

A theory of RPC calculi for client–server model

KWANGHOON CHOI

Chonnam National University, Gwangju, Republic of Korea
(e-mail: kwanghoon.choi@jnu.ac.kr)

BYEONG-MO CHANG

Sookmyung Women's University, Seoul, Republic of Korea
(e-mail: chang@sookmyung.ac.kr)

Abstract

With multi-tier programming languages, programmers can specify the locations of code to run in order to reduce development efforts for the web-based client–server model where programmers write client and server programs separately and test the multiple programs together. The RPC calculus, one of the foundations of those languages by Cooper and Wadler, has the feature of symmetric communication in programmer's writing arbitrarily deep nested client–server interactions. The feature of the calculus is fully implemented by asymmetric communication in trampolined style suitable for the client–server model. However, the existing research only considers a stateless server strategy in which all server states are encoded for transmission to the client so that server states do not need to be stored in the server. It cannot always correctly handle all stateful operations involving disks or databases. To resolve this problem, we first propose new stateful calculi that fully support both symmetric communication from the programmer's viewpoint and asymmetric communication in its implementation using trampolined style. All the existing calculi either provide only the feature of asymmetric communication or propose only symmetric implementation suitable for the peer-to-peer model, rather than the client–server model. Second, the method used to design our stateful server strategy is based on a new locative type system which paves the way for a theory of RPC calculi for the client–server model. Besides proposing the new stateful calculi, this theory can improve the existing stateless server strategy to construct new state-encoding calculi that eliminate runtime checks on remote procedure calls present in the existing strategy, and it enables us to design a new mixed strategy that combines the benefits of both kinds of strategies. As far as we know, there are no typed multi-tier calculi that offer programmers the feature of symmetric communication with the implementation of asymmetric communication under the three strategies together.

1 Introduction

Modern computing environments such as web systems involve programming not a single machine but several distributed machines together. For example, a web system basically consists of a web server that accesses databases and a web client that provides user interfaces, and they are connected by a network. Programmers have to develop two individual programs separately for the two machines, which increases the programmer's burden. Further, they need to test the two programs together, which is more complex than they do with one program on a single machine.

The programmer's task can be alleviated by the use of *multi-tier programming languages* (Murphy VII *et al.*, 2004; Neubauer & Thiemann, 2005; Balat, 2006; Serrano *et al.*,

2006; Cooper *et al.*, 2007; Murphy, 2008; Cooper & Wadler, 2009; Rastogi *et al.*, 2014; Chlipala, 2015; Serrano & Prunet, 2016). Such languages are equipped with the locations feature, which enables description of the location to describe where a specified part of the code should execute. Using this feature, programmers develop a single program that can be used freely in different locations such as the server and the client in one programming language. Then compilers will automatically generate separate programs for each location, guaranteeing communication integrity among the differently located programs.

Notably, Remote Procedure Call (RPC) calculus (Cooper & Wadler, 2009) offers the feature of symmetric communication between client and server. After adding to each function location annotations specifying where it must run, programmers can write arbitrarily deep nested client–server interactions by applying only the standard functions. Their compilation method will automatically map the deep interactions onto flat request–response interactions on the web system. The RPC calculus is the foundation of a practical multi-tier web programming language called Links (Cooper *et al.*, 2007).

The existing target language of the RPC calculus, called the Client-Server (CS) calculus, where a client and a server run separate programs, was designed for the server to maintain no session with individual clients. All server states during client–server interactions are appropriately encoded in the server for transmission to the client so that server states do not need to be stored in the server. The CS calculus uses asymmetric communication, which is almost free in the client–server model, and supports symmetric communication from the programmer’s viewpoint using *trampolined style* (Ganz *et al.*, 1999). Thus, the implementation of the RPC calculus was aligned with the well-known *RESTful* architecture of web systems, where web services consist of stateless operations.

However, the existing RPC and CS calculi cannot always correctly handle all stateful operations involving disks or databases. Some server states with disks or databases are not easily serialised, and therefore, when the server calls a client function, for example, between two subsequent database operations, it is not easy to encode all server states left after the client function call. Even when server states are serialisable, passing serialised server states between client and server repeatedly would increase communication overheads, giving rise to efficiency concerns. To address this problem, a stateful server strategy for the RPC calculus is necessary. This looks natural because supporting stateful interactions is common in web systems. For example, Java HttpSession provides a way to identify a user across more than one page request or visit to a web site and to store information about that user.

First, we propose new calculi, λ_{rpc}^{state} and λ_{cs}^{state} as shown in Figure 1, based on a stateful server strategy to resolve the problems of the existing calculi. The new calculi explicitly represent server states using the runtime stack, as in the conventional programming languages. Most importantly, the new calculi fully support the feature of symmetric communication in the RPC calculus by using trampolined style that is easy to implement in the client–server model. There have been several multi-tier web programming languages adopting stateful strategies, but none of them support both symmetric communication from the programmer’s viewpoint and asymmetric communication from the implementor’s viewpoint. Hop (Serrano *et al.*, 2006; Serrano & Prunet, 2016), Ur/Web (Chlipala, 2015) and Eliom (Balat, 2006; Radanne, 2017) only provide programmers asymmetric communication where only the client invokes server functions freely but, for the other

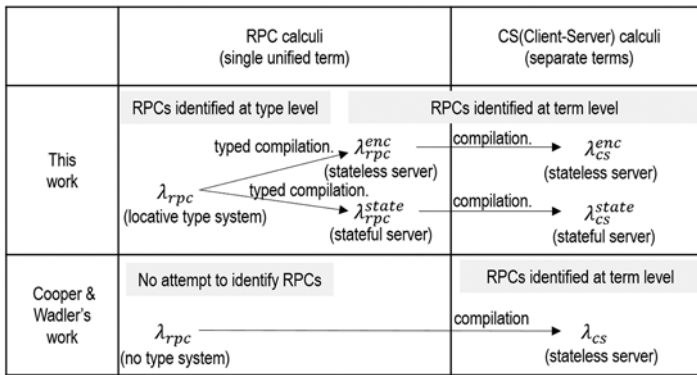


Fig. 1. Overview of a theory of RPC calculi.

way around, they must use a special network library, rather than a language construct. For example, Hop offers a web reactive programming library for server-to-client communication. The stateful strategies employed by the three works are different from our strategy fully supporting symmetric communication. Lambda5 (Murphy VII *et al.*, 2004; Murphy, 2008) does provide programmers symmetric communication but with distinct syntactic constructs for local and remote functions. The idea of its implementation is similar in that *continuations* span multiple worlds such as client and server, but it is based on stateful peer strategies. The semantic rules for all machines in a distributed system are the same, and so there is a gap between the semantic description for the peer-to-peer model and its implementation for the client–server model; to make it fit for the client–server model, they must employ some ad hoc method. Also note that to make it adopt a stateless server strategy, the semantic rules for server must be made different from those for client. A multi-tier calculus (Neubauer & Thiemann, 2005) is also for concurrently running processes, and so it suffers from the same limitation that Lambda5 has.

Second, the method used to design our stateful server strategy is based on a new locative type system, and it paves the way for a theory of RPC calculi for the client–server model. Under this theory, besides proposing the new stateful calculi, the theory can improve the existing stateless server strategy to construct new state-encoding calculi, λ_{rpc}^{enc} and λ_{cs}^{enc} as shown in Figure 1, that eliminate runtime checks on remote procedure calls in the existing strategy. The theory also enables us to design a new mixed strategy that combines the benefits of the state-encoding and stateful server strategy. We will elaborate more about our theory as follows.

A new RPC calculus λ_{rpc} in the theory is a typed version of the original RPC calculus with locative function types $\tau \xrightarrow{a} \tau'$ of functions that should run at location a . Typing with locative types can distinguish remote function applications from local ones at the type level, rather than at the term level by distinct constructs as in all the existing calculi except the RPC calculus. Using this type information, we are able to design a new type-directed compilation method. The new method is simpler than the existing untyped method because it has two kinds of compilation rules, one for client and the other for server, rather than eight kinds of rules, three for client and five for server, in the existing method. Also, the use of type information eliminates runtime checks on locations, which is present in the existing untyped compilation method. For example, due to the absence of location information in

function application terms, both a local function and a remote one can be invoked through the same function application term, e.g. $f M$. So, f is needed to check in runtime if it is a remote function because an invocation of remote functions is implemented differently from an invocation of local functions. With the locative types, either f is always a local function or it is always a remote function.

Our theory not only allows us to enjoy the benefits of the simple type-directed compilation method and the elimination of the runtime checks on locations, but also provides a framework to compare and to combine the two kinds of strategies. Under the same theory, it is easy to compare the two server strategies in terms of session management. In the state-encoding server strategy, one session corresponds to a single request–response interaction, whereas one session in the stateful server strategy can span multiple request–response interactions. A formal description of the comparison will be presented. Because our theory employs a typed approach, it is also possible to design a mixed strategy where the state-encoding server strategy is basically used to reduce the server resource consumption, but we switch to the stateful server strategy when necessary. This idea of a mixed strategy can be realised by adopting the monadic encapsulation of state (Launchbury & Peyton Jones, 1994; Timany *et al.*, 2017). Using the notion, stateful computations are encapsulated using monads, and they can be separated from the purity of functional language. In the design, we use the stateful calculi for the phase that uses stateful operations separated by the monadic type system, and we use the state-encoding calculi for the other phase that uses purely functional operations.

For the evaluation, we have implemented a prototype compiler of λ_{rpc} into two kinds of calculi with a locative type inference algorithm, and have implemented two evaluators running a client and a server communicating with the HTTP protocol.¹

Figure 1 shows an overview of our theory. As far as we know, there are no typed locative calculi with the feature of symmetric communication that are implemented with asymmetric communication using a stateless server strategy, a stateful server strategy and a mixed strategy together.

The contributions of this paper are as follows:

- We present a locative type system for the RPC calculus λ_{rpc} and prove its type soundness, guaranteeing that the locative information is preserved under the evaluation.
- We design the two RPC calculi λ_{rpc}^{enc} and λ_{rpc}^{state} with two semantically correct compilation methods of λ_{rpc} into λ_{rpc}^{enc} and λ_{rpc}^{state} , respectively.
- We also design two CS calculi λ_{cs}^{enc} and λ_{cs}^{state} with two semantically correct compilation methods of λ_{rpc}^{enc} into λ_{cs}^{enc} and of λ_{rpc}^{state} into λ_{cs}^{state} .
- We implement our stateless and stateful calculi to show their effectiveness by a location-type inference algorithm, two compilers and two HTTP-based evaluators.
- We formally compare the characteristics of session management of λ_{cs}^{enc} with those of λ_{cs}^{state} by extending the semantic rules with session annotations.
- We propose a method to design a mixed strategy of the two strategies by extending the RPC calculus with monadic encapsulation of state.

¹ <https://github.com/kwanghoon/rpccalculi>.

Syntax			
Location	a, b	$::=$	$\mathbf{c} \mid \mathbf{s}$
Term	L, M, N	$::=$	$V \mid LM$
Value	V, W	$::=$	$x \mid \lambda^a x.N$
Semantics			
$\frac{}{V \Downarrow_a V}$	(Value)	$\frac{L \Downarrow_a \lambda^b x.N \quad M \Downarrow_a W \quad N\{W/x\} \Downarrow_b V}{LM \Downarrow_a V}$	(Beta)

Fig. 2. The RPC calculus λ_{rpc} .

Section 2 introduces the RPC calculus λ_{rpc} with a locative type system and proves the type soundness of the calculus. Section 3 proposes two calculi λ_{rpc}^{enc} and λ_{cs}^{enc} as a new formulation of the server state-encoding strategy to implement the λ_{rpc} calculus. Section 4 extends the new formulation to two stateful calculi λ_{rpc}^{state} and λ_{cs}^{state} to implement the λ_{rpc} calculus with the stateful server strategy. Section 5 discusses related work. Section 6 concludes the paper. In the appendix, all proofs of theorems introduced in the sections are available.

2 The RPC calculus and its locative type system

Let us review the RPC calculus λ_{rpc} , proposed in Cooper & Wadler (2009). This is an ordinary call-by-value λ -calculus with location annotations on λ -abstractions. The annotations tell the locations where the λ -abstractions execute. This calculus was designed as the foundation of a practical multi-tier web programming language. The client–server model is assumed in the calculus, and so the location annotations are either \mathbf{c} denoting client or \mathbf{s} denoting server. The syntax of λ_{rpc} is thus defined as in Figure 2.

In Figure 2, λ_{rpc} has a big-step operational semantics with evaluation judgements,

$$M \Downarrow_a V$$

denoting the evaluation of a term M in the location a resulting in the value V . In the semantics, β -reduction is performed in the location annotated on the λ -abstraction. $N\{W/x\}$ is an ordinary substitution of W for x in N .

In λ_{rpc} , the evaluation of a term starts from the client. An example is:

$$(\lambda^{\mathbf{s}} f \cdot (\lambda^{\mathbf{s}} x \cdot x) (f \ c)) (\lambda^{\mathbf{c}} y \cdot (\lambda^{\mathbf{s}} z \cdot z) \ y)$$

Figure 3 depicts the flow between the client and server for the example. The client first requests the server to invoke a function $(\lambda^{\mathbf{s}} f \cdot \dots)$ at the application in a box named S1. During the evaluation of the body of the server function, the server calls a client function f in C1, which is $(\lambda^{\mathbf{c}} y \cdot \dots)$, with some constant c as an argument. This client function again requests the server to invoke an identity function $(\lambda^{\mathbf{s}} z \cdot z)$ in S2. Finally, a server function $(\lambda^{\mathbf{s}} x \cdot x)$ in S3 is locally applied to the result of $(f \ c)$.

2.1 A locative type system for λ_{rpc}

We propose a new type system for the untyped RPC calculus λ_{rpc} where function types

$$\tau \xrightarrow{a} \tau'$$




Fig. 3. Evaluation flow in λ_{rpc} starting from the client. (Labels C1, S1, S2 and S3 indicate the locations of the applications in the labelled boxes.)




Fig. 4. A locative type system for λ_{rpc} .

carry location annotations such that λ -abstractions of the function type run in the location a specified in the type annotation. We advocate the use of types because a typed RPC calculus is useful for compilation, which will be shown later. Figure 4 shows type terms and typing rules. Here, *base* denotes fundamental types such as *int*.

This type system uses typing judgements of the form $\Gamma \triangleright_a M : \tau$, which denotes that a term M has type τ under a type environment Γ at location a . By (T-Lam), a location specified in the λ -abstraction is defined to be carried by the function type, e.g. as $\Gamma \triangleright_c \lambda^s z.z : \tau \overset{s}{\rightarrow} \tau$ with location s . (T-Var) is defined as usual. There are three typing rules for applications classified by where the applications execute and where the functions of the applications execute. (T-App) pertains to local applications. The location where the functions of the application execute is the same as the location where the application executes. (T-Req) and (T-Call) are about remote applications. (T-Req) is a typing for applications of a server function requested by the client, whereas (T-Call) is a typing for applications of a client function called by server. Here is a typing derivation example:

$$(T-Req) \frac{\Gamma \triangleright_c \lambda^s z.z : \tau \overset{s}{\rightarrow} \tau \quad \Gamma \triangleright_c y : \tau}{\Gamma \triangleright_c (\lambda^s z.z) y : \tau}$$

The example λ_{rpc} term in Figure 3 can be typed under our type system using (T-Req) for the two function applications in the boxes S1 and S2, respectively, (T-App) for the function application in S3 and (T-Call) for the function application in C1.

For convenience, the typing derivation can be represented in an extended syntax as

$$(\lambda^s f \cdot (\lambda^s x \cdot x)^s (f^c c))^s (\lambda^c y \cdot (\lambda^s z \cdot z)^s y)$$

where f has type $base \xrightarrow{c} base$ and all of x, y, z and c have type $base$. In the extended syntax, each application term gets a location annotation. $L^a M$ at location a is for (T-App), $L^s M$ in the client is for (T-Req) and $L^c M$ in the server is for (T-Call). Then the extended syntax of the λ_{rpc} terms will uniquely determine typing derivations under location contexts. This representation will be used in our compilation later.

The proposed locative type system has type soundness in that location information is preserved under the evaluation in λ_{rpc} , as is proved by the following theorem.

