e_n) \rrbracket_{\Gamma}^{\vec{x}, z} &\stackrel{\text{def}}{=} (e'_0.m(\text{new } C_m(e'_1, \dots, e'_n)), \text{CT}_i \cup [C_m \mapsto \text{class } C_m \{ \vec{T} \vec{f}; K \}]) \\
&\text{where } \Gamma \vdash e_0 : C
\end{aligned}$$

The mapping of a configuration, threads and stores are defined as follows.

$$\begin{aligned}
\llbracket (\nu \vec{o}\vec{c}) \prod l_i[(\nu \vec{x}_i)(P_i, \sigma_i, \mathbf{CT}_i)] \rrbracket &\stackrel{\text{def}}{=} (\nu \vec{o}\vec{c}) \prod l_i[\llbracket (\nu \vec{x}_i)(P_i, \sigma_i, \mathbf{CT}_i) \rrbracket_{\vec{o}, \vec{c}}] \\
\llbracket (\nu \vec{x}\vec{x}_1 \dots \vec{x}_n)(Q \mid \prod_j P_j, \sigma \cup \bigcup_j \sigma_j, \mathbf{CT}) \rrbracket_\Gamma &\stackrel{\text{def}}{=} (\nu \vec{x}\vec{y}\vec{o}) (Q' \mid \prod_j P'_j, \sigma' \cup \bigcup_j \sigma'_j, \mathbf{CT} \cup \mathbf{CT}') \\
\text{with } \llbracket Q \rrbracket_{\Gamma, \vec{x}; \vec{T}} &= (Q', \mathbf{CT}_1) \quad \llbracket P_j \rrbracket_{\Gamma, \vec{x}_j; \vec{T}_j}^{x_j, y_j} = (P'_j, \mathbf{CT}_j) \quad \llbracket \sigma \rrbracket_{\Gamma, \vec{x}; \vec{T}} = (\sigma', \mathbf{CT}_0) \\
\llbracket \sigma_j \rrbracket_{\Gamma, \vec{x}_j; \vec{T}_j}^{x_j, o_j, m, C} &= (\sigma'_j, \mathbf{CT}_{0j}) \quad \mathbf{CT}' = \bigcup \mathbf{CT}_{0j} \cup \bigcup \mathbf{CT}_j \cup \mathbf{CT}_1 \cup \mathbf{CT}_0 \\
\llbracket \emptyset \rrbracket_\Gamma &\stackrel{\text{def}}{=} \llbracket \emptyset \rrbracket_{\Gamma}^{x, y, o, m, C} \stackrel{\text{def}}{=} (\emptyset, \emptyset) \\
\llbracket \sigma \cdot [x \mapsto v] \rrbracket_\Gamma &\stackrel{\text{def}}{=} (\sigma' \cdot [x \mapsto v'], \mathbf{CT}' \cup \mathbf{CT}) \\
\llbracket \sigma \cdot [o \mapsto (C, \vec{f} : \vec{v})] \rrbracket_\Gamma &\stackrel{\text{def}}{=} (\sigma' \cdot [o \mapsto (C, \vec{f} : \vec{v}')], \mathbf{CT}' \cup \mathbf{CT}) \\
\llbracket \sigma \cdot [\vec{x} \mapsto \vec{v}] \rrbracket_\Gamma^{x, y, o, m, C} &\stackrel{\text{def}}{=} (\sigma' \cdot [y \mapsto o] \cdot [o \mapsto (C_m, \vec{f} : \vec{v}')], \mathbf{CT}' \cup \mathbf{CT}) \\
\llbracket \sigma \cdot [o' \mapsto (C, \dots)] \rrbracket_\Gamma^{x, y, o, m, C} &\stackrel{\text{def}}{=} \llbracket \sigma \rrbracket_\Gamma^{x, y, o, m, C}
\end{aligned}$$

with $\llbracket v \rrbracket_\Gamma = (v', \mathbf{CT}')$ and $\llbracket \sigma \rrbracket_\Gamma = (\sigma', \mathbf{CT})$. Let $\mathcal{R} = \{((\nu \vec{u})(N \mid M), (\nu \vec{u})(\llbracket N \rrbracket_\Gamma \mid M))\}$. Then we prove $\mathcal{R}$ respects the action predicate and reduction-closed up to $\mapsto$, i.e.

- if $(\nu \vec{u})(N \mid M) \longrightarrow N_1$, then there exists N_2 s.t. $(\nu \vec{u})(\llbracket N \rrbracket_\Gamma \mid M) \mapsto^* \longrightarrow^* N_2$ with $N_1 \equiv (\nu \vec{u}')(N' \mid M')$ and $N_2 \equiv (\nu \vec{u}')(N' \mid M')$ for some $\vec{u}', N'$ and M' .
- (a) if $(\nu \vec{u})(\llbracket N \rrbracket_\Gamma \mid M) \longrightarrow N_2$, then there exists N_1 s.t. $(\nu \vec{u})(N \mid M) \longrightarrow N_1$ with $N_1 \equiv (\nu \vec{u}')(N' \mid M')$ and $N_2 \mapsto^* (\nu \vec{u}')(N' \mid M')$ for some $\vec{u}', N'$ and M' ; or
- (b) if $(\nu \vec{u})(\llbracket N \rrbracket_\Gamma \mid M) \mapsto N_2$, then there exists N_1 s.t. $(\nu \vec{u})(N \mid M) \longrightarrow N_1$ with $N_1 \equiv (\nu \vec{u}')(N' \mid M')$ and $N_2 \longrightarrow^* (\nu \vec{u}')(N' \mid M')$ for some $\vec{u}', N'$ and M'

