

Number 578



**UNIVERSITY OF
CAMBRIDGE**

Computer Laboratory

Linear types for packet processing (extended version)

Robert Ennals, Richard Sharp, Alan Mycroft

January 2004

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500
<http://www.cl.cam.ac.uk/>

© 2004 Robert Ennals, Richard Sharp, Alan Mycroft

Technical reports published by the University of Cambridge
Computer Laboratory are freely available via the Internet:

<http://www.cl.cam.ac.uk/TechReports/>

Series editor: Markus Kuhn

ISSN 1476-2986

Linear Types for Packet Processing

(extended version)

Robert Ennals

Richard Sharp

Alan Mycroft

Abstract

We present PACLANG: an imperative, concurrent, linearly-typed language designed for expressing packet processing applications. PACLANG’s linear type system ensures that no packet is referenced by more than one thread, but allows multiple references to a packet *within a thread*. We argue (i) that this property greatly simplifies compilation of high-level programs to the distributed memory architectures of modern Network Processors; and (ii) that PACLANG’s type system captures that style in which imperative packet processing programs are already written. Claim (ii) is justified by means of a case-study: we describe a PACLANG implementation of the IPv4 unicast packet forwarding algorithm.

PACLANG is formalised by means of an operational semantics and a Unique Ownership theorem formalises its correctness with respect to the type system.

1 Introduction

Network Processors (NPs) [1, 10, 19] are programmable, application-specific hardware architectures designed to facilitate high-speed packet processing. NPs typically contain multiple processor cores (allowing multiple network packets to be processed concurrently) and multiple memory banks. The majority of NPs are based on what is known as a *distributed memory architecture*: not all memory banks are accessible from all processor cores.

State-of-the-art programming techniques for NPs are invariably low-level, requiring programmers to specify explicitly how data will be mapped to, and dynamically moved between, complex arrays of available memory banks. This results in design flows which are difficult to master, error prone and inextricably tied to one particular architecture. Our research attempts to address this problem by developing higher-level languages and tools for programming NPs.

This paper formally defines an imperative language which provides an architecturally neutral concurrent programming model—multiple parallel threads accessing a shared heap—and considers the problem of compiling this language to NP-based architectures. We argue that although, *in general*, the efficient compilation of a high-level concurrent language to a diverse range of NP architectures is very difficult, *in the domain of packet processing* the problem is more tractable. Our argument is based on the observation that a great deal of network processing code is written in a restricted data-flow style, in which network packets can be seen to “flow” through a packet processing system. We formalise this restricted style of programming by means of a linear type system and prove a *unique*

ownership property: if a program is typeable then at any given point in its execution, each packet in the heap is referenced by exactly one thread. It is this property that we argue simplifies the process of compiling for network processors.

Contributions of this paper:

The contributions of this paper are:

- The formalisation of a concurrent, first-order, imperative language which we argue is well suited for compiling to a diverse range of NP architectures.
- A simple, intuitive linear type system in which no packet is referenced by more than one thread but multiple references to packets *within a thread* are allowed.
- A case-study that demonstrates both that the type system is applicable to the domain of packet processing (i.e. not overly restrictive) and that the resulting code is comprehensible to C/C++ programmers.

Previous research [17] has shown that (i) locking is unnecessary for linearly typed objects; and (ii) memory leaks can be detected statically as type errors. Although these are not direct contributions of this paper, they are nevertheless benefits enjoyed by PACLANG programmers.

Structure of this paper:

We outline the design of an NP-based platform and explain how our unique ownership property helps us target such architectures (Section 1.1). Section 2 gives formal syntax and types for PACLANG and Section 3 defines its linear typing judgements along with more intuitive explanation. An operational semantics for the language is presented and the *unique ownership property* formalised (Section 4) and proved (Appendix A). To justify that our linear type system is not overly restrictive in the domain of packet processing applications, we present a case study which shows how PACLANG can be used to implement an IPv4 packet forwarder (Section 5). Finally we consider related work (Section 6) and conclude, giving directions for future research/development (Section 7).

1.1 Compiling for Network Processors

Figure 1 gives a concrete example of an NP-based system: the Intel IXDP2400 [9]. The platform consists of two Intel IXP2400 NPs [10] connected in a pipeline configuration. Each NP contains 16 small RISC processors called *micro-engines* (labelled $\mu E-1, \dots, \mu E-16$), a single Intel X-Scale processor and some on-chip shared RAM. Each micro-engine¹ has its own local memory. The external SRAM and DRAMs can be accessed by all the cores on the NP to which they are connected. Micro-engines are equipped with *next neighbour registers*. The next neighbour registers of $\mu E-i$ are accessible to $\mu E-(i+1)$, forming a uni-directional datapath between consecutive micro-engines.

¹Note that each micro-engine actually supports 8 hardware-level threads. This level of detail is beyond the scope of the paper.