Theorem 2.1 (Type Soundness for λ_{rpc}). *If $\Gamma \triangleright_a M : \tau$ and $M \Downarrow_a V$, then $\Gamma \triangleright_a V : \tau$.*

Note that the typed RPC calculus is a strict subset of the untyped calculus. There are two kinds of untypeable terms. One kind is due to simple type incompatibility in which the base type is not compatible with any function types, which does not interest us much. We could filter out such incompatible terms with the simply typed system. The other kind of untyped terms is due to location incompatibility. Because we refine function types with location information, two lambda terms that would have the same function type under the simply typed system can now have different types. Consider an application term $(\lambda^a f \cdot \dots) M$ where f has type $base \xrightarrow{c} base$ and M has type $base \xrightarrow{s} base$. Obviously, this term cannot be typed under the type system for λ_{rpc} . However, it is not difficult to make it typed without changing the computation by a transformation of M into $\lambda^c x \cdot M x$, which now has the same type as f and the same computation.² This idea can be generalised as a transformation $\llbracket M \rrbracket^{\tau \rightsquigarrow \tau'}$ where a term M of type τ is transformed into τ' , and τ becomes identical to τ' on the removal of all location annotations in the types.

$$\begin{aligned} \llbracket M \rrbracket^{\tau \rightsquigarrow \tau} &= M \\ \llbracket M \rrbracket^{\tau_1 \xrightarrow{a} \tau_2 \rightsquigarrow \tau_3 \xrightarrow{b} \tau_4} &= \lambda^b x \cdot \llbracket M \llbracket x \rrbracket^{\tau_3 \rightsquigarrow \tau_1} \rrbracket^{\tau_2 \rightsquigarrow \tau_4} \end{aligned}$$

The above analysis establishes that every typed term in the simply typed system with arbitrary location annotations can be made typed under our locative type system for λ_{rpc} immediately or by the computation-preserving transformation. Our location-typed approach can thus deal with all terms in the existing location-untyped approach.

We have implemented a unification-based type inference algorithm for our locative type system. Although polymorphic types are beyond the scope of this research, we believe that it is possible to extend our type system with them.

² In a calculus with effects, the transformation may delay or eliminate effects because it transforms expressions with effects into values without them. In this situation, the transformation must involve more, for example, let $y = M$ in $\lambda^c x \cdot y x$ to have the same effects.

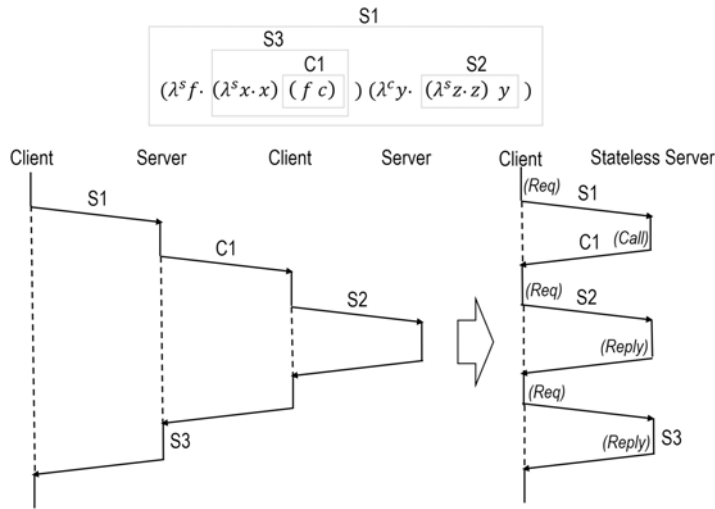


Fig. 5. Mapping of arbitrarily deep client–server interactions in λ_{rpc} onto flat request–responses interactions in λ_{rpc}^{enc} .

3 Locative calculi encoding server states

We present state-encoding calculi for the client–server model, following the stateless server strategy in Cooper & Wadler (2009). Under this strategy, when the server transfers control to the client, all the server states are encoded into a term so that the client can get the server to begin evaluating the term later. Unlike the original formulation, our calculi are new in that they are based on location information at the type level. This feature improves the original calculus, as we will see.

Figure 5 depicts the basic idea of the stateless server strategy on how to map every deep nesting of client–server interactions (on the left) into a series of request–response interactions between the client and the server (on the right). Each solid line indicates an active thread of control, and each dashed line indicates a stack frame which is waiting for a function to return. To support this mapping on the stateless server, we introduce the notion of *continuation* (Flanagan et al., 1993) to encode the server control to continue. For example, on invoking the client function in (C1), a continuation for (S3) is encoded into a term, and the encoded term is sent to the client. After the evaluation of (C1), the client requests the server to evaluate with the term to continue. Accordingly, there are two kinds of responses, as shown in Figure 5: one kind of response marked by (Call) involves the return of some encoded continuation, and the other kind marked by (Reply) simply returns values.

Under the stateless server strategy, it is assumed that every continuation carried by (Call)-type responses can be serialised so as to return to the client through the network. When this assumption does not hold, the continuation-based translation will not work. This observation leads to an alternative *stateful server* strategy in Section 4, which requires no such assumption. In this section, we develop our theory under the assumption of a serialisable continuation.

Now we introduce a new (state-encoding) RPC calculus λ_{rpc}^{enc} where local function applications are distinguished from remote function applications explicitly by the new term constructs req and call.

3.1 A state-encoding RPC calculus λ_{rpc}^{enc}

We first explore the state-encoding RPC calculus informally with an example, and then we will present its formal semantics in the next section. In λ_{rpc}^{enc} , there are three kinds of function applications: a local one $V_f(\overline{W})$ and two remote ones, $\text{req}(V_f, \overline{W})$ and $\text{call}(V_f, \overline{W})$. In the local function application, V_f is a local function. \overline{W} denotes a sequence of values that are its arguments. The terms $\text{req}(V_f, \overline{W})$ are used only in the client to transfer control to the server and then to invoke a server function V_f . The terms $\text{call}(V_f, \overline{W})$ are used only in the server to do the converse with a client function V_f .

Consider again the λ_{rpc} example term:

$$(\lambda^s f \cdot (\lambda^s x \cdot x) (f c)) \quad (\lambda^c y \cdot (\lambda^s z \cdot z) y)$$

whose evaluation begins in the client. This λ_{rpc} term can be implemented under the stateless server strategy by a λ_{rpc}^{enc} term as

$$\text{let } r = \text{req}(M_1, (M_2, ID)) \text{ in } r$$

where M_1 and M_2 implement the outermost function $(\lambda^s f \cdot (\lambda^s x \cdot x) (f c))$ and its argument $(\lambda^c y \cdot (\lambda^s z \cdot z) y)$, respectively. We used req to invoke the server function in the client. As previously explained, the notion of continuation was used to encode server states. In the stateless server strategy, server functions are implemented in continuation-passing style (CPS), whereas client functions are implemented in direct style (Flanagan *et al.*, 1993). M_1 is a CPS function converted from the server function. M_2 as a normal argument is converted from the original argument, and ID as a continuation argument denotes the identity continuation $\lambda^s w \cdot w$.

Now let us examine M_1 :

$$M_1 \triangleq \lambda^s (f, k) \cdot \text{call}(\lambda^c x \cdot \text{let } y = f(x) \text{ in } \text{req}(cont, y), \llbracket c \rrbracket)$$

where $cont$ is $\lambda^s w \cdot \llbracket \lambda^s x \cdot x \rrbracket(w, k)$, and $\llbracket M \rrbracket$ represents an implementation of the λ_{rpc} term M for the server. Here, f is the client function to be invoked, and k is a continuation, which is ID in the example. The body of M_1 is in the context of the server. We use call to invoke the client function ‘ $\lambda^c x \cdot \dots$ ’ in M_1 : it transfers control to the client to evaluate $f \llbracket c \rrbracket$ by ‘ $\text{let } y = f(x) \text{ in } \dots$ ’ above, and then it returns control to the server by req with the continuation $cont$ representing $(\lambda^s x \cdot x) [-]$ by ‘ $\dots \text{ in } \text{req}(cont, y)$ ’ above. The client function in call thus supports the behaviour of *commuting from the server to the client*, where $cont$ is regarded as an encoding of the server state. We call such a client function a *commuting function*, and our compiler will thus compose commuting functions.

M_2 has the following simple structure:

$$M_2 \triangleq \lambda^c y \cdot \text{let } d = \lambda^s (z, k) \cdot k(z) \text{ in} \\ \text{let } r = \text{req}(d, (y, ID)) \text{ in } r$$

where d is a CPS function converted from $\lambda^s z \cdot z$, and the client invokes this (server) function with its arguments y and ID .

3.2 The formal semantics of λ_{rpc}^{enc}

This section first introduces the syntax and semantics for the state-encoding RPC calculus, as shown in Figure 6. The syntax of M terms resembles A-normal form

Syntax

Term	M	$::= V \mid V_f(\overline{W}) \mid \text{req}(V_f, \overline{W}) \mid \text{call}(V_f, \overline{W}) \mid \text{let } x = M \text{ in } M$
Value	V, W	$::= x \mid \lambda^a \bar{x}. M$
Client context	Π	$::= \text{ctx } x M$
Server context	Δ	$::= \varepsilon$

Semantics

Client:	$M \mid \varepsilon \Rightarrow^{enc} M' \mid \varepsilon$ or $M \mid \varepsilon \Rightarrow^{enc} \Pi \mid M'$
(AppC)	$\text{let } y = (\lambda^c \bar{x}. M_0)(\overline{W}) \text{ in } M \mid \varepsilon \Rightarrow^{enc} \text{let } y = M_0\{\overline{W}/\bar{x}\} \text{ in } M \mid \varepsilon$
(Req)	$\text{let } x = \text{req}(\lambda^s \bar{x}. M_0, \overline{W}) \text{ in } M \mid \varepsilon \Rightarrow^{enc} \text{ctx } x M \mid (\lambda^s \bar{x}. M_0)(\overline{W})$
(ValC)	$\text{let } x = V \text{ in } M \mid \varepsilon \Rightarrow^{enc} M\{V/x\} \mid \varepsilon$
(LetC)	$\text{let } x = (\text{let } y = M_1 \text{ in } M_2) \text{ in } M \mid \varepsilon \Rightarrow^{enc} \text{let } y = M_1 \text{ in } (\text{let } x = M_2 \text{ in } M) \mid \varepsilon$
Server:	$\Pi \mid M \Rightarrow^{enc} \Pi \mid M'$ or $\Pi \mid M \Rightarrow^{enc} M' \mid \varepsilon$
(AppS)	$\Pi \mid (\lambda^s \bar{x}. M_0)(\overline{W}) \Rightarrow^{enc} \Pi \mid M_0\{\overline{W}/\bar{x}\}$
(Call)	$\text{ctx } x M \mid \text{call}(\lambda^c \bar{x}. M_0, \overline{W}) \Rightarrow^{enc} \text{let } x = (\lambda^c \bar{x}. M_0)(\overline{W}) \text{ in } M \mid \varepsilon$
(Reply)	$\text{ctx } x M \mid V \Rightarrow^{enc} \text{let } x = V \text{ in } M \mid \varepsilon$

Fig. 6. A state-encoding client–server calculus λ_{rpc}^{enc} .

(Flanagan *et al.*, 1993). Values are a variable or a location-annotated function and are denoted by V or W . Note that \bar{x} denotes a sequence of variables, and \overline{W} denotes a sequence of values. There are two kinds of remote function calls: $\text{call}(V_f, \overline{W})$ for invocation of a client function V_f in the server with arguments \overline{W} , and $\text{req}(V_f, \overline{W})$ for invocation of a server function in the client. The let constructs $\text{let } x = M_1 \text{ in } M_2$ bind x to some intermediate value from the evaluation of M_1 and continue to evaluate M_2 .

The client context Π , which is $\text{ctx } x M$, waits for a value to bind for x and to evaluate M . The server context Δ for the state-encoding RPC calculus is always empty.

In the semantics, *client* \mid *server* is the notation for snapshots in a client–server-based distributed system. We call such a snapshot a *configuration*. There are two kinds of configurations in this calculus:

- $M \mid \varepsilon$ for the client to evaluate a term M with the empty server context (ε)
- $\Pi \mid M$ for the server to evaluate a term M with a pending client context Π

The semantics is described as a relation on configurations as

$$\text{Conf} \Rightarrow^{enc} \text{Conf}'$$

Local evaluation is in the form of $M \mid \varepsilon \Rightarrow^{enc} M' \mid \varepsilon$ for the client and $\Pi \mid M \Rightarrow^{enc} \Pi \mid M'$ for the server. When control moves from the client to the server, the relevant semantic rule will have the form $M \mid \varepsilon \Rightarrow^{enc} \Pi \mid M'$. For the opposite direction, it is $\Pi \mid M \Rightarrow^{enc} M' \mid \varepsilon$. The notations R^+ and R^* for a relation R are defined as usual.

In the semantics, the evaluation of a term M begins in the client with the empty server state and normally finishes when it reaches the value V , as $M \mid \varepsilon \Rightarrow^{enc*} V \mid \varepsilon$, meaning that M evaluates to V .

For evaluation in the client, (AppC) performs a client function application. $M\{\overline{W}/\bar{x}\}$ denotes an ordinary parallel substitution of \overline{W} for \bar{x} in M . (ValC) binds an intermediate value with a let-variable and continues to evaluate a term in the let body. (LetC) takes

a nested let binding out for evaluation. With (Req), the client can request the server to execute a server function application, leaving the client context.

The server-side evaluation, which is initiated by (Req) in the client, continues with (AppS). It is reminiscent of (AppC) for the client. The evaluation in the server will reach either a value V or a call $\text{call}(f', \overline{W'})$. For the case of a value, (Reply) returns the value to the client context $\text{ctx } x M$ as $\text{let } x = V \text{ in } M$. For the case of a call, (Call) moves the client function call into the client context as $\text{let } x = f'(\overline{W'}) \text{ in } M$. In any case, the server state will be empty. This is the stateless server strategy. The typical steps in the server are as follows:

	[Client]		[Server]	
Initiation	$\text{let } x = \text{req}(f, \overline{W}) \text{ in } M$ $\Rightarrow^{enc} \text{ctx } x M$		ε $f(\overline{W})$	by (Req)
either	$\Rightarrow^{enc} + \text{ctx } x M$ $\Rightarrow^{enc} \text{let } x = V \text{ in } M$		V ε	by (Reply)
or	$\Rightarrow^{enc} + \text{ctx } x M$ $\Rightarrow^{enc} \text{let } x = f'(\overline{W'}) \text{ in } M$		$\text{call}(f', \overline{W'})$ ε	by (Call)

The proposed semantic rules are defined in a minimal way enough for λ_{rpc}^{enc} to be an intermediate calculus for compilation of λ_{rpc} , as will be seen. The semantic rules could be easily extended to cover all syntactic terms including, for example, local applications without let in (AppC) or remote applications without let in (Req), but this extension is not essential for our purpose. In the server side, the intermediate calculus for CPS conversion has only to support local applications, call, and values, all without let.

Note that the evaluation of call in the semantics is allowed only in tail positions because the server state right after the call must be empty. To map λ_{rpc} terms without such a restriction into λ_{rpc}^{enc} terms with it, compiler support is required as will be explained soon.

3.3 A typed compilation of λ_{rpc} into λ_{rpc}^{enc}

Now we are ready to present a typed compilation of the RPC calculus into the state-encoding RPC calculus. Figure 7 shows our compilation rules.

The compilation rules for client terms are described as $C[[M_{rpc}]] = M_{rpc}^{enc}$. Here, M_{rpc} actually denotes a typing derivation for a source term in the RPC calculus, rather than the source term itself. Such a typing derivation provides each application subterm location information where the subterm should be evaluated, as explained in Section 2. M_{rpc}^{enc} denotes a compiled target term.