The key cases are variable-read $\llbracket z \rrbracket_\Gamma^{x_1, \dots, x_n, z}$ and variable-write $\llbracket z := v \rrbracket_\Gamma^{x_1, \dots, x_n, z}$, for which we can directly use (cr) in the transformation rule in § 7.2. Since $\mapsto \subseteq \cong$, we have $N \cong \llbracket N \rrbracket_\Gamma$, as desired. $\square$

H. PROOFS OF THEOREM 7.2 (3)

This section proves that (RMI3) is equivalent to (Opt3) under the assumption there is no call-back. Let $e'_i = \text{deserialize}(v_i)$ where $v_i = \lambda(\text{unit } x).(\nu \vec{u})(l, a, \sigma_i)$ is a serialised value at line i in (Opt3) ($3 \leq i \leq 5$). We show the body of (Opt3) is equivalent with `return r.run(freeze{e[e'/b]; z})`. With out loss of generality, we firstly simplify the program with two arguments as follows:

```

1  int m3(RemoteObject r, MyObj a) {
2      return r.g(a, r.f(a));
3  }

```

The corresponding optimised program is:

```

1  int mOpt3(RemoteObject r, MyObj a){
2      ser<MyObj> b1 = serialize(a);

```

```

3 |   ser<MyObj> b2 = serialize(a);
4 |   thunk<int> t = freeze {
5 |     r.g(deserialize(b2), r.f(deserialize(b1)));
6 |   };
7 |   return r.run(t);
8 | }

```

First by Proposition 7.1 (1), we can ignore an effect of the class downloading. Since there is no call-back, we can also ignore an effect of the method invocation “ $r.f$ ” to the timing of the serialisation “ $b2 = \text{serialize}(a)$ ”. Hence there is no interleaving between “ $b1$ ” in Line 2 and “ $b2$ ” in Line 3 from the location m . Thus we apply (ni) to the optimised code, and it is equated to:

$$P \stackrel{\text{def}}{=} \text{return}(c) \text{ r.run}(\text{freeze}\{r.g(\text{deserialize}(v_2), r.f(\text{deserialize}(v_1)))\})$$

where $v_i = \lambda(\text{unit } x).(\nu \vec{u})(l, a, \sigma_i)$ is a serialised value at b_i in (Opt3) above. Note that if the method invocation $r.f$ does not terminate in $m3$, then the both programs $m3$ and $mOpt3$ diverge, thus they are trivially equivalent. Hence we assume the case the method invocation $r.f$ terminates.

First we execute the original program. We omit a surrounding context where it is unnecessary. We also assume the location m (server) contains a store which includes $[r \mapsto (C, \dots)]$ and a class table of class C which contains methods f and g . We also omit C and channel restriction of c and c_1 .

$$\begin{array}{ll}
l[\text{return}(c) \text{ r.g}(a, r.f(a))] \mid m[0] & \\
\mapsto l[\text{return}(c) \text{ r.g}(a, \text{await } c_1) \mid \text{go } r.f(\text{serialize}(a)) \text{ with } c_1] \mid m[0] & (\text{ni}) \\
\longrightarrow l[\text{return}(c) \text{ r.g}(a, \text{await } c_1) \mid \text{go } r.f(v_1) \text{ with } c_1] \mid m[0] & (s1) \\
\mapsto^+ l[\text{return}(c) \text{ r.g}(a, \text{await } c_1) \mid m[r.f(\text{deserialize}(v_1)) \text{ with } c_1]] & (\text{ni}) \\
\mapsto^+ l[\dots] \mid m[(\nu \vec{u})(r.f(a) \text{ with } c_1, \sigma_1)] & (\text{ni}) \\
\longrightarrow l[\dots] \mid m[\text{return}(c_1) \text{ } n] & (\dagger) \\
\mapsto^+ l[\text{return}(c) \text{ r.g}(a, n)] \mid m[0] & (\text{rm}, l2, \text{ni}) \\
\mapsto \longrightarrow \mapsto^+ l[\text{return}(c) \text{ await } c'] \mid m[r.g(\text{deserialize}(v_2), n) \text{ with } c'] & (\text{ni}, s2, \text{ni}) \\
\mapsto^+ l[\text{return}(c) \text{ await } c'] \mid (\nu \vec{u})(m[\text{return}(c') \text{ r.g}(a, n), \sigma_2]) & (\text{ni})
\end{array}$$

At $(\dagger)$, we assume the method invocation $r.f$ terminates and returns the answer n . If $r.f$ does not terminate in $m3$, as seen in the following, $mOpt3$ diverges, thus they are trivially equivalent. Now we look at execution of the optimised code.

$$\begin{array}{ll}
l[P] \mid m[0] & \\
\mapsto l[\text{return}(c) \text{ r.run}(\lambda().(r.g(\text{deserialize}(v_2), r.f(\text{deserialize}(v_1)))))] \mid m[0] & \\
\mapsto^+ l[\text{return}(c) \text{ await } c'] \mid (\nu \vec{u})(m[\text{return}(c') \text{ r.g}(\text{deserialize}(v_2), r.f(a)), \sigma_1]) & \\
\longrightarrow l[\text{return}(c) \text{ await } c'] \mid m[\text{return}(c') \text{ r.g}(\text{deserialize}(v_2), n)] & (\dagger) \\
\mapsto^+ l[\text{return}(c) \text{ await } c'] \mid (\nu \vec{u})(m[\text{return}(c') \text{ r.g}(a, n), \sigma_2]) &
\end{array}$$

In the above, we know that we can get n at $(\dagger)$ if and only if n is obtained at $(\dagger)$. Also the serialisation (s1,s2) in the original code can be derived with the assumption such that v_1 and v_2 are serialised values in Line 2 and Line 3 in the optimised code. Since $\mapsto \subseteq \cong$, (RMI3) is equivalent with (Opt3).