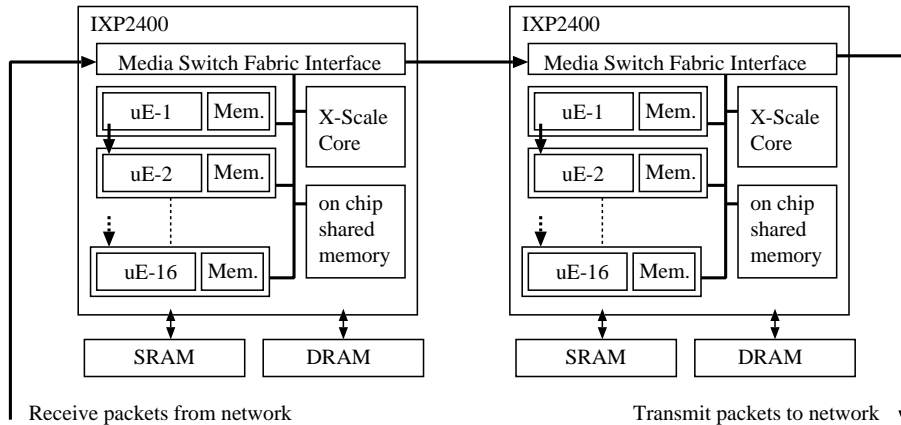


Figure 1: Simplified diagram of IXDP2400 NP evaluation board

When programming for architectures such as this a great deal of low-level knowledge is required to decide how packets should be represented in memory and how packets should be transferred between different memory banks². For example, if a small amount of data needs to be transferred between two threads running on $\mu\text{E-1}$ and $\mu\text{E-2}$, then the most efficient mechanism may well be to move the data across the serial datapath connecting $\mu\text{E-1}$ and $\mu\text{E-2}$. Conversely, if we have data already in DRAM that needs to be transferred between threads running on $\mu\text{E-1}$ and the X-Scale core then leaving the data where it is and passing a reference to it is the obvious choice. When passing data between the two separate IXP2400s we *must* pass-by-value since the two processors have no memories in common.

Clearly, if we are to succeed in our goal of designing a high-level language that can be compiled to a variety of different NP architectures then a compiler must be *free to choose* how data will be passed between threads (e.g. pass a reference pointing to shared memory, or pass the data itself over a shared bus). Unfortunately, conventional imperative languages (e.g. C/C++) are unsatisfactory in this respect since they force the programmer to *explicitly specify* whether data will be passed by value or passed by reference. Note that a compiler cannot simply transform pass-by-value into pass-by-reference (or vice versa); sophisticated *alias analysis* is required to determine whether this transformation is semantics-preserving—an analysis that is undecidable in general.

In contrast, in the PACLANG framework, the unique ownership property (provided by our linear type system) gives us two key benefits: (i) a thread can dynamically move an object it owns to a new location without worrying about dangling references; and (ii) pass-by-value and pass-by-reference are semantically indistinguishable when transferring data between threads. This gives a PACLANG compiler the freedom to make decisions regarding both the flow of packets between available memory banks and the low-level mechanisms employed to implement these transfers, on a per-architecture basis.

²Moving a packet between different memories as it “flows” through a packet processing application is crucial to achieving efficient implementation on the IXP (and other NPs). Data should be located “as close as possible” to the thread operating on it in order to reduce memory latency.

v	\leftarrow	i	<i>integer constant</i>	
		$ $	b	<i>boolean constant</i>
		$ $	x	<i>variable</i>
e	\leftarrow	if v then e_1 else e_2	<i>conditional</i>	
		$ $ let $(x_1, \dots, x_j) = f(v_1, \dots, v_k)$ in e	$(j, k \geq 0)$ <i>call</i>	
		$ $ return (v_1, \dots, v_k)	$(k \geq 0)$ <i>return</i>	
d	\leftarrow	$(\sigma_1, \dots, \sigma_j) f(\rho_1 x_1, \dots, \rho_k x_k)\{e\}$	$(j, k \geq 0)$ <i>function definition</i>	
		$ $ <i>packetQueue</i> q ;	<i>queue definition</i>	
p	\leftarrow	$d_1 \dots d_n$	<i>top-level program</i>	
ρ, σ, τ	\leftarrow	<i>int</i> $ $ <i>bool</i> $ $ <i>packet</i> $ $ <i>!packet</i>	<i>types</i>	

Figure 2: Syntax of PACLANG, our simple first-order CBV language.

2 A Packet Processing Language

PACLANG is a concurrent, first-order, imperative language with a linear type system; its syntax is presented in Figure 2. Concurrency is supported by the simple expedient of treating functions whose name is spelt `main1, ..., mainn`, as representing the entry points of n static threads ($n \geq 1$)³. Such functions take no arguments and return no results. Statically allocated, synchronous (blocking) packet queues are provided to facilitate inter-thread communication.

Expressions, e , comprise conditionals, function calls and **returns**. Note that PACLANG syntax forces all function return values to be explicitly bound to variables (cf. A-Normal Form [6]); this simplifies subsequent typing and semantic rules. In addition to function names defined by *function definitions* we assume some functions to be built-in, e.g. arithmetic and queue manipulation primitives (see later); both forms may be called using the **let** construct. We use v to range over simple expressions (constants and variables). A program, p , consists of a sequence of function and queue definitions, d . Note that functions take multiple arguments and return multiple results.

In contrast to most languages in the linear types literature, PACLANG has imperative features. Given a packet reference p , an integer offset i and a value v , the built-in function `update(p, i, v)` updates p to contain v at offset i . The packet read function `read(p, i)` returns the value stored at offset i in packet p .

Global queues also provide side-effecting enqueue and dequeue operations. To simplify the presentation we assume built-in `enq` and `deq` functions replicated over queue names. We write `q.enq(p)` to enqueue a packet p in the queue declared with name `q`. Similarly, `q.deq()` returns a packet dequeued from `q`. The semantics of queues are formalised in Section 4.