(VarC) compiles a variable term simply into a variable. (LamCC) compiles each client function using direct style. (LamCS) compiles each server function using CPS conversion. The reason for using the conversion will be explained later. (AppCC) is a compilation rule for a local client function application typed under the typing rule (T-App), denoted by $L^c M$. It is a conventional compilation rule that compiles L into a term generating a client function, compiles M into another term generating an argument and ends with a local application term. (AppCS) compiles a remote server function application typed under the typing rule (T-Req), denoted by $L^s M$. It is almost the same as (AppCC). The

Client: $C[[M_{rpc}]] = M_{rpc}^{enc}$	
(VarC)	$C[[x]] = x$
(LamCC)	$C[[\lambda^c x.M]] = \lambda^c x.C[[M]]$
(LamCS)	$C[[\lambda^s x.M]] = \lambda^s(x, k).S[[M]] k$
(AppCC)	$C[[L^c M]] = \text{let } f = C[[L]] \text{ in}$ $\text{let } x = C[[M]] \text{ in}$ $\text{let } r = f(x) \text{ in } r$
(AppCS)	$C[[L^s M]] = \text{let } f = C[[L]] \text{ in}$ $\text{let } x = C[[M]] \text{ in}$ $\text{let } r = \text{req}(f, (x, \lambda^s y.y)) \text{ in } r$
Server: $S[[M_{rpc}]] K = M_{rpc}^{enc}$	
(VarS)	$S[[x]] K = K(x)$
(LamSC)	$S[[\lambda^c x.M]] K = K(\lambda^c x.C[[M]])$
(LamSS)	$S[[\lambda^s x.M]] K = K(\lambda^s(x, k).S[[M]] k)$
(AppSC)	$S[[L^c M]] K = S[[L]] (\lambda^s f.$ $S[[M]] (\lambda^s x.$ $\text{call}(\lambda^c x.\text{let } y = f(x) \text{ in req}(K, y), x)))$
(AppSS)	$S[[L^s M]] K = S[[L]] (\lambda^s f.$ $S[[M]] (\lambda^s x.$ $f(x, K)))$

Fig. 7. A typed compilation for λ_{rpc}^{enc} .

only difference is the replacement of a local application term ($f x$) with a remote application term $\text{req}(f, (x, \lambda^s y \cdot y))$. Our compilation rules ensure that f is a CPS-converted function. Therefore, it will take not only a function argument x but also a continuation. The identity continuation ($\lambda^s x \cdot x$) is given to the remote application term because it begins the evaluation of a CPS term.

(AppCC) and (AppCS) introduce the extra layer of ‘let $r = \dots$ in r ’ surrounding local and remote applications to adapt our compilation to the minimal semantic rules (AppC) and (Req), respectively.

The basic idea of compilation for the server is CPS conversion. The compilation rules for server terms are described as $S[[M_{rpc}]] K = M_{rpc}^{enc}$, where K is (a value for) a continuation. In the compiled term, we first evaluate M_{rpc} and then continue with K and the evaluation result.

(VarS) compiles a variable into a term applying a continuation K to a value bound to the variable. (LamSC) and (LamSS) have a similar structure as (VarS). They only replace the compiled variable x in (VarS) with $C[[\lambda^c x \cdot M]]$. (AppSS) compiles a local server function application typed under (T-App), and it is the same as the CPS conversion of an application. (AppSC) is the most important compilation rule. A typing derivation denoted by $L^c M$ typed under (T-Call) is compiled as this. Both L and M are first compiled into terms whose evaluation will be bound to a function f and an argument x in the same way as in (AppSS). Then the compilation rule generates $\text{call}(\lambda^c x \cdot \text{let } y = f(x) \text{ in req}(K, y), x)$. By this remote client function application, control is first transferred to the client to evaluate $f(x)$. Control is then returned to the server by $\text{req}(K, y)$ to continue K in the server with the application result y . Our compiler thus supports the behaviour of commuting from the server to the client by composing this client function, which we call a commute function.

In regard to a client function call from the server, the RPC calculus can allow it in non-tail positions but the state-encoding RPC calculus allows it only in the tail position. This mismatch is resolved as follows. The commute function above is composed to have

a request with a continuation after the client function call. The continuation will hold the context of the client function call in a non-tail position. In this way, (AppSC) ensures that $\text{call}(-, -)$ can be always in a tail position in the state-encoding RPC calculus no matter where client function calls from the server appear in the RPC calculus.

Now, we prove the correctness of our compilation for λ_{rpc}^{enc} as follows:

Theorem 3.1 (Correctness of Compilation for λ_{rpc}^{enc}). *Assume a well-typed λ_{rpc} term M under the locative type system:*

- If $M \Downarrow V_c$, then $C\llbracket M \rrbracket \mid \varepsilon \Rightarrow^{enc*} C\llbracket V \rrbracket \mid \varepsilon$.
- If $M \Downarrow V_s$, then $\Pi \mid S\llbracket M \rrbracket K \Rightarrow^{enc*} \Pi \mid S\llbracket V \rrbracket K$ for all Π and K .

The basic idea of our proof is this. A sequence of evaluation steps is built in λ_{rpc}^{enc} matched to each subtree of evaluation in λ_{rpc} by induction on the height of the evaluation trees, and then the sequences are put together to build what is matched to the whole evaluation tree in the conditions. As a natural consequence, the proof also guarantees that client–server communication in λ_{rpc} is preserved in λ_{rpc}^{enc} , though the theorem itself does not state this explicitly. Figure 5 shows an example of the correspondence of client–server communication between λ_{rpc} and λ_{rpc}^{enc} with identical box names S1, C1 and S2 labelled on the related flows.

3.4 Separating client and server terms in λ_{cs}^{enc}

Although the state-encoding RPC calculus has the potential to make the λ_{rpc} terms run on the client and the server separately, in the λ_{rpc}^{enc} terms, the client part and the server part are still together. In this section, we propose a state-encoding CS calculus λ_{cs}^{enc} where the two parts are clearly separated, and names are used in one part to refer to functions in the other part.

Figure 8 shows the syntax and semantics of the state-encoding CS calculus. Terms in the calculus are now denoted by m , and values are variables or closures $\text{clo}(F, \bar{v})$ denoted by v or w . F is a name for a closed function. A closed function, $\bar{z}\lambda^a \bar{x} \cdot m$, is a function with a list of free variables \bar{z} . In a closure, \bar{v} consists of values to substitute for the free variables \bar{z} in the closed function. A substitution of \bar{v} for \bar{x} in m is written as $m\{\bar{v}/\bar{x}\}$. The function store ϕ_a is a mapping of names into closed functions for each location a . π and δ are client and server contexts, respectively.

Configurations in the semantics of the state-encoding CS calculus are as follows:

- $m \mid \varepsilon$ for the client to evaluate a term m with the empty server context ε
- $\pi \mid m$ for the server to evaluate a term m with a pending client context π

The semantic rules for the state-encoding CS calculus are isomorphic to those for the state-encoding RPC calculus under the aforementioned changes.

Figure 8 also shows the compilation rules of the state-encoding RPC calculus into the state-encoding CS calculus. The compilation rules are actually closure conversion gathering closed functions for each location. For a given λ_{rpc}^{enc} term M , the following compilation generates a λ_{cs}^{enc} term m :

$$CC\llbracket M \rrbracket = m$$

Syntax

Term	m	$::= v \mid v_f(\bar{w}) \mid \text{req}(v_f, \bar{w}) \mid \text{call}(v_f, \bar{w}) \mid \text{let } x = m \text{ in } m$
Value	v, w	$::= x \mid \text{clo}(F, \bar{v})$
Function store	ϕ_a	$::= \{ \dots, F = \bar{z}\lambda^a \bar{x}.m, \dots \}$
Client context	π	$::= \text{ctx } x m$
Server context	δ	$::= \varepsilon$

Semantics

Client:	$m \mid \varepsilon \Rightarrow^{enc} m' \mid \varepsilon$ or $m \mid \varepsilon \Rightarrow^{enc} \pi \mid m'$
(AppC)	$\text{let } y = \text{clo}(F, \bar{v})(\bar{w}) \text{ in } m \mid \varepsilon \Rightarrow^{enc} \text{let } y = m_0\{\bar{v}/\bar{z}, \bar{w}/\bar{x}\} \text{ in } m \mid \varepsilon$ if $\phi_c(F) = \bar{z}\lambda^c \bar{x}.m_0$
(Req)	$\text{let } x = \text{req}(\text{clo}(F, \bar{v}), \bar{w}) \text{ in } m \mid \varepsilon \Rightarrow^{enc} \text{ctx } x m \mid \text{clo}(F, \bar{v})(\bar{w})$ if $\phi_s(F) = \bar{z}\lambda^s \bar{x}.m_0$
(ValC)	$\text{let } x = v \text{ in } m \mid \varepsilon \Rightarrow^{enc} m\{v/x\} \mid \varepsilon$
(LetC)	$\text{let } x = (\text{let } y = m_1 \text{ in } m_2) \text{ in } m \mid \varepsilon \Rightarrow^{enc} \text{let } y = m_1 \text{ in } (\text{let } x = m_2 \text{ in } m) \mid \varepsilon$
Server:	$\pi \mid m \Rightarrow^{enc} \pi \mid m'$ or $\pi \mid m \Rightarrow^{enc} m' \mid \varepsilon$
(AppS)	$\pi \mid \text{clo}(F, \bar{v})(\bar{w}) \Rightarrow^{enc} \pi \mid m_0\{\bar{v}/\bar{z}, \bar{w}/\bar{x}\}$ if $\phi_s(F) = \bar{z}\lambda^s \bar{x}.m_0$
(Call)	$\text{ctx } x m \mid \text{call}(\text{clo}(F, \bar{v}), \bar{w}) \Rightarrow^{enc} \text{let } x = \text{clo}(F, \bar{v})(\bar{w}) \text{ in } m \mid \varepsilon$ if $\phi_c(F) = \bar{z}\lambda^c \bar{x}.m_0$
(Reply)	$\text{ctx } x m \mid v \Rightarrow^{enc} \text{let } x = v \text{ in } m \mid \text{emp}$

Compilation

$CC[[M_{rpc}^{enc}]]$	$= m$
$CC[[x]]$	$= x$
$CC[[\lambda^a x.m]]$	$= \text{clo}(F, \bar{z})$ if $\bar{z} = fv(\lambda^a x.m)$ and $\phi_a(F) = \bar{z}\lambda^a x.CC[[m]]$
$CC[[v(\bar{w})]]$	$= CC[[v]](CC[[\bar{w}]])$
$CC[[\text{let } x = m_1 \text{ in } m_2]]$	$= \text{let } x = CC[[m_1]] \text{ in } CC[[m_2]]$
$CC[[\text{req}(v, \bar{w})]]$	$= \text{req}(CC[[v]], CC[[\bar{w}]])$
$CC[[\text{call}(v, \bar{w})]]$	$= \text{call}(CC[[v]], CC[[\bar{w}]])$

Fig. 8. Syntax and semantics of λ_{cs}^{enc} with a compilation of λ_{rpc}^{enc} .

The compilation rules traverse a source term. Whenever a location-annotated function is encountered, a fresh name F is created, a list of free variables \bar{z} is collected over the function, and a mapping of this name into a constructed closed function is added to the function store at the location. fv is the standard function for free variable collection trivially extended with the located lambda calculus. Note that the compilation rules utilise implicit global function stores, with which the mapping of names and closed functions encountered during the compilation is checked. Actually, rather than checking the presence of closed functions in the function stores, we intend to collect the closed functions to construct the function stores.

Our main compilation method using $CC[-]$ compiles a closed state-encoding RPC term into a closed state-encoding CS term. As a side effect, it generates a function store ϕ_c for the client and a function store ϕ_s for the server. Now we can prove that a combined compilation of $C[-]$ with $CC[-]$ of RPC terms will generate state-encoding CS terms that correctly implement the RPC terms as follows.

Theorem 3.2 (Correctness of Compilation for λ_{cs}^{enc}). *Assume a well-typed term M under the locative type system. Given ϕ_c and ϕ_s :*

- If $M \Downarrow_c V$, then $CC[[C[[M]]]] \mid \varepsilon \Rightarrow^{enc*} CC[[C[[V]]]] \mid \varepsilon$.

By this compilation correctness theorem, it is demonstrated that the feature of the RPC calculus allowing arbitrarily deep nesting of control contexts between the client and the server can be correctly implemented by a series of flat request–responses under the stateless server strategy.

$$\begin{aligned}
 \phi_c & : \text{main} = \text{let } r_3 = \text{req}(\text{clo}(g_7, \{\}), \text{clo}(g_{10}, \{\}), \text{clo}(g_{11}, \{\})) \text{ in } r_3 \\
 & \quad g_2 = \{f_7, f_5, k_4\} \lambda^c z_9. \text{let } r_{10} = f_7 z_9 \text{ in } \text{req}(\text{clo}(g_1, \{f_5, k_4\}), r_{10}) \\
 & \quad g_{10} = \{\} \lambda^c y. \text{let } r_{14} = \text{req}(\text{clo}(g_8, \{\}), y, \text{clo}(g_9, \{\})) \text{ in } r_{14} \\
 \phi_s & : g_1 = \{f_5, k_4\} \lambda^s x_6. f_5(x_6, k_4) \\
 & \quad g_3 = \{f_7, f_5, k_4\} \lambda^s x_8. \text{call}(\text{clo}(g_2, \{f_7, f_5, k_4\}), x_8) \\
 & \quad g_4 = \{f_5, k_4\} \lambda^s f_7. \text{clo}(g_3, \{f_7, f_5, k_4\}) c \\
 & \quad g_5 = \{k_4, f\} \lambda^s f_5. \text{clo}(g_4, \{f_5, k_4\}) f \\
 & \quad g_6 = \{\} \lambda^s x, k_{11}. k_{11} x \\
 & \quad g_7 = \{\} \lambda^s f, k_4. \text{clo}(g_5, \{k_4, f\}) (\text{clo}(g_6, \{\})) \\
 & \quad g_8 = \{\} \lambda^s z, k_{15}. k_{15} z \\
 & \quad g_9 = \{\} \lambda^s x_{16}. x_{16} \\
 & \quad g_{11} = \{\} \lambda^s x_{17}. x_{17}
 \end{aligned}$$

Fig. 9. Compilation of the RPC term in Figure 3.

We have implemented a parser of λ_{rpc} , a location-type inference procedure for our type system, a compiler and an evaluator of λ_{rpc}^{enc} and λ_{cs}^{enc} on HTTP between the client and the server. Figure 9 shows a compilation of the RPC term in Figure 3, generated by the compiler. The evaluator starts with *main* in the client. For readability, each list of free variables is enclosed by $\{\}$, and so is each list of values for such free variables in closures.

Our state-encoding calculi are higher-order, whereas the original CS calculus remains first-order. This is because the original CS calculus employed an implementation method called *defunctionalisation* (Reynolds, 1972), which converts higher-order functions into first-order ones. However, the state-encoding calculi offer more freedom in the implementation of higher-order functions. They do not actually depend on any specific implementation methods such as closure conversion; they could even utilise defunctionalisation as well.

3.5 Discussion

3.5.1 Trampoline style

Our state-encoding calculi use a new trampolined style embedded in the semantics. This contrasts with the original CS calculus (Cooper & Wadler, 2009), which has the trampolined style (Ganz *et al.*, 1999) formulated using terms that introduces a special loop function called *trampoline* in the client as follows:

$$\begin{aligned}
 \text{Line 1: } & \text{trampoline}(x) = \text{case } x \text{ of} \\
 \text{Line 2: } & \quad | \text{Call}(f, x, k) \rightarrow \text{trampoline}(\text{req}(k, f(x))) \\
 \text{Line 3: } & \quad | \text{Return}(x) \rightarrow x \\
 \text{Line 4: } & \text{trampoline}(\text{req}(g, v))
 \end{aligned}$$

In Line 4, an invocation of a server function g from the client is surrounded by the trampoline function. It is like the semantic behaviour defined by (Req) in the state-encoding calculi. On receiving a datum like $\text{Call}(f, x, k)$ from the server in Line 2, the client applies f to x , and then invokes a continuation k with the application result. This is surrounded by the trampoline function again for repetition as many times as is necessary. For commuting between the server and the client like this, our compilation method composes commute functions, as explained in Section 3.3:

$$\text{commute} \triangleq \lambda^c(f, x, k) \cdot \text{let } y = f(x) \text{ in } \text{req}(k, y)$$

so that $Call(f, x, k)$ can be regarded as $commute(f, x, k)$. In the trampoline function, two branches in Lines 2 and 3, one for $Call$ and the other for $Return$, correspond to what are supported by (Call) and (Reply), respectively.

3.5.2 A location-typed approach

A difference between our state-encoding calculi and the original calculi (Cooper & Wadler, 2009) is the use of locative type information for compiling function applications. In the original location-untyped approach, location information is not available for function applications, whereas functions have location annotations. Both the client function and the server function can therefore be invoked in the same function application. For example, in the λ_{rpc} term:

$$\text{let } g = \lambda^a f \cdot f \ M \text{ in } \dots g (\lambda^c x \cdot x) \dots g (\lambda^s y \cdot y) \dots$$

both $(\lambda^c x \cdot x)$ and $(\lambda^s y \cdot y)$ are invoked in the same application term $f \ M$. As a result, every function application must check at runtime the location of the function to decide whether to use the trampoline function in case of the server function. In our location-typed approach, local function applications are clearly separated from remote function applications, and so no such dynamic check at runtime is required. We also showed that all terms in the location-untyped approach can be supported in the location-typed approach by the transformation explained in Section 2.

In addition to the absence of runtime checks on location, our location-typed approach results in a simple form of compilation where there are only two kinds of compilation rules, one for the client and the other for the server. The original compilation method in the location-untyped approach consists of eight kinds of compilation rules: three for the client and five for the server. The simple form helps explain the essential structure of our compilation rules: the compilation of λ_{rpc} into λ_{rpc}^{enc} is a combination of direct-style compilation for the client and CPS-style compilation for the server, and the compilation of λ_{rpc}^{enc} into λ_{cs}^{enc} is a variant of closure conversion classifying closed functions on locations. The simple compilation rules using locative types facilitate the development of new stateful calculi in the next section, which was not achieved by the previous work.

4 Locative calculi with explicit server states

The stateless server strategy in the state-encoding calculi is not always satisfactory. When the server states involve disks or databases that do not permit serialisation, it is difficult to encode them. Let us see an example:

$$\begin{aligned} \lambda^s query \cdot & \text{let } cursor = executeOnDatabase(query) \text{ in} \\ & \text{let } name = getNameFromRecord(cursor) \text{ in} \\ & \text{let } r = f_{client}(name) \text{ in} \\ & \text{let } cursor = nextRecord(cursor) \text{ in } \dots \end{aligned}$$

This is a server function that has a query as an argument. It is assumed to invoke a client function repeatedly for each of a list of records obtained from a database query. A remote client function invocation $f_{client}(name)$ resides between two subsequent database operations. We should serialise a list of records referenced by the database cursor before

the invocation. However, in general, it is difficult to pack the records into a continuation leaving nothing in the server. Even when it is possible to serialise such server states, moving the serialised records between the client and the server repeatedly will increase communication overheads. This will cause some efficiency concerns. This observation motivates us to design new stateful calculi that eliminate the need to encode server states.

In this section, we develop new calculi that save and restore server states using the stack in the underlying semantics, and therefore there is no need to encode them as in the state-encoding calculi. In fact, using the runtime stack is similar to how conventional programming languages have approached the problem of saving and restoring control contexts.

We first introduce the syntax and semantics for a stateful RPC calculus λ_{rpc}^{state} using stack-based control states in the server. Then we propose a typed compilation method of compiling λ_{rpc} terms into λ_{rpc}^{state} terms and subsequently another compilation method for separating client and server terms in λ_{cs}^{state} .

4.1 A stateful RPC calculus λ_{rpc}^{state}

Let us explore the features of a stateful RPC calculus λ_{rpc}^{state} before we explain the details of its syntax and semantics. An important feature is a stack Δ in the server:

$$\Delta ::= \varepsilon \mid \text{ctx } x M \cdot \Delta$$

where the empty stack is denoted by ε , and any non-empty stack consists of contexts concatenated with the dot operator. A context $\text{ctx } x M$ can be viewed as a let term with a hole such as $\text{let } x = [-]$ in M denoting a control state.

We illustrate two new constructs `call` and `ret` in the stateful RPC calculus with the following steps under the evaluation relation \Rightarrow^{state} :

	[Client]		[Server]
	$\text{ctx } z M_z$		$\Delta; \text{let } x = \text{call}(f_{client}, V_{arg}) \text{ in } M$ (1)
\Rightarrow^{state}	$\text{let } z = f_{client}(V_{arg}) \text{ in } M_z$		$\text{ctx } x M \cdot \Delta$ (2)
$\Rightarrow^{state} +$	$\text{let } z = \text{ret}(W) \text{ in } M_z$		$\text{ctx } x M \cdot \Delta$ (3)
\Rightarrow^{state}	$\text{ctx } z M_z$		$\Delta; \text{let } x = W \text{ in } M$ (4)

In Configuration (1), a server is about to switch to a client in order to invoke a client function. The new call construct does this after pushing the server context ($\text{ctx } x M$) onto a stack Δ . In Configuration (2), the waiting client context ($\text{ctx } z M_z$) is open to become the let with a hole filled with the client function application $f_{client}(V_{arg})$, and the server context stack grows from Δ to $\text{ctx } x M \cdot \Delta$. This application is assumed to evaluate eventually to $\text{ret}(W)$ using the other new construct. In Configuration (3), the new `ret` construct is about to switch to the server to return the value W after popping the top server context from the server stack and restoring the server context saved by `call` previously. In addition, the `ret` construct leaves the client context surrounding itself in the client. This will lead to Configuration (4).

Note that $\text{call}(f, \overline{W})$ in λ_{rpc}^{state} is different from $\text{call}(f, \overline{W})$ in λ_{rpc}^{enc} in regard to whether or not the server stack is involved. Accordingly, the organisation of client functions for calling by the constructs is different, as we will see soon.

The configuration ‘let $x = \text{ret}(V)$ in $M \mid \Delta$ ’ in the stateful RPC calculus is reminiscent of the configuration ‘let $x = \text{req}(K, V)$ in $M \mid \varepsilon$ ’ in the state-encoding RPC calculus. The server stack Δ in the former can be regarded as the continuation K encoded for a server control state in the latter. As we apply K to V in the server and bind the application result to x in the latter, we return V to the server with Δ and bind the evaluation result to x in the former.

Consider again the λ_{rpc} example term:

$$(\lambda^s f \cdot (\lambda^s x \cdot x) (f \ c)) (\lambda^c y \cdot (\lambda^s z \cdot z) \ y)$$

For comparison with the state-encoding RPC calculus, this term can be implemented in the stateful RPC calculus as

$$\text{let } r = \text{req}(M_1, M_2) \text{ in } r$$

where M_1 and M_2 are supposed to implement the outermost function $(\lambda^s f \cdot (\lambda^s x \cdot x) (f \ c))$ and its argument $(\lambda^c y \cdot (\lambda^s z \cdot z) \ y)$, respectively. The use of req is the same as for λ_{rpc}^{enc} . The notable difference is that it does not require any continuation argument. In λ_{rpc}^{state} , it is not necessary to use CPS conversion for server functions, because the underlying semantics of the stateful calculus explicitly manages a server stack that continuations would substitute for in the state-encoding calculus.

Now let us examine M_1 :

$$M_1 \triangleq \lambda^s f \cdot \text{let } w = \text{call}(\lambda^c x \cdot \text{let } y = f(x) \text{ in } \text{ret}(y)) \ , \ \llbracket c \rrbracket \text{ in} \\ \text{let } r = \llbracket \lambda^s x \cdot x \rrbracket(w) \text{ in } r$$

where $\llbracket M \rrbracket$ represents an implementation of a λ_{rpc} term M for the server. Given f as a client function, call switches to the client to invoke a commute function $(\lambda^c x \cdot \text{let } y = f(x) \text{ in } \text{ret}(y))$, pushing a server context $(\text{ctx } w \ (\text{let } r = \llbracket \lambda^s x \cdot x \rrbracket(w) \text{ in } r))$ on a server stack. The commute function applies f to x (i.e. $\llbracket c \rrbracket$) in the client. It then returns the result value y to the server by $\text{ret}(y)$, popping the server context. We continue to evaluate the term from the restored context with the result value bound to w in the server.

One implementation of M_2 is as follows:

$$M_2 \triangleq \lambda^c y \cdot \text{let } r = \text{req}(\llbracket \lambda^s z \cdot z \rrbracket, \ y) \text{ in } r$$

When compared with the implementation for λ_{rpc}^{enc} , this is a simple structure with no continuation argument in the use of req to invoke a server function.

4.2 The formal semantics of λ_{rpc}^{state}

Figure 10 shows the syntax and semantics of the stateful RPC calculus λ_{rpc}^{state} in the same style of the state-encoding RPC calculus. Basically, the syntax of the calculus has A-normal form with three constructs: req , call and ret . We intend this new construct $\text{ret}(V)$ to return a value to the server from the client, and a symmetric construct of $\text{call}(f, \overline{W})$ invokes a client function from the server.

Values are either a variable or a locative function. Client contexts Π are the same for λ_{rpc}^{enc} . Server contexts now have the form of a stack denoted by Δ .

Syntax

Term	$M ::= V \mid V_f(\overline{W}) \mid \text{req}(V_f, \overline{W}) \mid \text{call}(V_f, \overline{W}) \mid \text{ret}(V) \mid \text{let } x = M \text{ in } M$
Value	$V, W ::= x \mid \lambda^a \overline{x}. M$
Client context	$\Pi ::= \text{ctx } x M$
Server context stack	$\Delta ::= \varepsilon \mid \text{ctx } x M \cdot \Delta$

Semantics

Client:	$M \mid \Delta \Rightarrow^{state} M' \mid \Delta \quad \text{or} \quad M \mid \Delta \Rightarrow^{state} \Pi \mid \Delta'; M'$
(AppC)	$\text{let } y = (\lambda^c \overline{x}. M_0)(\overline{W}) \text{ in } M \mid \Delta \Rightarrow^{state} \text{let } y = M_0\{\overline{W}/\overline{x}\} \text{ in } M \mid \Delta$
(Req)	$\text{let } x = \text{req}(\lambda^s \overline{x}. M_0, \overline{W}) \text{ in } M \mid \Delta \Rightarrow^{state} \text{ctx } x M \mid \Delta; \text{let } r = (\lambda^s \overline{x}. M_0)(\overline{W}) \text{ in } r$
(ValC)	$\text{let } x = V \text{ in } M \mid \Delta \Rightarrow^{state} M\{V/x\} \mid \Delta$
(LetC)	$\text{let } x = (\text{let } y = M_1 \text{ in } M_2) \text{ in } M \mid \Delta \Rightarrow^{state} \text{let } y = M_1 \text{ in } (\text{let } x = M_2 \text{ in } M) \mid \Delta$
(Ret)	$\text{let } y = \text{ret}(V) \text{ in } M_2 \mid \text{ctx } x M_1 \cdot \Delta \Rightarrow^{state} \text{ctx } y M_2 \mid \Delta; \text{let } x = V \text{ in } M_1$
Server:	$\Pi \mid \Delta; M \Rightarrow^{state} \Pi \mid \Delta; M' \quad \text{or} \quad \Pi \mid \Delta; M \Rightarrow^{state} M' \mid \Delta'$
(AppS)	$\Pi \mid \Delta; \text{let } y = (\lambda^s \overline{x}. M_0)(\overline{W}) \text{ in } M \Rightarrow^{state} \Pi \mid \Delta; \text{let } y = M_0\{\overline{W}/\overline{x}\} \text{ in } M$
(Call)	$\text{ctx } y M_2 \mid \Delta; \text{let } x = \text{call}(\lambda^c \overline{x}. M_0, \overline{W}) \text{ in } M_1 \Rightarrow^{state} \text{let } y = (\lambda^c \overline{x}. M_0)(\overline{W}) \text{ in } M_2 \mid \text{ctx } x M_1 \cdot \Delta$
(Reply)	$\text{ctx } x M \mid \Delta; V \Rightarrow^{state} \text{let } x = V \text{ in } M \mid \Delta$
(ValS)	$\Pi \mid \Delta; \text{let } x = V \text{ in } M \Rightarrow^{state} \Pi \mid \Delta; M\{V/x\}$
(LetS)	$\Pi \mid \Delta; \text{let } x = (\text{let } y = M_1 \text{ in } M_2) \text{ in } M \Rightarrow^{state} \Pi \mid \Delta; \text{let } y = M_1 \text{ in } (\text{let } x = M_2 \text{ in } M)$

Fig. 10. A stateful client–server calculus λ_{rpc}^{state} .

There are two kinds of configurations in the stateful RPC calculus:

- $M \mid \Delta$ for the client to evaluate a term M with a stack Δ in the server separated by |
- $\Pi \mid \Delta; M$ for the server to evaluate a term M on a stack Δ with a pending client context Π in the client separated by |

The semantics of the calculus is described by the evaluation relation \Rightarrow^{state} on the configurations. The semantic rules are defined in a minimal way for λ_{rpc}^{state} enough to be an intermediate calculus for compilation of λ_{rpc} as those for λ_{rpc}^{enc} . In the semantics, (AppC), (ValC) and (LetC) are semantic rules for locally running the client in the form $M \mid \Delta \Rightarrow^{state} M' \mid \Delta$. Note that the three client-side rules are different from those rules in λ_{rpc}^{enc} only in that the server context stack Δ replaces the empty server context of the rules in λ_{rpc}^{enc} . (AppS), (ValS) and (LetS) are semantic rules for locally running the server in the form $\Pi \mid \Delta; M \Rightarrow^{state} \Pi \mid \Delta; M'$. The three server-side rules introduce the server context stack Δ when compared with those rules in λ_{rpc}^{enc} .

The remaining four semantic rules, (Req), (Ret), (Call) and (Reply), change active running from the client to the server or vice versa. We intend a request initiated by (Req) to finish by (Reply), and we also intend a call that (Call) begins to finish with (Ret). Note that both (Req) and (Reply) are exactly the same as those in λ_{rpc}^{enc} except for the introduction of a server context stack Δ instead of the empty server context.

(Call) in λ_{rpc}^{state} exhibits an important difference from that in λ_{rpc}^{enc} . The call construct can now be used in a non-tail position, e.g. $\text{let } y = [-] \text{ in } M_0$. (Call) pushes it onto the current server context stack Δ as $\text{ctx } y M_0 \cdot \Delta$, switching to the client. (Ret) reverses this procedure: it returns a value computed in the client back to the server, popping the top server context from the server stack $\text{ctx } y M_0 \cdot \Delta$.

Client: $C[[M_{rpc}]] = M_{rpc}^{state}$	
(VarC)	$C[[x]] = x$
(LamCC)	$C[[\lambda^c x.M]] = \lambda^c x.C[[M]]$
(LamCS)	$C[[\lambda^s x.M]] = \lambda^s x.S[[M]]$
(AppCC)	$C[[L^c M]] = \text{let } f = C[[L]] \text{ in}$ $\text{let } x = C[[M]] \text{ in}$ $\text{let } r = f(x) \text{ in } r$
(AppCS)	$C[[L^s M]] = \text{let } f = C[[L]] \text{ in}$ $\text{let } x = C[[M]] \text{ in}$ $\text{let } r = \text{req}(f, x) \text{ in } r$
Server: $S[[M_{rpc}]] = M_{rpc}^{state}$	
(VarS)	$S[[x]] = x$
(LamSC)	$S[[\lambda^c x.M]] = \lambda^c x.C[[M]]$
(LamSS)	$S[[\lambda^s x.M]] = \lambda^s x.S[[M]]$
(AppSC)	$S[[L^c M]] = \text{let } f = S[[L]] \text{ in}$ $\text{let } x = S[[M]] \text{ in}$ $\text{let } r = \text{call}(\lambda^c x.\text{let } y = f(x) \text{ in } \text{ret}(y), x) \text{ in } r$
(AppSS)	$S[[L^s M]] = \text{let } f = S[[L]] \text{ in}$ $\text{let } x = S[[M]] \text{ in}$ $\text{let } r = f(x) \text{ in } r$

Fig. 11. A typed compilation for λ_{rpc}^{state} .

Interestingly, we find that the client–server interactions in λ_{rpc}^{state} are analogous to the behaviour of coroutines (Conway, 1963). When the client is assumed to be an initiator and the server is assumed to be a coroutine, (Req) is an initiation of a coroutine, (Call) is a suspension of the execution of the coroutine to return to its initiator, (Ret) is a resumption of the coroutine and (Reply) is a terminator of the coroutine.

The stateful RPC calculus is also based on the semantic-based trampoline style, which was explained for the state-encoding calculi in Section 3.5, by (Req), (Call) and (Reply) together with the support for composing commute functions by its compiler, which will be explained in the next section.

4.3 A typed compilation of λ_{rpc} for λ_{rpc}^{state}

Now we present a typed compilation of the RPC calculus into the stateful RPC calculus. Our compilation takes a typing derivation of a source term under our locative type system in the same way as before.

Figure 11 shows our compilation rules. The compilation rules for the client and the server have the same form:

$$C[[M_{rpc}]] = M_{rpc}^{state} \quad \text{and} \quad S[[M_{rpc}]] = M_{rpc}^{state}$$

The compilation rules are in direct style. This contrasts with the use of CPS conversion in the compilation rules for the state-encoding RPC calculus, which does not support saving and restoring server contexts in its semantics. Therefore, it is a compiler that supports this by CPS conversion. The stateful calculus, however, does support saving and restoring server contexts in the underlying semantics, and so the compiler for this calculus does not need to track any server contexts to manage. As a result, it is not necessary to adopt CPS conversion for compilation of server terms anymore.

In the compilation rules, (VarC), (LamCC) and (LamCS) are the same as (VarS), (LamSC) and (LamSS) for variable, client function and server function. The compilation rules only follow the structure of the terms.