We also provide built-in functions `new(i)` which given an integer size i , creates a new

³In practice a compiler may wish to transform a single programmer-specified thread into several finer-grained threads for efficient implementation on certain NP architectures. Such *automated partitioning* is beyond the scope of this paper.

packet buffer of size i , returning its associated *packet reference*; and `kill(p)`, which takes a packet reference p , consumes it, deallocates the associated packet and returns nothing. In PACLANG packets are simply dynamically created arrays. As well as representing network packets, they can also be used as containers for arbitrary (non-pointer) data within a PACLANG program.

Syntactic sugaring is employed to increase readability for those familiar with C. We write “ $f(v_1, \dots, v_k); e$ ” to abbreviate “`let () = $f(v_1, \dots, v_k)$ in e` ”. We allow nested function calls to be written directly, assuming a left-to-right evaluation order. For example, the expression “`let $x = f(g(x), h(y))$ in e` ”, becomes “`let $z_1 = g(x)$ in let $z_2 = h(y)$ in let $x = f(z_1, z_2)$ in e` ” where $z_1, z_2 \notin fv(e)$. We add types to the variables following `let` (corresponding to the return type of the called function) as usual in C, and also drop the `let` keyword and parentheses entirely. We write `void` in place of `()` in the signature of functions that return an empty tuple of values and adopt the convention that an expression without an explicit `return` implicitly returns a value of type `void`. We also introduce a `switch/case` construct which is translated into nested ifs in the usual way. We write `update(p, a, v)` using C-like syntax, $p[a] = v$, and use array subscript notation, $p[a]$, for `read(p, a)`. To avoid the extra syntax required for explicit while-loops, we use recursion to model iteration in this presentation. (Of course, a PACLANG compiler would implement tail-call optimisation to ensure the generated code is as efficient as if it had been written with while-loops.)

3 Linear Type System

An industrial NP programming textbook [11] observes: “*In most designs processing of a packet by multiple threads simultaneously is rare, but using multiple threads to process packets at different points in the packet’s lifetime is common.*” This quote summarises the philosophy behind PACLANG’s linear type system: we (i) *guarantee* that, in a well typed program, a packet, p , cannot be accessed by multiple threads simultaneously; whilst (ii) allowing p to be transferred between threads as it flows through an application. Unlike many linear type systems, we allow p to be referenced multiple times (aliasing). However, we require that all references to p reside in the same thread. Allowing intra-thread aliasing is crucial to the readability of PACLANG code. It is this property that allows PACLANG programmers to write in the familiar imperative (pass-by-reference) style—see Section 5. In many ways, PACLANG’s type system is inspired by Kobayashi’s work on *Quasi-Linear Types* [12].

We start by defining some terminology. A variable, x , that holds a reference to a packet p , will either be of type *packet* or *!packet*:

- If x has type *packet* then, at any point where x is in scope, there are no other references to p . We say that x is an *owning variable* and holds an *owning reference* for p . If one has an owning reference to p then one can safely move p to another thread without creating dangling references.
- If x has type *!packet* then other references to p may exist, but they must be on the same thread as x . We say that x is an *alias* for p . An alias cannot escape from its enclosing scope.

Value types, *int* and *bool*, are non-linear and behave as is normal.

Our type rules use arrow-notation for vectors. We write $\vec{\tau}$ to represent a vector of types and similarly \vec{x} and \vec{v} for vectors of variables and values respectively. We let ρ, σ, τ range over types. Function types are written $\rho_1 \times \dots \times \rho_k \rightarrow \sigma_1 \times \dots \times \sigma_j$ or more briefly as $\vec{\rho} \rightarrow \vec{\sigma}$. Given a function name f , $\mathcal{F}(f)$ gives its type signature. For user-defined functions this comes from the textual definition; the types of built-in functions are given in Figure 4.

Typing judgements for expressions (\vdash^e) take the form $\Gamma \vdash^e e : \vec{\tau}$ where Γ is a typing environment, e is an expression (generally returning a tuple) and $\vec{\tau}$ is the return type (generally a tuple type). Γ can be regarded as the set of *capabilities* that e requires during its execution. Each capability in Γ takes the form $x:\rho$ specifying that x is used as ρ in e . We adopt the usual notation of writing $\Gamma, x:\tau$ to represent Γ extended to include the capability $x:\tau$ (note that scope shadowing is problematic in linear type systems so we will simply presume that all variable names are distinct).

The typing rules for PACLANG, rely on two partial binary operators, “+” and “;” which operate on type environments. An expression which requires capabilities in Γ_1 and then subsequently the capabilities in Γ_2 , can equivalently be said to require the capabilities in environment $\Gamma_1; \Gamma_2$. Similarly, if an expression requires capabilities in both Γ_1 and Γ_2 in an unordered fashion, it requires $\Gamma_1 + \Gamma_2$. We first define, “+” and “;” over *types*:

$$\begin{array}{ll} \textit{int} + \textit{int} & = \textit{int} & \textit{int} ; \textit{int} & = \textit{int} \\ \textit{bool} + \textit{bool} & = \textit{bool} & \textit{bool} ; \textit{bool} & = \textit{bool} \\ \textit{!packet} + \textit{!packet} & = \textit{!packet} & \textit{!packet} ; \textit{!packet} & = \textit{!packet} \\ & & \textit{!packet} ; \textit{packet} & = \textit{packet} \end{array}$$