(AppCC) and (AppSS) are for local applications in the client and server, respectively. They generate local application terms $f(x)$ with f and x as the compiled terms of the function and the argument of the applications.

(AppCS) and (AppSC) are the most interesting compilation rules. (AppCS) compiles a term of invocation of a server function in the client. It simply generates a target term with $\text{req}(f, x)$ with f as a compiled server function and x as a compiled argument. Note that there is no identity continuation once introduced in compilation for the state-encoding RPC calculus. (AppSC) compiles a term of invocation of a client function in the server, into a target term with $\text{call}(\lambda^c x \cdot \text{let } y = f(x) \text{ in } \text{ret}(y), x)$ composing a commute function with call and ret .

We now prove a slightly stronger version of the correctness of the compilation for λ_{rpc}^{state} . For the proof, we introduce a definition of *call-return balanced* evaluation steps. call and ret are the only constructs to push onto and to pop from a server stack. Intuitively, this implies that in the evaluation steps, every invocation of a client function from the server will return to the server without changing a client context and a server stack.

Definition 4.1. A sequence of evaluation steps, $\text{Conf}_1 \Rightarrow^{state*} \text{Conf}_n$, is said to be *call-return balanced* if the evaluation sequence is derived from the following grammar:

- $\text{bal} ::= \text{Conf}_{\text{call}} \text{ bal } \text{Conf}_{\text{ret}} \mid \text{Conf}_{\text{else}} \mid \text{bal}_1 \text{ bal}_2$ where
 - $\text{Conf}_{\text{call}}$ and Conf_{ret} : configurations paired in the form of
 - (1) $\text{ctx } y \text{ } M_y \mid \Delta; \text{ let } x = \text{call}(V\bar{W}) \text{ in } M_x$ and
 - (2) $\text{let } y = \text{ret}(V_r) \text{ in } M_y \mid \text{ctx } x \text{ } M_x \cdot \Delta$, respectively.
 - $\text{Conf}_{\text{else}}$: configurations in the form of neither $\text{Conf}_{\text{call}}$ nor Conf_{ret}

Note that it is easy to construct $\text{Conf}_{\text{call}} \Rightarrow^{state*} \text{ctx } y \text{ } M_y \mid \Delta; \text{ let } x = V_r \text{ in } M_x$ in the call-return balanced evaluation steps, because $\text{Conf}_{\text{ret}} \Rightarrow^{state} \text{ctx } y \text{ } M_y \mid \Delta; \text{ let } x = V_r \text{ in } M_x$ by the definition of (Ret). The client context $\text{ctx } y \text{ } M_y$ and the server stack Δ are preserved during the client function call.

Theorem 4.1 (Correctness of Compilation for λ_{rpc}^{state}). Assume a well-typed term $\lambda_{rpc} M$ under the locative type system:

- If $M \Downarrow_c V$, then $C[M] \mid \Delta \Rightarrow^{state*} C[V] \mid \Delta$ for all Δ , which is call-return balanced.
- If $M \Downarrow_s V$, then $\Pi \mid \Delta; S[M] \Rightarrow^{state*} \Pi \mid \Delta; S[V]$ for all Π and Δ , which is call-return balanced.

The basic idea of our proof has the same structure as the proof for λ_{rpc}^{enc} explained in Section 3.3. The proof builds a sequence of evaluation steps in λ_{rpc}^{state} matched to each subtree of evaluation in λ_{rpc} by induction on the height of evaluation trees, and then it puts the sequences together to build what is matched to the whole evaluation tree in the conditions. The proof therefore guarantees that client–server communication in λ_{rpc} is preserved in λ_{rpc}^{enc} , though the theorem itself does not tell this explicitly. Later, we will

Syntax

Term	$m ::= v \mid v_f(\bar{w}) \mid \text{req}(v_f, \bar{w}) \mid \text{call}(v_f, \bar{w}) \mid \text{ret}(v) \mid \text{let } x = m \text{ in } m$
Value	$v, w ::= x \mid \text{clo}(F, \bar{v})$
Client context	$\pi ::= \text{ctx } x m$
Server context	$\delta ::= \varepsilon \mid \text{ctx } x m \cdot \delta$

Semantics

Client:	$m \mid \delta \Rightarrow^{state} m' \mid \delta \quad \text{or} \quad m \mid \delta \Rightarrow^{state} \pi \mid \delta'; m'$
---------	---

(AppC) $\text{let } y = v_F(\bar{w}) \text{ in } m \mid \delta \Rightarrow^{state} \text{let } y = m_0\{\bar{v}/\bar{z}, \bar{w}/\bar{x}\} \text{ in } m \mid \delta$
if $v_F = \text{clo}(F, \bar{v})$ and $\phi_c(F) = \bar{z}\lambda^c \bar{x}.m_0$

(Req) $\text{let } x = \text{req}(v_f, \bar{w}) \text{ in } m \mid \delta \Rightarrow^{state} \text{ctx } x m \mid \delta; \text{let } r = v_f(\bar{w}) \text{ in } r$
if $v_f = \text{clo}(F, \bar{v})$ and $\phi_s(F) = \bar{z}\lambda^s \bar{x}.m_0$

(ValC) $\text{let } x = v \text{ in } m \mid \delta \Rightarrow^{state} m\{v/x\} \mid \delta$

(LetC) $\text{let } x = (\text{let } y = m_1 \text{ in } m_2) \text{ in } m \mid \delta \Rightarrow^{state} \text{let } y = m_1 \text{ in } (\text{let } x = m_2 \text{ in } m) \mid \delta$

(Ret) $\text{let } y = \text{ret}(v) \text{ in } m_2 \mid \text{ctx } x m_1 \cdot \delta \Rightarrow^{state} \text{ctx } x m_2 \mid \delta; \text{let } x = v \text{ in } m_1$

Server:	$\pi \mid \delta; m \Rightarrow^{state} \pi \mid \delta; m' \quad \text{or} \quad \pi \mid \delta; m \Rightarrow^{state} m' \mid \delta'$
---------	---

(AppS) $\pi \mid \delta; \text{let } y = v_F(\bar{w}) \text{ in } m \Rightarrow^{state} \pi \mid \delta; \text{let } y = m_0\{\bar{v}/\bar{z}, \bar{w}/\bar{x}\} \text{ in } m$
if $v_F = \text{clo}(F, \bar{v})$ and $\phi_s(F) = \bar{z}\lambda^s \bar{x}.m_0$

(Call) $\text{ctx } x m_2 \mid \delta; \text{let } x = \text{call}(v_f, \bar{w}) \text{ in } m_1 \Rightarrow^{state} \text{let } y = v_f(\bar{w}) \text{ in } m_2 \mid \text{ctx } x m_1 \cdot \delta$
if $v_f = \text{clo}(F, \bar{v})$ and $\phi_c(F) = \bar{z}\lambda^c \bar{x}.m_0$

(Reply) $\text{ctx } x m \mid \delta; v \Rightarrow^{state} \text{let } x = v \text{ in } m \mid \delta$

(ValS) $\pi \mid \delta; \text{let } x = v \text{ in } m \Rightarrow^{state} \pi \mid \delta; m\{v/x\}$

(LetS) $\pi \mid \delta; \text{let } x = (\text{let } y = m_1 \text{ in } m_2) \text{ in } m \Rightarrow^{state} \pi \mid \delta; \text{let } y = m_1 \text{ in } (\text{let } x = m_2 \text{ in } m)$

Compilation

$CC[[M_{rpc}^{state}]] = m$

$CC[[\text{ret}(v)]] = \text{ret}(CC[[v]])$ (The rest of the compilation rules are the same as for λ_{cs}^{enc})

Fig. 12. Syntax and semantics of λ_{cs}^{state} with a compilation of λ_{rpc}^{state} .

present an example of a correspondence of client–server communication between λ_{rpc} and λ_{rpc}^{state} indirectly by comparing related flows between λ_{rpc}^{enc} and λ_{rpc}^{state} in Section 4.5.1.

4.4 Separating client and server terms in λ_{cs}^{state}

As was done for the state-encoding calculi, we propose a stateful CS calculus. Again, the basic idea is to use closure conversion to decompose a term into closed functions and to classify them according to location annotations.

In Figure 12, the syntax and semantics of λ_{cs}^{state} are shown. The syntax of λ_{cs}^{state} is identical to that of λ_{cs}^{enc} except for $\text{ret}(v)$. The terms in the calculus are denoted by m . The calculus explicitly deals with closures v as values, rather than lambda terms. F is a name for a closed function. ϕ_a is a mapping of names into closed functions in the location a . π and δ are client and server contexts, respectively.

Configurations in the semantics of λ_{cs}^{state} are changed accordingly:

- $m \mid \delta$ for the client to evaluate a term m with a stack δ in the server separated by |
- $\pi \mid \delta; m$ for the server to evaluate a term m on a stack δ with a pending client context π in the client separated by |

The semantic rules for the stateful CS calculus are isomorphic to those for the stateful RPC calculus under the aforementioned changes.

$$\begin{aligned}
 \phi_c & : \text{main} = \text{let } r_3 = \text{req}(\text{clo}(g_3, \{\}), \text{clo}(g_5, \{\})) \text{ in } r_3 \\
 & \quad g_2 = \{f_7\} \lambda^c z_{10}. \text{let } y_9 = f_7 z_{10} \text{ in } \text{ret}(y_9) \\
 & \quad g_5 = \{\} \lambda^c y. \text{let } r_{14} = \text{req}(\text{clo}(g_4, \{\}), y) \text{ in } r_{14} \\
 \phi_s & : g_1 = \{\} \lambda^s x. x \\
 & \quad g_3 = \{\} \lambda^s f. \text{let } x_5 = (\text{let } r_{11} = \text{call}(\text{clo}(g_2, \{f\}), c) \text{ in } r_{11}) \text{ in } \text{let } r_6 = \text{clo}(g_1, \{\}) x_5 \text{ in } r_6 \\
 & \quad g_4 = \{\} \lambda^s z. z
 \end{aligned}$$

Fig. 13. A compilation of the RPC term in Figure 3.

The compilation rules for λ_{cs}^{state} are the same as those for λ_{cs}^{enc} except for $\text{ret}(v)$ in Figure 12. For a given closed λ_{rpc}^{state} term M , $CC\llbracket M \rrbracket = m$ means that it generates a λ_{cs}^{state} term m for the client, producing a function store ϕ_c for the client and another function store ϕ_s for the server.

We now prove the correctness of the compilation for λ_{cs}^{state} by assuming a similar definition of call–return balance for evaluation steps as Definition 4.1.

Theorem 4.2 (Correctness of Compilation for λ_{cs}^{state}). *Assume a well-typed term M under the locative type system. Given ϕ_c and ϕ_s :*

- If $M \Downarrow_c V$, then $CC\llbracket C\llbracket M \rrbracket \rrbracket \mid \varepsilon \Rightarrow^{state*} CC\llbracket C\llbracket V \rrbracket \rrbracket \mid \varepsilon$, which is call–return balanced.

While we reused the parser of λ_{rpc} and the location-type inference procedure for our type system mentioned in Section 3.4, we have implemented a compiler and an evaluator of λ_{rpc}^{state} and λ_{cs}^{state} on HTTP between the client and the server. Figure 13 shows a compilation of the RPC term in Figure 3, generated by the compiler. The evaluator starts with *main* in the client.

4.5 Discussion

4.5.1 Session management

Because the RPC calculus is designed for the web-based client–server model, it is necessary to discuss whether or not the new stateful CS calculus fits the model. In the state-encoding CS calculus, it is rather straightforward to associate the semantics with the *request–response* interaction of the web-based model, thanks to the nature of the *stateless* server which is aligned well with the RESTful web. However, in the stateful CS calculus, it is not so obvious how the semantics is connected with the web-based interaction, because the web-based model does not directly support the server states that the semantics deals with.

To explain how the stateful CS calculus manages server states through a series of multiple request–response interactions, we introduce the notion of *session* to capture one or more request–response interactions. In the notation, we place a session annotation over the bar separating a client and a server as

- ‘ $\overset{session}{client} \mid server$ ’ to denote the client and the server connected under a session identified by a unique number *session*, or
- ‘ $\overset{nothing}{client} \mid server$ ’ to denote that no session is established between the two.

Note that this notion also turns out to be useful for explicitly explaining how the semantics in the state-encoding CS calculus is related to request–response interactions.

Let us begin with the semantics of λ_{cs}^{enc} extended with session annotations where one session exactly corresponds to one request–response interaction on the web.

$$\begin{aligned}
 (\text{Req}) \quad & \text{let } x = \text{req}(v_F, \bar{w}) \text{ in } m \mid \varepsilon \xRightarrow{enc} \text{ctx } x \ m \mid \text{let } r = v_F(\bar{w}) \text{ in } r \\
 & \text{where } v_F = \text{clo}(F, \bar{v}), \phi_s(F) = \bar{z}\lambda^s \bar{x} \cdot m_0, \text{ and fresh session} \\
 (\text{Call}) \quad & \text{ctx } x \ m \mid \text{call}(v_F, \bar{w}) \xRightarrow{enc} \text{let } x = v_F(\bar{w}) \text{ in } m \mid \varepsilon \\
 & \text{where } v_F = \text{clo}(F, \bar{v}), \phi_c(F) = \bar{z}\lambda^c \bar{x} \cdot m_0 \\
 (\text{Reply}) \quad & \text{ctx } x \ m \mid v \xRightarrow{enc} \text{let } x = v \text{ in } m \mid \varepsilon
 \end{aligned}$$

(Req) corresponds to the creation of a new session. Both (Call) and (Reply) correspond to the closing of the session. For the other local rules, the status of the session remains unchanged:

$$\begin{aligned}
 \text{Client (AppC), (ValC), (LetC)} & : m \mid \varepsilon \xRightarrow{enc} m' \mid \varepsilon \\
 \text{Server (AppS), (ValS), (LetS)} & : \pi \mid m \xRightarrow{enc} \pi \mid m'
 \end{aligned}$$

where m, m' and π are those from the semantic rules in λ_{cs}^{enc} .

Now let us discuss the semantics of λ_{cs}^{state} extended with session annotations where one session can not only correspond to one request–response interaction on the web but can also be expanded to correspond to more than one interaction:

$$\begin{aligned}
 (\text{Req1}) \quad & \text{let } x = \text{req}(v_F, \bar{w}) \text{ in } m \mid \varepsilon \xRightarrow{state} \text{ctx } x \ m \mid \varepsilon; \text{let } r = v_F(\bar{w}) \text{ in } r \\
 & \text{where } v_F = \text{clo}(F, \bar{v}), \phi_s(F) = \bar{z}\lambda^s \bar{x} \cdot m_0, \text{ and fresh session} \\
 (\text{Req2}) \quad & \text{let } x = \text{req}(v_F, \bar{w}) \text{ in } m \mid \delta \xRightarrow{state} \text{ctx } x \ m \mid \delta; \text{let } r = v_F(\bar{w}) \text{ in } r \\
 & \text{where } v_F = \text{clo}(F, \bar{v}), \phi_s(F) = \bar{z}\lambda^s \bar{x} \cdot m_0, \text{ and } \delta \neq \varepsilon \\
 (\text{Reply1}) \quad & \text{ctx } x \ m \mid \varepsilon; v \xRightarrow{state} \text{let } x = v \text{ in } m \mid \varepsilon \\
 (\text{Reply2}) \quad & \text{ctx } x \ m \mid \delta; v \xRightarrow{state} \text{let } x = v \text{ in } m \mid \delta \text{ if } \delta \neq \varepsilon \\
 (\text{Call}) \quad & \text{ctx } y \ m_2 \mid \delta; \text{let } x = \text{call}(v_F, \bar{w}) \text{ in } m_1 \xRightarrow{state} \text{let } y = v_F(\bar{w}) \text{ in } m_2 \mid \text{ctx } x \ m_1 \cdot \delta \\
 & \text{where } v_F = \text{clo}(F, \bar{v}) \text{ and } \phi_c(F) = \bar{z}\lambda^c \bar{x} \cdot m_0 \\
 (\text{Ret}) \quad & \text{let } y = \text{ret}(v) \text{ in } m_2 \mid \text{ctx } x \ m_1 \cdot \delta \xRightarrow{state} \text{ctx } y \ m_2 \mid \delta; \text{let } x = v \text{ in } m_1
 \end{aligned}$$

The original semantic rule (Req) in λ_{cs}^{state} is divided into two rules. (Req1) with the empty server stack corresponds to the creation of a new session, and (Req2) with the non-empty stack δ corresponds to a continuation of the previously created session. The original semantic rule (Reply) in λ_{cs}^{state} is also divided into two rules. (Reply1) with the empty server stack is matched with (Req1) to close the created session, and (Reply2) with the non-empty stack δ is matched with (Req2) to maintain the previously created session. Both (Call) and (Ret) are extended to continue to maintain the status of the session. For the other local rules, the status of session remains unchanged as

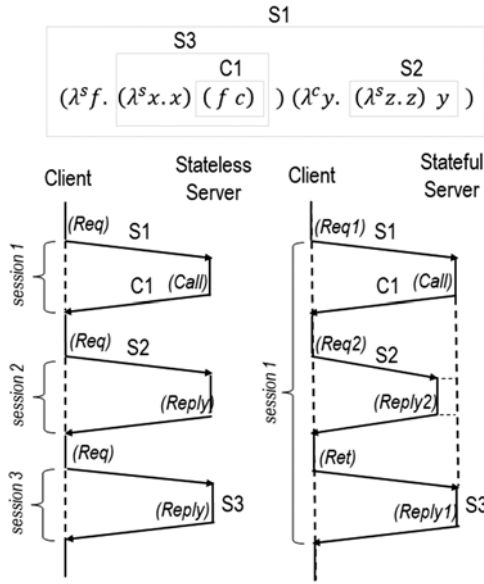


Fig. 14. Comparison of session management between λ_{cs}^{enc} and λ_{cs}^{state} .

$$\begin{aligned}
 \text{Client } (\text{AppC}), (\text{ValC}), (\text{LetC}) &: m \mid \delta \xRightarrow{enc} m' \mid \delta \\
 \text{Server } (\text{AppS}), (\text{ValS}), (\text{LetS}) &: \pi \mid \delta; m \xRightarrow{enc} \pi \mid \delta; m'
 \end{aligned}$$

where m, m', π and δ are those from the semantic rules in λ_{cs}^{state} . For the client-side local rules, it is confirmed that either $optSession = nothing$ and $\delta = \varepsilon$ or $optSession = session$ and $\delta \neq \varepsilon$ holds.

The session-annotated semantics of λ_{cs}^{state} clearly explains how the stateful CS calculus using request–response interactions is implemented on the web-based model. Also, the idea of applying session annotation to the semantics enables us compare λ_{cs}^{enc} with λ_{cs}^{state} in terms of the web-based implementation using request–response interaction.

Figure 14 shows a comparison of session management between λ_{rpc}^{enc} and λ_{rpc}^{state} when the λ_{rpc} term in the example is evaluated. First, the figure shows that client–server communication in λ_{rpc}^{state} is the same as that in λ_{rpc}^{enc} by using the same box names S1, C1 and S2 on the related flows. Second, with the stateful server, the active thread of control described by the solid line with (S2) is above the pending thread of control described by the dashed line with (S3), meaning that S2 is being evaluated with the server stack holding the server context of S3.

4.5.2 A mixed strategy

In the previous comparison between λ_{cs}^{enc} and λ_{cs}^{state} , we clearly understood how they manage sessions differently. In λ_{cs}^{enc} , the server does not need to retain anything once it sends the client a response. This is advantageous because the server resource consumption is reduced. In λ_{cs}^{state} , the server can naturally support storing a stateful context that the server must have before and after a client function call. This is necessary when it is not easy or

not possible to encode the stateful context, for example, in database operations as seen in Section 4, and to send the encoded context to the client. The question then arises if we can use both strategies: a mixed strategy basically employs the state-encoding strategy to reduce the server resource consumption but switches to the stateful strategy when necessary. We here argue that our theory can be extended to support this mixed strategy.

For the mixed strategy, the problem is how to distinguish one phase of λ_{rpc} terms in which the compilation rules and semantic rules of λ_{rpc}^{enc} and λ_{cs}^{enc} are applied, from the other phase of the terms in which those of λ_{rpc}^{state} and λ_{cs}^{state} are applied. A simple idea is to use the notion of monadic encapsulation of state (Launchbury & Peyton Jones, 1994; Timany *et al.*, 2017), where stateful computations are encapsulated using monads and the purity of the functional language embracing them is preserved.

To explain how monadic encapsulation of state can help in the design of our mixed strategy, we first introduce the standard monadic operations and stateful operations to the term syntax of λ_{rpc} , and a new (state) type is introduced to the type syntax:

$$\begin{aligned} M & ::= \dots \mid runST\ M \mid thenST\ M\ (\lambda x \cdot M) \mid returnST\ M \\ & \quad \mid newVar\ M \mid readVar\ M \mid writeVar\ M\ M \\ \tau & ::= \dots \mid ST\ state\ \tau \mid \forall state \cdot \tau \end{aligned}$$

A key construct $runST\ M$ encapsulates a stateful computation M where states may be cascaded through $newVar$, $readVar$ and $writeVar$ by the monadic binder $thenST$, which allows a value to escape from the monad by the monadic stopper $returnST$ at the end. Note that $state$ is a (state) type variable, which is essential to the monadic encapsulation of state in a pure functional language. A new state type $ST\ state\ \tau$ can be read intuitively as follows: a stateful computation of this type can allocate, read and write in a memory region indexed by $state$ and then produce a value of type τ . The type of $runST$ in the literature (Launchbury & Peyton Jones, 1994; Timany *et al.*, 2017) is

$$(\forall state \cdot ST\ state\ \tau) \rightarrow \tau.$$

Having a universally quantified state type variable in the type guarantees that every memory region allocated inside the stateful computation is not accessible outside the $runST$ anymore by limiting the scope of the state type variable, which never occurs in τ . This is a key property of the monadic encapsulation of state.

For example, in an evaluation of the following extended λ_{rpc} term:

$$\begin{aligned} & clientFun2\ (\\ & \quad runST\ (\\ & \quad \quad (thenST\ (newVar\ 0)\ (\lambda l. \\ & \quad \quad \quad (thenST\ (writeVar\ l\ (clientFun1()))\ (\lambda(). \\ & \quad \quad \quad \quad (thenST\ (readVar\ l)\ (\lambda x. \\ & \quad \quad \quad \quad \quad (returnST\ x))))))))) \) \) \end{aligned}$$

a new piece of memory is allocated at an address l and is initialised with 0 by $newVar\ 0$, the server calls a client function to obtain a value by $clientFun1()$, the memory is updated with the value by $writeVar\ l\ (\dots)$, the written value is read again by $readVar\ l$, and it escapes the $runST$ monad by $returnST\ x$. Finally, the other client function $clientFun2$ is called from the server with the escaped value as an argument.

It is important to note that *clientFun1* is called in the middle of the stateful computations, whereas *clientFun2* is called after all the stateful computations. Therefore, the former should be implemented in the stateful server strategy, but the latter can be implemented in the state-encoding server strategy. To capture these different phases, we introduce two forms of typing judgements as follows:

$$\Gamma \triangleright_a M : \tau \quad \Gamma \blacktriangleright_a M : \tau$$

where the typing judgements with the black triangle denote that a term under evaluation is in the middle of stateful computations, and the typing judgements with the white triangle denote the absence of stateful computations. We also change the function type to include the colour as $\tau \xrightarrow{a, colour} \tau'$, where it is white or black. Functions of the black function type are invoked in the middle of the stateful computations, whereas functions of the white function type are never invoked there.

The typing rules for λ_{rpc} are then duplicated, and therefore we make one set have the typing rules with only the white triangle and make the other set have them with only the black triangle. In (T-App-B) of the black set, the colour of the function type in the premise is made the same as the colour of the triangle. In (T-Lam-B), the colour of the function type in the conclusion is the same as the colour of the triangle. These coloured typing rules will *get a function infected* whenever the function is invoked within a stateful expression.

$$(T\text{-App-B}) \frac{\Gamma \blacktriangleright_a L : \tau \xrightarrow{a, black} \tau' \quad \Gamma \blacktriangleright_a M : \tau}{\Gamma \blacktriangleright_a LM : \tau'} \quad (T\text{-Lam-B}) \frac{\Gamma \{x : \tau\} \blacktriangleright_b M : \tau'}{\Gamma \blacktriangleright_a \lambda^b x.M : \tau \xrightarrow{b, black} \tau'}$$

The two sets of the typing rules are combined through two new typing rules:

$$(Gen\&Run) \frac{fresh\ state, \ state \notin FTV(\Gamma) \cup FTV(\tau) \quad \Gamma \blacktriangleright_s M : ST\ state\ \tau}{\Gamma \blacktriangleright_s runST\ M : \tau}$$

$$(Purify) \frac{FTV(\Gamma) \cup FTV(\tau) = \emptyset \quad \Gamma \blacktriangleright_s runST\ M : \tau}{\Gamma \triangleright_s runST\ M : \tau}$$

(Gen&Run) is a typing rule combining (Gen) and (Run) in the literature, which generalises a state type variable and is a typing rule for *runST*, respectively. Those typing rules for the monadic constructs and state manipulation constructs are from the literature, and they are included in the black set. (Purify) is a typing rule that encapsulates a stateful world from pure functional expressions whenever no free state type variables escape, i.e. $FTV(\Gamma) \cup FTV(\tau) = \emptyset$.

With these typing rules, *clientFun1* in the example must be typed by (T-App-B); it *gets infected* because it is invoked within a *runST* expression, which is typed by (Gen&Run). However, *clientFun2* can be typed by (T-App-W), which is a white-coloured version of (T-App-B). The scope of states of the *runST* expression is thus limited inside it by (Gen&Run). Then it can be treated as a pure expression by (Purify) because there are no other *runST* expressions enclosing it.

After typing an extended λ_{rpc} term with the coloured typing rules, the phases using the white typing rules are compiled and evaluated under the state-encoding server strategy, and the phases using the black typing rules are compiled and evaluated under the stateful

server strategy. Monadic encapsulation of state has been proved to separate the stateful world from the pure expression in the literature (Launchbury & Peyton Jones, 1994; Timany *et al.*, 2017). Therefore, we can utilise the idea to separate where we can use the state-encoding server strategy from where we should use the stateful server strategy. What remains to be seen is how to compile (Purify). This typing rule is compiled as a direct style invocation of a CPS-style expression compiled from the typing in the premise by applying the initial continuation to the expression. This can be viewed as a border between the two strategies.

This completes the explanation of how we designed a hybrid strategy, which is another benefit of our theory of RPC calculi.

5 Related Work

Links (Cooper *et al.*, 2007) is a practical multi-tier web programming language that employs the RPC calculus as the foundation for client–server communication. It is a typed functional programming language incorporating Kleisli-based database query optimisation, continuations for web interactions, concurrency with message passing and XML programming. It would be interesting to extend the implementation of Links using the proposed theory of RPC calculi.

Lambda5 (Murphy VII *et al.*, 2004; Murphy, 2008) is a modally typed lambda calculus in which modal type systems can control local resources safely in distributed systems. It is extended to a full-fledged web programming language. From the programmer’s perspective, there are two things to compare. First, while Lambda5 is similar to the RPC calculus in that it allows programmers to write symmetric communication, it distinguishes the syntax of remote function applications, say, ‘get(s)(here $M N$)’, from that of local ones, say, ‘ $M N$ ’. In this respect, Lambda5 is like λ_{rpc}^{enc} and λ_{rpc}^{state} because λ_{rpc} offers a uniform syntax for both kinds of applications.

Second, the modal type system for Lambda5 introduces at modality, ‘ τ at a ’ in order to be able to refer to a specific location. The use of at-types for some term M and typing environment Γ can be viewed as typing judgements $\Gamma \triangleright_a M : \tau$. One interesting difference is that our type system introduces locative types $\tau \xrightarrow{a} \tau'$ with specific location a . The type system for Lambda5 uses box types $\Box\tau$ and diamond types $\Diamond\tau$, but locations in the modal types are too implicit to be usable; box types are associated with all arbitrary locations, and diamond types are associated with some unknown location. On the basis of this observation, we believe that our type system is a modal type system using locative types that internalise at modality. A formal comparison on how our type system is related to the modal type system is left as a future work. Another interesting question is to investigate if the state-encoding and stateful calculi could have any modal type systems beyond the type system for the RPC calculus.

From the implementation perspective, Lambda5 designed for peer-to-peer distributed systems is based on symmetric communication, whereas the RPC calculus designed for client–server systems is based on asymmetric communication. The semantic rules for client and server in Lambda5 are therefore the same. Although Lambda5 could be exploited for a stateless server strategy, the same semantic rules have to be revised in a nontrivial way

to achieve a stateful client and stateless server, where each is defined by different semantic rules. For a stateful server strategy, the Lambda5 semantic rules should also be revised to make them usable with the client–server model. The idea used in Lambda5 of continuations spanning multiple worlds such as client and server is similar. However, in λ_{rpc}^{state} , the semantic rules (Req) and (Call) for client-to-server RPCs and vice versa are asymmetric; the former is defined to initiate a new server control context, and the latter is defined to return to a client context waiting for this server response. This fits the implementation based on trampolined-style well. Thus, the semantics of Lambda5 has this subtle gap between its description and its web-based implementation. Lambda5 is implemented on a web-based client–server model in an ad hoc way, for example to support channels on top of HTTP.

A multi-tier calculus by Neubauer and Thiemann (Neubauer & Thiemann, 2005) automatically constructs communications for concurrently running processes. The calculus employs session types (Gay & Hole, 1999) to enforce the integrity of communications automatically. They proposed a series of transformations as compilation to convert a source program into separate programs at different locations determined by the use of primitives that run only at specific locations. The calculus can be called a *stateful peer strategy* because each process manages its own state, but there is no concept of a central server as in Lambda5.

There are several programming languages that have adopted the feature of multi-tier web programming: Hop (Serrano *et al.*, 2006; Serrano & Berry, 2012) extending Scheme; Hop.js extending JavaScript (Serrano & Prunet, 2016); Eliom, a multi-tier ML programming language, featuring module systems extended with location annotations (Radanne, 2017) in the project Ocsigen (Balat, 2006); and Ur/Web (Chlipala, 2015) with a dependently typed system; and a multi-tier functional reactive programming framework for Scala (Reynders *et al.*, 2014). They are all practical web programming languages featuring concurrency, reactive UI, database and XML.

Hop, Eliom and Ur/Web provide programmers only asymmetric communication where basically only the client can invoke server functions; the server resorts to a network library to call client functions back. All of them use a special syntax for RPC distinct from that for local function calls. As far as we know, they are all based on a stateful server approach as described for HOP in Serrano & Queinnec (2010) and Boudol *et al.* (2012) and for Eliom in Radanne (2017), though Ur/Web never presented any formal semantics on their RPC implementation. These languages do not need to implement the mapping of symmetric communication onto asymmetric communication as is done in our theory. It is therefore less meaningful to compare our theory with them in terms of stateful server approaches. In Serrano & Queinnec (2010), a denotational continuation-based semantics for a core subset of HOP is present for the purpose of elaboration of how to generate client-side code by the server-side code. In Boudol *et al.* (2012), more versatile semantics is proposed in order to reason about the safety of the same-origin policy in the web application security model, whereas we use our semantics for proving the correctness of compilation for the two RPC implementation strategies.

The PLT Scheme community has studied the construction of interactive server-side web applications. In Graunke *et al.* (2001) and Matthews *et al.* (2004), the authors have pointed

out the mismatch problem between the structure of *sequential* programs and the structure of the corresponding web interactions. They suggest an automatic structuring method of sequential programs so that a restructured web program and a consumer can construct a pair of coroutines where each interaction point can be resumed arbitrarily often. The method stores a continuation object at the server, referring to it by ID, which can be viewed as a stateful server strategy. In the implementation of our stateful calculi, a client and a server play exactly the same role as coroutines, as explained in Section 4.3, though we use session numbers sent from the client on requests to resume the suspension. In Krishnamurthi *et al.* (2007) and McCarthy (2009, 2010), the authors have addressed the practical problem of how scalable, continuation-based web programs can achieve modular compilation and higher-order interaction with third-party libraries by stack inspection and manipulation techniques known as continuation marks and delimited continuations. Although the techniques used are different, their solution can be viewed as a mixed strategy from the standpoint of our theory.

Using continuations as a language-level technique to structure the client–server interactions of web programs is well known, as has been discussed by Queinnec (2004). Although it is not about multi-tier web programming, it is worth noting that this work discusses the advantages and disadvantages of storing continuations on the client side and server side.

Wysteria (Rastogi *et al.*, 2014) is another related work. This system is in the different domain of secure multi-party computation but shares some similar technical machinery and development. It indexes the types of data and computations by participating principals, which can be regarded as locations. Further, Wysteria also proposes abstractions for private states among participating principals and among subsets of these principals, acting as a secure computation. Similar to our research, it also proposes a translation whose targets consist of ‘sliced’ code for each participant.