If no rule is given then that signifies a type error. Note that “;” is a strict superset of “+”. We extend these operators to type environments by the following definitions (here \otimes ranges over “+” and “;”):

$$\begin{aligned} \Gamma_1 \otimes \Gamma_2 & \stackrel{\textit{def}}{=} \{x : \tau \mid x : \tau \in \Gamma_1 \wedge x \notin \text{dom } \Gamma_2\} \cup \\ & \{x : \tau \mid x : \tau \in \Gamma_2 \wedge x \notin \text{dom } \Gamma_1\} \cup \\ & \{x : \tau_1 \otimes \tau_2 \mid x : \tau_1 \in \Gamma_1 \wedge x : \tau_2 \in \Gamma_2\} \end{aligned}$$

Figure 3 presents the linear type rules. Typing rules for definitions (\vdash^d), programs (\vdash^p) and simple expressions (\vdash^v) are self-explanatory, although note that, in contrast to non-linear type systems, the (*var*) rule only holds if the environment contains the *single* assumption, $x:\tau$.

The typing rules for expressions (\vdash^e) deserve more detailed explanation. Essentially they ensure (*i*) that an owning variable x for a packet p goes out of scope for the lifetime of any aliases to p ; and (*ii*) that an owning variable x may no longer be used if its owning reference is passed to another owning variable.

For a variable x to be of type *!packet*, x must be bound as a formal parameter of its enclosing function definition. When calling a function, f , the (*apseq*) rule allows aliases to be created from an owning variable provided that the owning variable goes out of scope⁴

⁴In PACLANG this is ensured by the absence of nested function definitions and the absence of global packet variables.

$$\begin{array}{c}
\frac{}{\{x : \tau\} \vdash^v x : \tau} \textit{(var)} \qquad \frac{}{\emptyset \vdash^v i : \textit{int}} \textit{(int)} \quad (i \in \mathbb{Z}) \qquad \frac{}{\emptyset \vdash^v b : \textit{bool}} \textit{(bool)} \quad (i \in \mathbb{B}) \\
\\
\frac{\Gamma_1 \vdash^v v_1 : \rho_1 \quad \dots \quad \Gamma_k \vdash^v v_k : \rho_k \quad \Gamma_0, x_1 : \sigma_1, \dots, x_j : \sigma_j \vdash^e e : \vec{\tau}}{(\Gamma_1 + \dots + \Gamma_k); \Gamma_0 \vdash^e \textit{let } (x_1, \dots, x_j) = f(v_1, \dots, v_k) \textit{ in } e : \vec{\tau}} \textit{(apseq)} \\
\text{provided } \mathcal{F}(f) = (\rho_1 \times \dots \times \rho_k) \rightarrow (\sigma_1 \times \dots \times \sigma_j) \\
\\
\frac{\Gamma_1 \vdash^v v_1 : \tau_1 \quad \dots \quad \Gamma_j \vdash^v v_j : \tau_j}{\Gamma_1 + \dots + \Gamma_j \vdash^e \textit{return } (v_1, \dots, v_j) : \tau_1 \times \dots \times \tau_j} \textit{(ret)} \\
\text{provided no } \tau_j \textit{ is } !\textit{packet} \\
\\
\frac{\Gamma_1 \vdash^v v : \textit{bool} \quad \Gamma_2 \vdash^e e_1 : \vec{\tau} \quad \Gamma_2 \vdash^e e_2 : \vec{\tau}}{\Gamma_1; \Gamma_2 \vdash^e \textit{if } v \textit{ then } e_1 \textit{ else } e_2 : \vec{\tau}} \textit{(if)} \qquad \frac{\Gamma \vdash^e e : \vec{\tau}}{\Gamma, x : \rho \vdash^e e : \vec{\tau}} \textit{(weak)} \\
\text{provided } \rho \textit{ is not } !\textit{packet} \\
\\
\frac{\{x_1 : \rho_1, \dots, x_k : \rho_k\} \vdash^e e : \sigma_1 \times \dots \times \sigma_j}{\overset{d}{\vdash} (\sigma_1, \dots, \sigma_j) f(\rho_1 x_1, \dots, \rho_k x_k) \{e\}} \textit{(fundef)} \qquad \frac{}{\overset{d}{\vdash} \textit{packetQueue } q;} \textit{(qdef)} \\
\\
\frac{\overset{d}{\vdash} d_1 \quad \dots \quad \overset{d}{\vdash} d_n}{\overset{p}{\vdash} d_1 \dots d_n} \textit{(prog)}
\end{array}$$

Figure 3: Linear typing rules

during the call, and that the alias dies before the function returns. It is the use of the environment sequencing operator “;” in the consequent of the *(apseq)* rule that facilitates the creation of aliases from an owning variable. (Recall that *!packet*; *packet* = *packet*).

The *(ret)* rule types its subexpressions in environments $\Gamma_1, \dots, \Gamma_j$ respectively, and uses the “+” operator to combine these environments in the antecedent. The “+” operator is used in this context to prevent owning references being duplicated. We prevent values of type *!packet* being returned because the enclosing scope may contain the corresponding owning variable.