6 Conclusion

We have proposed a theory of RPC calculi that deals with the implementation of arbitrarily deep nested client–server interactions using the feature of symmetric communication on the web-based flat request–response interactions in trampolined style under the state-encoding server strategy and the stateful server strategy. In this theory, a new typed RPC calculus λ_{rpc} is proposed where typing with locative types identifies remote function applications at the type level. Our new theory of RPC calculi has advantages over the existing research work. With the help of type-level identification of remote functions, the structure of compilation of λ_{rpc} into λ_{rpc}^{enc} and λ_{rpc}^{state} is simplified compared to the structure of the original compilation without using type information. This simplicity allows our theory to explore three design choices: a state-encoding server strategy, a stateful server strategy and, furthermore, a mixed strategy of these two strategies.

In future, we hope to design a full-fledged multi-tier web programming language based on our theory of RPC calculi. The design must consider many more features such as HTML constructs, reactive UI, databases and concurrency because our theory addresses only client–server communication.

Acknowledgment

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIP) (No. 2017R1A2B4005138).

References

- Balat, V. (2006) Ocsigen: Typing web interaction with Objective Caml. In ML'06: Proceedings of the 2006 Workshop on ML, pp. 84–94.
- Boudol, G., Luo, Z., Rezk, T. & Serrano, M. (2012) Reasoning about web applications: An operational semantics for HOP. *ACM Trans. Program. Lang. Syst.* **34**(2), 10:1–10:40.
- Chlipala, A. (2015) Ur/Web: A simple model for programming the web. In POPL'15: Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, pp. 153–165.
- Conway, M. E. (1963) Design of a separable transition-diagram compiler. *Commun. ACM* **6**(7), 396–408.
- Cooper, E. K. & Wadler, P. (2009) The RPC calculus. In PPDP'09: Proceedings of the 11th ACM SIGPLAN Conference on Principles and Practice of Declarative Programming. New York, NY, USA: ACM, pp. 231–242.
- Cooper, E. K., Lindley, S., Wadler, P. & Yallop, J. (2007) Links: Web programming without tiers. In FMCO'06: Proceedings of the 5th International Conference on Formal Methods for Components and Objects. Berlin, Heidelberg: Springer-Verlag, pp. 266–296.
- Flanagan, C., Sabry, A., Duba, B. F. & Felleisen, M. (1993) The essence of compiling with continuations. In PLDI'93: Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation. New York, NY, USA: ACM, pp. 237–247.
- Ganz, S. E., Friedman, D. P. & Wand, M. (1999) Trampolined style. In ICFP'99: Proceedings of the Fourth ACM SIGPLAN International Conference on Functional Programming. New York, NY, USA: ACM, pp. 18–27.
- Gay, S. J. & Hole, M. (1999) Types and subtypes for client–server interactions. In ESOP'99: Proceedings of the 8th European Symposium on Programming Languages and Systems. London, UK, UK: Springer-Verlag, pp. 74–90.
- Graunke, P., Findler, R. B., Krishnamurthi, S. & Felleisen, M. (2001) Automatically restructuring programs for the web. In ASE'01: Proceedings of the 16th IEEE International Conference on Automated Software Engineering. Washington, DC, USA: IEEE Computer Society, pp. 211–222.
- Krishnamurthi, S., Hopkins, P. W., McCarthy, J., Graunke, P., Pettyjohn, G. & Felleisen, M. (2007) Implementation and use of the PLT scheme web server. *Higher-Order Symb. Comput.* **20**(4), 431–460.
- Launchbury, J. & Peyton Jones, S. L. (1994) Lazy functional state threads. In PLDI'94: Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation. Orlando, Florida, USA, pp. 24–35.
- Matthews, J., Findler, R. B., Graunke, P., Krishnamurthi, S. & Felleisen, M. (2004). Automatically restructuring programs for the web. *Autom. Softw. Eng.* **11**(4), 337–364.
- McCarthy, J. A. (2009) Automatically restful web applications: Marking modular serializable continuations. In ICFP'09: Proceedings of the 14th ACM SIGPLAN International Conference on Functional Programming. New York, NY, USA: ACM, pp. 299–310.
- McCarthy, J. A. (2010) The two-state solution: Native and serializable continuations accord. In OOPSLA'10: Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications. New York, NY, USA: ACM, pp. 567–582.
- Murphy VII, T. (2008) *Modal Types for Mobile Code*. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, USA.

- Murphy VII, T., Cray, K., Harper, R. & Pfenning, F. (2004) A symmetric modal lambda calculus for distributed computing. In LICS'04: Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science. Washington, DC, USA: IEEE Computer Society, pp. 286–295.
- Neubauer, M. & Thiemann, P. (2005) From sequential programs to multi-tier applications by program transformation. In POPL'05: Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. New York, NY, USA: ACM, pp. 221–232.
- Queinnee, C. (2004) Continuations and web servers. *Higher-Order Symb. Comput.* **17**(4), 277–295.
- Radanne, G. (2017) *Tierless Web Programming in ML*. Ph.D. thesis, University Paris Diderot.
- Rastogi, A., Hammer, M. A. & Hicks, M. (2014) Wysteria: A programming language for generic, mixed-mode multiparty computations. In SP'14: Proceedings of the 2014 IEEE Symposium on Security and Privacy. Washington, DC, USA: IEEE Computer Society, pp. 655–670.
- Reynders, B., Devriese, D. & Piessens, F. (2014) Multi-tier functional reactive programming for the web. In Onward!'14: Proceedings of the 2014 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming & Software, pp. 55–68.
- Reynolds, J. C. (1972) Definitional interpreters for higher-order programming languages. In ACM'72: Proceedings of the ACM Annual Conference - Volume 2. New York, NY, USA: ACM, pp. 717–740.
- Serrano, M. & Berry, G. (2012) Multi-tier programming in Hop. *Commun. ACM* **55**(8), 53–59.
- Serrano, M. & Prunet, V. (2016) A glimpse of Hopjs. In ICFP'16: Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, pp. 180–192.
- Serrano, M. & Queinnee, C. (2010) A multi-tier semantics for Hop. *Higher-order Symb. Comput.* **23**(4), 409–431.
- Serrano, M., Gallezio, E. & Loitsch, F. (2006) Hop: A language for programming the web 2.0. In DLS'06: Proceedings of the 1st Dynamic Languages Symposium, pp. 975–985.
- Timany, A., Stefanescu, L., Krogh-Jespersen, M. & Birkedal, L. (2017) A logical relation for monadic encapsulation of state: Proving contextual equivalences in the presence of runST. In PACMPL'17: Proceedings of ACM Programming Languages, 2(POPL), pp. 64:1–64:28.

Appendix: A Proof of Theorem for the RPC Calculus

Theorem A.1 (Type Soundness for λ_{rpc}). *If $\Gamma \triangleright_a M : \tau$ and $M \Downarrow_a V$ then $\Gamma \triangleright_a V : \tau$.*

Proof. We prove this theorem by induction on the height of the typing derivation on M . For (T-Var) as a base case, it is provable trivially. For inductive cases, we do case analysis on the last typing rule used. For (T-Lam), it is straightforward. For (T-App), (T-Req) and (T-Call), we use inductive arguments and the substitution [Lemma A.1](#). In order to prove the cases with (T-Req) and (T-Call), we need an additional [Fact A.1](#) saying that the types of values are the same wherever the values are. \square

Fact A.1. *If $\Gamma \triangleright_a V : \tau$, then $\Gamma \triangleright_b V : \tau$.*

Proof. This is provable due to the definition of (T-Var) and (T-Lam) which are typing rules for V , either a variable and a function. \square

Lemma A.1 (Substitution). *If $\Gamma, x : \tau_1 \triangleright_a M : \tau_2$ and $\Gamma \triangleright_b W : \tau_1$, then $\Gamma \triangleright_a M\{W/x\} : \tau_2$.*

Proof. The proof is straightforward by induction on the height of the typing derivation on M . \square

Appendix B: Proofs of Theorems for the State-Encoding Calculi

The main theorem states the correctness of compilation for λ_{rpc}^{enc} as follows.

Theorem B.1 (Correctness of Compilation for λ_{rpc}^{enc}). *Assume a well-typed λ_{rpc} term M under the locative type system.*

- If $M \Downarrow_c V$ then $C\llbracket M \rrbracket \mid \varepsilon \Rightarrow^{enc*} C\llbracket V \rrbracket \mid \varepsilon$.
- If $M \Downarrow_s V$ then $\Pi \mid S\llbracket M \rrbracket K \Rightarrow^{enc*} \Pi \mid S\llbracket V \rrbracket K$ for all Π and K .

Proof. The proof consists of two parts. The proof firstly builds a sequence of evaluation of the compiled term in λ_{rpc}^{enc} matched to each subtree of evaluation of a given term in λ_{rpc} by induction on the height of the evaluation of M . Then the proof puts the sequences together to construct a longer one matched to the whole evaluation. The composition lemma, Lemma B.1, supports this by showing that, given an λ_{rpc}^{enc} evaluation sequence, it is possible to construct the same sequence but under arbitrary let bindings. The proof uses this lemma as glue in the client. But for the same purpose in the server, the lemma is not used because the proof uses continuations as glue.

In addition, the proof must derive the use of substitutions in λ_{rpc}^{enc} from the use of those in λ_{rpc} to complete the construction of the λ_{rpc}^{enc} evaluation sequence. The substitution lemma, Lemma B.2, supports this by showing that a compilation of a substituted term in λ_{rpc} is the same as a substitution in λ_{rpc}^{enc} of compiled terms.

In the RPC calculus, every base case involves only (Value) where the height is 1, and every inductive case must use (Beta) to make the height at least higher than 1. For the base cases with (Value), both of M and V are $\lambda^b x \cdot M_0$. The theorem can be easily verified since $C\llbracket M \rrbracket = C\llbracket V \rrbracket$ is true.

In the inductive cases, M is $L M_{arg}$. The inductive cases are also analysed into four by the similar reason since (Beta) involves two locations a and b : the application is being evaluated at the location a , and L evaluates to a function of the location b .

i. When a is **s**: Firstly, by the induction hypotheses with three subtrees in the evaluation of the λ_{rpc} term M , three subsequences in λ_{rpc}^{enc} can be constructed with universally quantified variables K_1, K_2, K_3 . Secondly, to combine these subsequences in order to make what the theorem demands, the proof instantiates the three variables with a continuation to $S\llbracket L \rrbracket$ (see Figure 7), a continuation to $S\llbracket M_{arg} \rrbracket$ (see Figure 7) and K in the compilation rules (AppSS) or (AppSC) at Figure 7, respectively. Note that the proof for these inductive cases (a is **s**) is thus proved without using Lemma B.1, which is necessary for the following inductive cases (a is **c**). Finally, the proof uses Lemma B.2 to complete building the evaluation sequence in λ_{rpc}^{enc} as the theorem demands.

ii. When a is **c**: This proof firstly uses the induction hypotheses with the three subtrees in the evaluation of the λ_{rpc} term M to build λ_{rpc}^{enc} evaluation subsequences. Secondly, the proof has to combine these subsequences in order to form a longer one that the theorem demands. Each of $C\llbracket L \rrbracket$, $C\llbracket M_{arg} \rrbracket$ and a local or remote application is surrounded by a let term of the form $\text{let } var = \dots \text{ in } term$ in the compilation rules (AppCC) and (AppCS) at Figure 7. So, the proof has to resort to Lemma B.1 to lift an evaluation subsequence for $C\llbracket L \rrbracket$ to one for $\text{let } var = C\llbracket L \rrbracket \text{ in } term$, and to lift another evaluation subsequence for

$C[[M_{arg}]]$ to one for let $var' = C[[M_{arg}]]$ in $term'$. Finally, the proof takes the third evaluation subsequence for the local or remote application to finish by using [Lemma B.2](#). \square

To prove [Theorem B.1](#), we use [Lemmas B.1](#) and [B.2](#). The former lemma states that the evaluation of terms in λ_{rpc}^{enc} is preserved under the let binding.

Lemma B.1 (Composition). *In the state-encoding RPC calculus λ_{rpc}^{enc} ,*

- *If $M \mid \varepsilon \Rightarrow^{enc*} V \mid \varepsilon$
then $let\ x = M\ in\ M_0 \mid \varepsilon \Rightarrow^{enc*} let\ x = V\ in\ M_0 \mid \varepsilon$*

Proof. This is proved by induction on the length of the evaluation in the condition. In the base case, the length is zero, and so M is identical to V . Therefore, $let\ x = M\ in\ M_0$ is identical to $let\ x = V\ in\ M_0$, which proves the base case.

In inductive cases where the length is one or longer, the proof does a case analysis on kinds of the first rule used in the evaluation of $M \mid \varepsilon \Rightarrow^{enc*} V \mid \varepsilon$ and $\Pi \mid M \Rightarrow^{enc*} \Pi \mid V$, to make the evaluation progress by applying a rule for the kind. Then the proof will have an evaluation subsequence shorter than one in the condition, with which we can apply induction to finish the proof.

Hence, there are four cases for consideration.

- i. (AppC) and (ValC): The lemma is provable simply by induction as explained above.
- ii. (Req): Let M be $let\ y = req(V_f, \bar{W})$ in M_y . Then there exist steps such that

$$\begin{aligned} & let\ y = req(V_f, \bar{W})\ in\ M_y \mid \varepsilon \\ \Rightarrow^{enc} & ctx\ y\ M_y \mid V_f(\bar{W}) \\ \Rightarrow^{enc} & Conf_1 \Rightarrow^{enc} \dots \Rightarrow^{enc} Conf_k \Rightarrow^{enc*} V \mid \varepsilon \end{aligned}$$

where either (Reply) or (Call) must be used in some step, $Conf_{k-1} \Rightarrow^{enc} Conf_k$, and neither (Reply) nor (Call) is used before the step. This is because, in order to move from the server configuration $(ctx\ y\ M_y \mid V_f(\bar{W}))$ to a client configuration $(V \mid \varepsilon)$, there is no other way than using (Reply) or (Call) by the definition of the semantic rules for λ_{rpc}^{enc} .

After the use of either (Reply) or (Call), it is easy to construct,

$$Conf_k \Rightarrow^{enc+} Conf_{k+n} \Rightarrow^{enc*} V \mid \varepsilon$$

where all except $Conf_{k+n}$ are for the server and $Conf_{k+n} \Rightarrow^{enc*} V \mid \varepsilon$ is a condition that we can apply induction with. This finishes the proof for this case.

iii. (LetC): In this case, M is $let\ x_1 = (let\ x_2 = M_1\ in\ M_2)$ in M_3 . This is proved by induction after using the intermediate value lemma, [Lemma B.3](#), and the context-independent intermediate value lemma, [Lemma B.4](#). By applying (LetC) two times, we have:

$$let\ x = M\ in\ M_0 \mid \varepsilon \Rightarrow^2 let\ x_2 = M_1\ in\ (let\ x_1 = M_2\ in\ (let\ x = M_3\ in\ M_0)) \mid \varepsilon$$

where induction cannot be applied immediately. After the evaluation of both M_1 and M_2 , we will reach a configuration as: $let\ x = M_3\{V_1/x_1\}$ in $M_0 \mid \varepsilon$ where V_1 is a value of $(let\ x_2 = M_1\ in\ M_2)$. Then we can apply induction with $M_3\{V_1/x_1\} \mid \varepsilon \Rightarrow^* V \mid \varepsilon$. We thus need to make the evaluation progress over M_1 and M_2 . To do this, we apply the two additional lemmas to lead the evaluation to the configuration above. This finishes the proof in the case. \square

Lemma B.2 shows the compilation of a substituted term yields the same term as the compilation of a term substituted with compiled values.

Lemma B.2 (Substitution). *The compilation $C[-]$ and $S[-]$ of λ_{rpc} into λ_{rpc}^{enc} preserve substitution operations in λ_{rpc} and λ_{rpc}^{enc} . Let $V[\lambda^a x \cdot M] = C[\lambda^a x \cdot M]$.*

- $C[N\{W/x\}] = C[N]\{V[W]/x\}$.
- $S[N\{W/x\}] K = (S[N] K)\{V[W]/x\}$.

Proof. This lemma is true because of the definition of the compilation rules. □

To prove **Lemma B.1**, we introduce **Lemmas B.3** and **B.4**, as follows.

Lemma B.3 (Intermediate Value).

- If $\text{let } x = M \text{ in } M_0 \mid \varepsilon \Rightarrow^{enc*} V \mid \varepsilon$
then $\text{let } x = M \text{ in } M_0 \mid \varepsilon \Rightarrow^{enc*} \text{let } x = V_x \text{ in } M_0 \mid \varepsilon \Rightarrow^{enc*} V \mid \varepsilon$.