The type signatures of built-in functions are given in Figure 4. Note that inter-thread communication primitives, *q.enq* and *q.deq*, enqueue and dequeue values of type *packet*

```

int <arith-op> (int x, int y);
bool <rel-op> (int x, int y);
packet new(int size);
packet q.deq();
int read(!packet p, int offset);
void update(!packet p, int offset, int value);

bool <bool-op> (bool x, bool y);
void kill(packet p);
void q.enq(packet p);

```

Figure 4: Type signatures of the built-in functions expressed in C-like syntax

(i.e. owning references). Thus when transferring a packet between threads, the packet can be physically relocated.

We assume the existence of special queues, `receive` and `transmit` which represent the external network interface. `receive.deq()` reads a packet from the network; `transmit.enq(p)` schedules packet `p` for transmission. In reality a network appliance would typically contain multiple network interfaces and one would expect a variety of more sophisticated range of read/write calls. However, since this does not affect the fundamental concepts expressed in PACLANG, we stick to the simple interface above for the purposes of this paper.

The following two examples (in de-sugared form) help to illustrate PACLANG’s linear type system. First consider the following valid PACLANG expression:

```
let packet x = new(10) in
let () = update(x,0,5) in
let () = update(x,1,10) in
let () = transmit.enq(x) in
return ()
```

This code will type check successfully as, although `x` is accessed multiple times, `update` only requires a local alias to `x` (*!packet*) and so does not prevent `x` being passed as type *packet* to `transmit.enq`.

Conversely, consider the following *invalid* PACLANG expression:

```
let packet x = new(10) in
let () = kill(x) in
let () = transmit.enq(x) in
return ()
```

The type system rejects this program since `x` is passed twice as type *packet*: once to the `kill` function and then again to the `transmit.enq` function. Once `x` has been passed as type *packet*, it effectively disappears and cannot be used again.

4 Operational Semantics

In this section we define the language formally by means of a small-step operational semantics modelled on the Chemical Abstract Machine [2]. States of the machine are an unordered multiset of reacting molecules. We let W range over machine states. The “|” operator separates molecules in the multiset solution (cf. parallel composition). Molecules take the following forms:

$$\begin{array}{ll}
 M \leftarrow & \{ \{ \vec{x} \} e \}_f \quad (\text{function definition}) \\
 & | \langle \langle Q \rangle \rangle_q \quad (\text{queue}) \\
 & | \langle H \rangle \quad (\text{shared heap}) \\
 & | (e, \Sigma) \quad (\text{thread: expression and its stack frames})
 \end{array}$$

We let α range over an infinite set of heap-bound packet reference values⁵. We extend the syntactic category, v , allowing values also to be packet references, α , in addition to integer and boolean constants.

⁵There is no syntactic (compile-time) form of such values.

Core language primitives:

$$\begin{aligned}
\llbracket \{\vec{x}\}e_1 \rrbracket_f \mid \llbracket \text{let } \vec{y} = f(\vec{v}) \text{ in } e_2, \Sigma \rrbracket &\rightsquigarrow \llbracket \{\vec{x}\}e_1 \rrbracket_f \mid \llbracket e_1\{\vec{v}/\vec{x}\}, (\{\vec{y}\}e_2) \bullet \Sigma \rrbracket && \text{(user-call)} \\
\llbracket \text{let } x = v_1 \text{ op } v_2 \text{ in } e, \Sigma \rrbracket &\rightsquigarrow \llbracket e\{v/x\}, \Sigma \rrbracket \text{ where } v = v_1 \text{ op } v_2 && \text{(basic-op)} \\
\llbracket \text{return } \vec{v}, (\{\vec{x}\}e) \bullet \Sigma \rrbracket &\rightsquigarrow \llbracket e\{\vec{v}/\vec{x}\}, \Sigma \rrbracket && \text{(ret)} \\
\llbracket \text{if true then } e_1 \text{ else } e_2, \Sigma \rrbracket &\rightsquigarrow \llbracket e_1, \Sigma \rrbracket && \text{(if}_1\text{)} \\
\llbracket \text{if false then } e_1 \text{ else } e_2, \Sigma \rrbracket &\rightsquigarrow \llbracket e_2, \Sigma \rrbracket && \text{(if}_2\text{)}
\end{aligned}$$

Packet manipulation:

$$\begin{aligned}
\langle H \rangle \mid \llbracket \text{let } x = \text{new}(i) \text{ in } e, \Sigma \rrbracket &\rightsquigarrow \langle H[\alpha \mapsto \text{empty}_i] \rangle \mid \llbracket e\{\alpha/x\}, \Sigma \rrbracket && \text{(new)} \\
&&& \text{where } \alpha \text{ is a fresh name} \\
\langle H \rangle \mid \llbracket \text{let } () = \text{kill}(\alpha) \text{ in } e, \Sigma \rrbracket &\rightsquigarrow \langle H \setminus \alpha \rangle \mid \llbracket e, \Sigma \rrbracket && \text{(kill)} \\
\langle H \rangle \mid \llbracket \text{let } x = \text{read}(\alpha, i) \text{ in } e, \Sigma \rrbracket &\rightsquigarrow \langle H \rangle \mid \llbracket e\{H(\alpha)(i)/x\}, \Sigma \rrbracket && \text{(read)} \\
\langle H \rangle \mid \llbracket \text{let } () = \text{update}(\alpha, i, v) \text{ in } e, \Sigma \rrbracket &\rightsquigarrow \langle H[\alpha \mapsto H(\alpha)[i \mapsto v]] \rangle \mid \llbracket e, \Sigma \rrbracket && \text{(update)}
\end{aligned}$$

Queues:

$$\begin{aligned}
\langle \langle \alpha \bullet Q \rangle \rangle_q \mid \llbracket \text{let } x = q.\text{deq}() \text{ in } e, \Sigma \rrbracket &\rightsquigarrow \langle \langle Q \rangle \rangle_q \mid \llbracket e\{\alpha/x\}, \Sigma \rrbracket && \text{(deq)} \\
\langle \langle Q \rangle \rangle_q \mid \llbracket \text{let } () = q.\text{enq}(\alpha) \text{ in } e, \Sigma \rrbracket &\rightsquigarrow \langle \langle Q \bullet \alpha \rangle \rangle_q \mid \llbracket e, \Sigma \rrbracket && \text{(enq)}
\end{aligned}$$

Figure 5: Operational Semantics for PACLANG

A *function molecule*, $\llbracket \{\vec{x}\}e \rrbracket_f$, corresponds to a source-level function definition for f , with formal parameters, \vec{x} , and body e . Function molecules themselves do not perform computation, they exist merely to provide their bodies to callers.

A *queue molecule*, $\langle \langle Q \rangle \rangle_q$ has name q and contents Q , where Q takes the form of a list of elements separated by the associative operator \bullet . An empty queue is written ϵ . In the semantics presented here, queues are unbounded with synchronous (blocking) behaviour.

A *heap molecule*, $\langle H \rangle$, maps *packet references*, α , onto packets, p . A packet, p , is a finite mapping from integer offsets, i , to integer values, $p(i)$ ⁶. We define empty_k as the empty packet of size k : $\text{empty}_k = \{0 \mapsto 0, \dots, (k-1) \mapsto 0\}$. We write $\langle H \setminus \alpha \rangle$ for the heap which is as $\langle H \rangle$ but no longer contains an entry for α . The heap, which is as $\langle H \rangle$ but maps α to p , is written $\langle H[\alpha \mapsto p] \rangle$. Thus, updating the i^{th} element of the packet referenced by α to value 42 is written $\langle H[\alpha \mapsto H(\alpha)[i \mapsto 42]] \rangle$. Note that the heap molecule need not be implemented as a single shared memory. The central result of this paper is that, due to PACLANG's unique ownership property, a compiler is free to map the packets in $\langle H \rangle$ to the multiple memory banks of a distributed memory architecture.

A *thread molecule*, $\llbracket e, \Sigma \rrbracket$, corresponds to one of the **main**-functions (see Section 2) where e is an *evaluation state* and Σ is a stack of (function-return) continuations as in Landin's SECD machine. As with queue molecules, we use the associative operator, \bullet , to separate elements on the stack and write ϵ for the empty stack. The syntax of evaluation states, e , is essentially the same as the syntax of expressions given in Figure 2, save that as noted above unbound variables do not occur but values may also be packet references, α .

⁶If the offset is outside the domain of the packet the result is undefined.

We write “ $\{\vec{x}\}e$ ” for the continuation that binds variables \vec{x} and continues as evaluation state e .

We write $init(P)$ to denote the initial state for program P . It consists of a function molecule, $\{\{\vec{x}\}e\}_f$, for each function definition, $f(\vec{x})\{e\}$, in P ; a thread molecule, $\langle e_j, \epsilon \rangle$, for each definition $main_j()\{e_j\}$ in P ; an empty queue molecule $\langle\langle \epsilon \rangle\rangle_q$ for each queue definition in P ; and an empty heap $\langle \rangle$.

The semantics is given as a transition relation, \rightarrow , on states (multisets of molecules). For convenience, we define \rightarrow in terms of an operator, \rightsquigarrow , which operates on only a subset of the molecules in a complete state:

$$\frac{M_1 \mid \dots \mid M_n \rightsquigarrow M'_1 \mid \dots \mid M'_n}{\Delta_1 \mid \dots \mid \Delta_k \mid M_1 \mid \dots \mid M_n \rightarrow \Delta_1 \mid \dots \mid \Delta_k \mid M'_1 \mid \dots \mid M'_n} (\rightsquigarrow)$$

where $\Delta_1, \dots, \Delta_k$ are molecules that form part of the state, but are not of interest to the \rightsquigarrow -transition. Figure 5 gives transition rules for PACLANG. Capture-free substitution of values \vec{v} for variables \vec{x} in expression e is written $e\{\vec{v}/\vec{x}\}$.

The outside world interacts with the program by placing network packets on the **receive** queue and removing network packets from the **transmit** queue. In reality these queues would be filled and emptied by hardware. However, to avoid extra semantic rules to express this, we can model a generic network interface as if it were two system-provided threads, one of which dequeues packets from **transmit** and frees them; the other of which constructs and allocates appropriate packets and enqueues them on **receive**.

4.1 Unique Ownership

Recall from Section 1 that the primary motivation for our type-system is to ensure that every packet in the heap is referenced by exactly one thread. We have argued that this *unique ownership* property makes practical the mapping of a high-level program to a variety of diverse hardware architectures (see Section 1.1). In this section we define unique ownership formally.

Let, $pr(M)$, be the set of packet references, α , that occur in molecule M (This is empty for function molecules, but generally non-empty for thread and queue molecules; $pr(\cdot)$ is not applied to the heap molecule). We write $S_1 \uplus S_2$ to be the union of sets S_1 and S_2 under the condition that the elements of S_1 and S_2 are disjoint ($S_1 \cap S_2 = \emptyset$).