Proof. This is proved by induction on the length of the evaluation in the conditions. □

Lemma B.4 (Context-Independent Intermediate Value).

- If $\text{let } x = M \text{ in } M_1 \mid \varepsilon \Rightarrow^{enc*} \text{let } x = V \text{ in } M_1 \mid \varepsilon$
then $\text{let } x = M \text{ in } M_2 \mid \varepsilon \Rightarrow^{enc*} \text{let } x = V \text{ in } M_2 \mid \varepsilon$.

Proof. This is proved by induction on the length of the evaluation in the conditions. In (ValC), (AppC) and (Req), the proof is straightforward. In (LetC), when M is $(\text{let } y = M_y \text{ in } M_x)$, we have $(\text{let } x = M \text{ in } M_1) \mid \varepsilon \Rightarrow^{enc} \text{let } y = M_y \text{ in } (\text{let } x = M_x \text{ in } M_1) \mid \varepsilon$. The proof uses an inner induction on the number of the let constructs enclosing $\text{let } x = M_x \text{ in } M_1$ such as ‘ $\text{let } y = M_y \text{ in } -$ ’. The use of the inner induction will eventually lead to a form of $\text{let } x = M_x\{V_y/y\} \text{ in } M_1 \mid \varepsilon$ where V_y is a value of M_y . Then we can apply the outer induction to it to finish the proof. □

The theorem of the correctness of compilation for λ_{cs}^{enc} immediately follows the theorem for λ_{rpc}^{enc} .

Theorem B.2 (Correctness of Compilation for λ_{cs}^{enc}). *Assume a well-typed term M under the locative type system. Given ϕ_c and ϕ_s ,*

- If $M \Downarrow_c V$ then $CC[C[M]] \mid \varepsilon \Rightarrow^{enc*} CC[C[V]] \mid \varepsilon$.

Proof. By **Theorem B.1**, it holds that if $M \Downarrow_c V$, then $C[M] \mid \varepsilon \Rightarrow^{enc*} C[V] \mid \varepsilon$. From such a series of evaluation steps in λ_{rpc}^{enc} , it is easy to construct evaluation steps in λ_{cs}^{enc} as $CC[C[M]] \mid \varepsilon \Rightarrow^{enc*} CC[C[V]] \mid \varepsilon$. This is because of two reasons, as follows.

First, the compilation rules $CC[-]$ preserve the structure of terms, for example, mapping a variable in λ_{rpc}^{enc} into a variable in λ_{cs}^{enc} , an application in λ_{rpc}^{enc} into an application in λ_{cs}^{enc} , and so on. The only exception is to map a function in λ_{rpc}^{enc} into a closure in λ_{cs}^{enc} , but a

closure has a role as a function in λ_{cs}^{enc} . Second, the semantic rules in λ_{rpc}^{enc} are isomorphic to the semantic rules in λ_{cs}^{enc} .

To construct evaluation steps starting from $CC[C[M]] \mid \varepsilon$ in λ_{cs}^{enc} , we have only to choose the isomorphic semantic rule corresponding to the semantic rule used in each of the evaluation steps in λ_{rpc}^{enc} . \square

Appendix C: Proofs of Theorems for the Stateful Calculi

The structure of proofs of theorems for the stateful calculi is very similar to that of those for the state-encoding calculi.

The main theorem states the correctness of compilation for λ_{rpc}^{state} as follows.

Theorem C.1 (Correctness of Compilation for λ_{rpc}^{state}). *Assume a well-typed term $\lambda_{rpc} M$ under the locative type system.*

- If $M \Downarrow_c V$, then $C[M] \mid \Delta \Rightarrow^{state*} C[V] \mid \Delta$ for all Δ , which is call-return balanced.
- If $M \Downarrow_s V$ then $\Pi \mid \Delta; S[M] \Rightarrow^{state*} \Pi \mid \Delta; S[V]$ for all Π and Δ , which is call-return balanced.

Proof. This proof has the same structure as for [Theorem B.1](#). It firstly builds a subsequence of evaluation of the compiled term in λ_{rpc}^{state} to each subtree of evaluation of a given term in λ_{rpc} , and then it puts the subsequences together into a longer one for the whole evaluation. The composition lemma, [Lemma C.1](#), supports the second part. The substitution lemma, [Lemma C.2](#), is required in the second part to show that the λ_{rpc}^{state} substitution after a λ_{rpc} -to- λ_{rpc}^{state} compilation produces the same term as one by a λ_{rpc} -to- λ_{rpc}^{state} compilation after a λ_{rpc} substitution.

Now let us prove by induction on the height of the evaluation of M . The base cases with height 1 use (Value) in the RPC calculus. For base cases with (Value), both M and V are the same as $\lambda^b x \cdot M_0$, and so the theorem is trivially true by zero step evaluation.

The inductive cases with higher than height 1 must use (Beta) in the condition of this theorem. Every case is provable by induction and the use of [Lemmas C.1](#) and [C.2](#) in the same way as for [Theorem B.1](#). For the case with $L^c M_{arg} \Downarrow_c V$, each evaluation subsequence in λ_{rpc}^{state} obtained from each subtree of evaluation in λ_{rpc} is balanced by induction. The three subsequences are combined by having a single configuration using (ValC) between the first and the second, and by having another configuration using (AppC) between the second and the third. (ValC) is used in the end. By the definition of call-return balance, the resulting evaluation sequence in λ_{rpc}^{state} is call-return balanced. The case with $L^s M_{arg} \Downarrow_s V$ is verified exactly in the same way but with (ValS), (AppS) and (ValS).

For the case with $L^s M_{arg} \Downarrow_c V$, the three call-return balanced evaluation subsequences are combined with intermediate configurations using (ValC) between the first and the second, using (Req) and (AppS) between the second and the third, and using (ValS), (Reply) and (ValC) in the end.

The case with $L^c M_{arg} \Downarrow_s V$ is most interesting because a new pair of call and ret will be added to the resulting evaluation sequence explicitly, not just by induction. The first subsequence is followed by configurations using (ValS), which is followed by the second

subsequence, which is followed by other configurations using (Call), (AppC), (LetC) and (AppC), which is followed by the third subsequence, which is lastly followed by a configuration using (Ret). The whole evaluation sequence is call-return balanced by the definition of $bal ::= Conf_{call} \text{ bal } Conf_{ret}$. \square

To prove [Theorem C.1](#), we use a lemma stating that the evaluation of terms in λ_{rpc}^{enc} is preserved under the let binding, as follows.

Lemma C.1 (Composition). *In the stateful RPC calculus λ_{rpc}^{state} ,*

- *If $M \mid \Delta \Rightarrow^{state*} V \mid \Delta$ and it is call-return balanced, then $\text{let } x = M \text{ in } M_0 \mid \Delta \Rightarrow^{state*} \text{let } x = V \text{ in } M_0 \mid \Delta$, which is call-return balanced.*
- *If $\Pi \mid \Delta; M \Rightarrow^{state*} \Pi \mid \Delta; V$ and it is call-return balanced, then $\Pi \mid \Delta; \text{let } x = M \text{ in } M_0 \Rightarrow^{state*} \Pi \mid \Delta; \text{let } x = V \text{ in } M_0$, which is call-return balanced.*

Proof. This is proved by induction on the length of the evaluation in the conditions. In the base case, the length is zero, and so M is identical to V , which proves the lemma.

In inductive cases where the evaluation takes one or more steps, the proof is done by case analysis on kinds of the first rule used in the evaluation of $M \mid \Delta \Rightarrow^{state*} V \mid \Delta$ and $\Pi \mid \Delta; M \Rightarrow^{state*} \Pi \mid \Delta; V$. Hence, there are 10 cases for consideration.

i. (AppC), (AppS), (ValC), (ValS): The lemma is proved simply by induction. The resulting evaluation sequence preserves the call-return balance because it is just the inductively obtained balanced evaluation subsequences appended in the front or in the end with $Conf_{else}$ using some of (AppC), (LetC), (ValC), (AppS), (LetS) and (ValS).

ii. (LetC), (LetS): In these cases, M is $\text{let } x_1 = (\text{let } x_2 = M_1 \text{ in } M_2) \text{ in } M_3$. These cases are proved by induction together with [Lemma C.3](#) and [Lemma C.4](#). The two additional lemmas are used to uncover values which M_1 and M_2 evaluate to and then to make the evaluation of $\text{let } x = M \text{ in } M_0$ progress until it reaches a configuration $\text{let } x = M_3 \{V_{1,2}\}$ in M_0 , to which induction can be applied. The condition call-return balance is verified by $bal ::= balbal$ because all the evaluation subsequences obtained from applying induction and the two lemmas are all call-return balanced and only some extra configurations, which are not $Conf_{call}$ nor $Conf_{ret}$, are introduced.

iii. (Ret): The call-return balance is not satisfied in the condition of this case. So, it is vacuously true.

iv. (Reply): By the definition of (Reply), M is a value. When $M = V$, it is trivially true with the zero evaluation step. When $M \neq V$, the evaluation sequence in the condition of this case is impossible. So, it is vacuously true.

v. (Call): Let M be $\text{let } y = \text{call}(V_f, \overline{W})$ in M_y and let Π be $\text{ctx } z \ M_z$.

$$\begin{aligned}
 & \Pi \mid \Delta; \text{let } y = \text{call}(V_f, \overline{W}) \text{ in } M_y && \text{by (Call)} \\
 \Rightarrow^{enc} & \text{let } z = V_f(\overline{W}) \text{ in } M_z \mid \text{ctx } y \ M_y \cdot \Delta && \text{by the balance} \\
 \Rightarrow^{enc*} & \text{let } z = \text{ret}(V_{call}) \text{ in } M_z \mid \text{ctx } y \ M_y \cdot \Delta && \text{by (Ret)} \\
 \Rightarrow^{enc} & \Pi \mid \Delta; \text{let } y = V_{call} \text{ in } M_y && \text{by (Val)} \\
 \Rightarrow^{enc} & \Pi \mid \Delta; M_y \{v_{call}/y\} && \text{by the condition} \\
 \Rightarrow^{enc*} & \Pi \mid \Delta; V &&
 \end{aligned}$$

It is possible to mock the evaluation steps except the last steps above with $\Pi \mid \Delta$; let $x = M$ in M_0 because all the steps between $Conf_{call}$ and $Conf_{ret}$ above use only the evaluation rules that do not depend on the server stack. For the last steps of evaluation above, we prove it by induction. The call-return balance is verified by $bal ::= bal\ bal$.

vi. (Req): Let M be let $y = req(V_f, \overline{W})$ in M_y . We do an analysis to look for a configuration as follows.

$$\begin{aligned} & \text{let } y = req(V_f, \overline{W}) \text{ in } M_y \mid \Delta \\ \Rightarrow^{enc} & \text{ctx } y M_y \mid \Delta; \text{ let } r = V_f(\overline{W}) \text{ in } r \\ \Rightarrow^{enc} & Conf_1 \Rightarrow^{enc} \dots \Rightarrow^{enc} Conf_k \Rightarrow^{enc*} V \mid \Delta \end{aligned}$$

where either (Reply) or (Call) is used in a step, $Conf_{k-1} \Rightarrow^{enc} Conf_k$, and neither (Reply) nor (Call) is used before the step. There should exist such a configuration. This is because, in order to move from a server configuration ($\text{ctx } y M_y \mid \Delta; \text{ let } r = V_f(\overline{W}) \text{ in } r$) to a client configuration ($V \mid \Delta$), there is no other way than using (Reply) or (Call) by the definition of the semantic rules for λ_{rpc}^{enc} .

For use of (Reply), it is easy to construct, after the step,

$$Conf_k \Rightarrow^{enc+} Conf_{k+n} \Rightarrow^{enc*} V \mid \Delta$$

where $Conf_{k+n} \Rightarrow^{enc*} V \mid \Delta$ is a condition to apply induction with. Since the subsequence from $Conf_1$ to $Conf_{k-1}$ does not contain any configurations of the form $Conf_{call}$ or $Conf_{ret}$, the condition of the lemma implies that the subsequence from $Conf_k$ to $V \mid \Delta$ is balanced.

For use of (Call), an inner induction is needed to make it progress to reach a configuration that induction of the lemma is applicable. Let $Conf_k$ be

$$\begin{aligned} & \text{ctx } y M_y \mid \Delta; \text{ let } z = \text{call}(V_g, \overline{W}_{garg}) \text{ in } M_z \\ & Conf_k && \text{by (Call)} \\ \Rightarrow^{enc} & \text{let } y = V_g(\overline{W}_{garg}) \text{ in } M_y \mid \text{ctx } z M_z \cdot \Delta && \text{by the balance} \\ \Rightarrow^{enc*} & \text{let } y = \text{ret}(V_{call}) \text{ in } M_y \mid \text{ctx } z M_z \cdot \Delta && \text{by (Ret)} \\ \Rightarrow^{enc} & \text{ctx } y M_y \mid \Delta; \text{ let } z = V_{call} \text{ in } M_z && \text{by (Val)} \\ \Rightarrow^{enc} & \text{ctx } y M_y \mid \Delta; M_z\{v_{call}/z\} \end{aligned}$$

The subsequence after the configuration ($\text{ctx } y M_y \mid \Delta; \text{ let } z = V_{call} \text{ in } M_z$) all the way to ($V \mid \Delta$) is balanced by the following argument. The subsequence before the configuration is constructed to be balanced. Since the whole of the evaluation sequence in the condition of the lemma is balanced, the definition of call-return balance implies that the subsequence after the configuration is balanced.

We repeat the analysis again from the last configuration above until we find any use of (Reply), which allows to apply induction of the lemma. Since the steps are finite by the condition, the analysis will be terminated in finite time. This finishes the proof. \square

Lemma C.2 (Substitution). *Let $V[\lambda^a x \cdot M] = C[\lambda^a x \cdot M]$.*

- $C[N\{W/x\}] = C[N]\{V[W]/x\}$.
- $S[N\{W/x\}] K = (S[N] K)\{V[W]/x\}$.

Proof. This lemma is true because of the definition of the compilation rules. \square

To prove [Lemma C.1](#), we introduce two lemmas as follows.

Lemma C.3.

- If $\text{let } x = M \text{ in } M_0 \mid \Delta \Rightarrow^{state*} V \mid \Delta$ and it is call-return balanced, then $\text{let } x = M \text{ in } M_0 \mid \Delta \Rightarrow^{state*} \text{let } x = V_x \text{ in } M_0 \mid \Delta \Rightarrow^{state*} V \mid \Delta$.
- If $\Pi \mid \text{let } x = M \text{ in } M_0 \Rightarrow^{state*} \Pi \mid V$ and it is call-return balanced, then $\Pi \mid \text{let } x = M \text{ in } M_0 \Rightarrow^{state*} \Pi \mid \text{let } x = V_x \text{ in } M_0 \Rightarrow^{state*} \Pi \mid V$.

In each conclusion of the two conditional statements above, the former evaluation sequence until M evaluates to V_x and the latter one after that are both call-return balanced.

Proof. This is proved by induction on the length of the evaluation in the conditions. \square

Lemma C.4.

- If $\text{let } x = M \text{ in } M_1 \mid \Delta \Rightarrow^{state*} V_1 \mid \Delta$ and it is call-return balanced, then $\text{let } x = M \text{ in } M_2 \mid \Delta \Rightarrow^{state*} \text{let } x = V \text{ in } M_2 \mid \Delta$ and it is also call-return balanced.
- If $\Pi \mid \Delta; \text{let } x = M \text{ in } M_1 \Rightarrow^{state*} \Pi \mid \Delta; V_1$ and it is call-return balanced, then $\Pi \mid \Delta; \text{let } x = M \text{ in } M_2 \Rightarrow^{state*} \Pi \mid \Delta; \text{let } x = V \text{ in } M_2$ and it is also call-return balanced.

Proof. The lemma is proved by induction on the length of the evaluation in the conditions. \square

The theorem of the correctness of compilation for λ_{cs}^{state} immediately follows the theorem for λ_{rpc}^{state} .

Theorem C.2 (Correctness of Compilation for λ_{cs}^{state}). *Assume a well-typed term M under the locative type system. Given ϕ_c and ϕ_s ,*

- If $M \Downarrow_c V$, then $CC[C[M]] \mid \varepsilon \Rightarrow^{state*} CC[C[V]] \mid \varepsilon$.

Proof. This theorem is proved by the same argument used in the proof of [Theorem B.2](#). The key idea is that the compilation rules $CC[-]$ preserve the structure of λ_{rpc}^{state} terms. \square