Definition 1 (Unique Ownership)

$$UniqueOwn(M_1 \mid \dots \mid M_n \mid \langle H \rangle) \iff pr(M_1) \uplus \dots \uplus pr(M_n) = \text{dom } H$$

Theorem 1 (Preservation of Unique Ownership) *If a program, P , types then any state, W , resulting from execution of P satisfies unique ownership:*

$$\Vdash^P P \wedge init(P) \xrightarrow{*} W \Rightarrow UniqueOwn(W)$$

A proof of this theorem is presented at the end of this document (see Appendix A).

5 Case Study: An IPv4 Packet Forwarder

In this section we outline the PACLANG implementation of the IP (version 4) uni-cast packet forwarding algorithm. The precise details of IP packet forwarding are not the focus of this section and, indeed, due to space constraints are not described here (the interested reader is referred to the IETF standards [13, 14]). Instead, the purpose of the example is to outline the implementation of a realistic application, demonstrating that (i) packets are indeed treated as linear objects in real-life applications; and (ii) that the PACLANG type-system is intuitive and easy to use.

Figure 6 shows the top-level structure of our PACLANG IPv4 packet forwarder; we omit many of the function bodies, giving only their signatures. Function `eth_readAndClassify()` reads ethernet packets from the network interface and classifies them according to their type. Function `ip_lookup` performs a route lookup (on an externally defined routing table) returning both the next hop IP address and the network interface on which to forward the packet. (Note that since the external routing table component only contains `ints`, we do not have to worry about it breaking unique ownership.)

The program is divided into two separate threads (`main1` and `main2`). The `main1` thread sits in a tight loop, repeatedly calling the read/classify function. Packets which require ICMP processing are written to the `icmpQ` queue. The `main2` thread reads packets from `icmpQ` and performs the ICMP processing. The details of this thread are omitted. An example of a side-effecting function that manipulates packets is given by the `ip_dec_ttl` function which decrements the IP Time-To-Live field.

The PACLANG code given here demonstrates that our linear type system naturally models the way packets flow through the system as a whole. Note particularly that aliases (*!packet* types) allow us to express imperative operations on packets without requiring a packet’s owning reference to be tediously threaded through all imperative operations (cf. the Clean language [8]).

Let us now consider compiling the PACLANG IPv4 packet forwarder to the architecture shown in Figure 1.1. Due to PACLANG’s unique ownership property a compiler is free to choose between either (say) implementing both threads on a single IXP2400, implementing the queue using a shared memory; or placing `main1` on the first IXP2400, and `main2` on the second, passing packets by value between the two. In the latter case the queue itself could be located on either the first or second network processor and, further, as soon as a packet is transmitted from the first to the second IXP it can be immediately deallocated from the memory of the first. Although PACLANG makes these tradeoffs very easy to reason about, recall that this is not the case in the framework of traditional imperative languages (e.g. C/C++)—see Section 1.1.

6 Related Work

There has been a great deal of work in the area of linear types. Of this, our research is most closely related to Kobayashi’s quasi-linear type system [12]. Aliases in our work (*!packet*) correspond to δ -uses in the quasi-linear framework. Similarly, owning references (of type *packet*) correspond to 1-use types. There are three major differences between our work and Kobayashi’s. Firstly, we do not deal with higher-order functions: these add much of the perceived complication of linear types; hence this makes our type system

```

int eth_type(!packet p); // read type-field of ethernet frame
void eth_gen_arp_response(packet p); // generate response to ARP query
void eth_proc_arp_response(packet p); // process ARP response packet
void ip_local_packet(packet p); // send packet to local network stack
void ip_process_options(!packet p); // process IP options
void ip_transmit(packet p, int tx_interface, int ip_nexthop);
    // do MAC-level processing, get MAC address of next hop and transmit
int eth_rx_iface(!packet p); // which interface was packet from?

packetQueue icmpQ; // queue for packets which require ICMP processing

(int, int) ip_lookup(int dest_ip_addr);
    // lookup IP address and network interface for next-hop in route

void eth_readAndClassify() {
    packet p = receive.deq();
    switch (eth_type(p)) of {
        case ARP_QUERY: eth_gen_arp_response(p);
        case ARP_RES: eth_proc_arp_response(p);
        case IP: ip_forward(p);
        default: kill(p); }}

void ip_forward(packet p) {
    if (ip_packet_for_localhost(p)) ip_local_packet(p);
    else {
        (int nexthop, int iface) = ip_lookup(eth_ip_dest_addr(p));
        if (nexthop == 0) icmpQ.enq(p);
        else { if (iface==eth_rx_iface(p)) icmpQ.enq(p);
                else ip_forward_packet(p,iface,nexthop); }}}

void ip_forward_packet(Packet p, int iface, int nexthop) {
    if (ip_dec_ttl(p)) icmpQ.enq(p);
    else { ip_process_options(p);
            ip_transmit(p,tx_iface,ip_nexthop); }}}

bool ip_dec_ttl(!packet p) {
    if (p[ip_ttl]==0) return true;
    else { p[ip_ttl] = p[ip_ttl]-1; return false; }}

// thread 1 -- read, process, transmit loop:
void main1() { eth_readAndClassify(); main1(); }

// thread 2 -- process icmp packets:
void main2() { packet p = icmpQ.deq(); ... ICMP processing ...; main2(); }

```

Figure 6: Top-level structure of an IPv4-over-ethernet packet forwarding engine

much simpler and, we argue, more accessible to industrial software engineers. Secondly, we work with a concurrent language. Thirdly, whereas Kobayashi is primarily concerned with eliminating garbage collection, we use linear types to assist in the mapping of high-level specifications to distributed memory architectures.

Clarke et al. have proposed *ownership types* [5, 4] for object-oriented languages. Each ownership type is annotated with a *context declaration*. The type system ensures that types with different declared contexts cannot alias. Ownership types have been applied to the detection of data races and deadlocks [3].

The `auto_ptr` template in the C++ standard library [15] exploits the notion of “unique ownership” *dynamically*. Suppose `p1` and `p2` are of type `auto_ptr`, then the assignment `p2 = p1` copies `p1` to `p2` as usual but also nulls `p1` so it cannot be used again. Similarly, if an `auto_ptr` is passed to a function it is also nulled. In this framework violation of the unique ownership constraint is only detected at run-time by a null pointer check. However, the widespread use of `auto_ptr`s in practical C++ programming suggests that working software engineers find the notion of object ownership intuitive and useful.

Region-based memory management allows a programmer to declare *regions*: unbounded areas of memory where values can be allocated. When the objects in a region, r , are no-longer required, r (and thus all the objects it contains) is deallocated in one fell swoop. A number of researchers have shown how type-systems which control the allocation/deallocation of regions can support safe explicit memory management without the need for garbage collection [16, 7]. Although regions provide a structured way to deal with explicit memory allocation/deallocation, they do not enforce the linear “unique ownership” property that PACLANG exploits. Walker and Watkins have explored the relationship between using regions and linear types for memory management in more detail [18].

7 Conclusions and Future Work

We have formally defined the PACLANG language and claimed, with reference to a case-study, that it is capable of specifying realistic packet processing applications. Our linear type system enforces a unique ownership property whilst allowing software engineers to program in a familiar style. We have argued that unique ownership greatly simplifies the compilation of a high-level program to a diverse range of NP architectures.

There are a number of ways in which PACLANG must be extended if it is to be used on an industrial-scale. For example, the type system needs to be able to differentiate between different types of *packet* (e.g. IP packets, UDP packets etc.); support for structured types (such as records and tuples) and a module system are also required. However, much of this work amounts to engineering rather than research effort. We believe that PACLANG, as presented in this paper, captures the essence of packet-processing, providing a solid foundation for the development of an industrial-scale packet processing framework.

Our current work focuses on the implementation of a PACLANG compiler which targets a variety of platforms, including Network Processors. This implementation effort will provide insights that can be fed back into the PACLANG language design.

We hope that the ideas presented in this paper can be applied to the automatic partitioning of high-level code across multi-core architectures more generally (i.e. not just Network Processors). Since industrial trends suggest that such architectures will

become more prevalent (as silicon densities continue to increase) we believe that this is an important topic for future research.

Acknowledgements

The authors wish to thank Tim Griffin and Lukas Kencl for their valuable comments and suggestions. We are grateful for the generous support of Microsoft Research, who funded the studentship of the first author.

A Preservation of Unique Ownership

In Section 4.1 we claimed that if a PACLANG program is accepted by the type-checker of Section 3 and executed using the operational semantics of Section 4 then all intermediate program states must be such that every packet in the heap is referenced by exactly one molecule. In this appendix, we prove that this is indeed the case.

The *preservation of unique ownership theorem* can be stated formally as follows:

$$\vdash^p P \wedge \text{init}(P) \rightarrow^* W \Rightarrow \text{UniqueOwn}(W)$$

where *UniqueOwn* is defined in Section 4.1.

To prove this theorem, we first strengthen *UniqueOwn* to a stronger *WellFormed* property (Section A.1). We show that any well-formed state must satisfy unique ownership (Section A.1.1). Finally, we show that wellformedness is preserved by each of the transitions given in the operational semantics of Section 4 and that the initial state of a program will always be well-formed. We thus conclude that the preservation of unique ownership theorem is indeed correct.

A.1 Type-Checking States

We start by extending the type system of Section 3 so that it can type check states of the operational semantics. Recall from Section 5 that evaluation states of thread molecules extend expressions by allowing packet references, α , as values. We extend the typing rules of Figure 3 accordingly, by allowing type environments Γ to include capabilities of the form $\alpha : \sigma$ in addition to capabilities of the form $x : \sigma$. In the case of a packet reference, α , the type will always be *packet* or *!packet*. Just as a capability, $x : \sigma$ allows x to be used as a type σ , a capability, $\alpha : \sigma$ allows packet references α to be used as type σ . To reflect this, we add a typing rule for packet references as follows:

$$\frac{}{\{\alpha : \tau\} \vdash^v \alpha : \tau} \text{ (p-reference)} \quad \text{where } \tau \in \{\text{packet}, \text{!packet}\}$$

To check that a thread molecule (e, Σ) is well typed, we have to type check the stack Σ as well as the evaluation state e . Stacks can be type-checked using the following rules:

$$\frac{\Gamma_1, \vec{x} : \vec{\rho} \vdash^e e : \vec{\sigma} \quad \Gamma_2 \vdash^\Sigma \vec{\sigma} \triangleright \Sigma}{\Gamma_1; \Gamma_2 \vdash^\Sigma \vec{\rho} \triangleright \{\vec{x}\}e \bullet \Sigma} \text{ (cont-step)} \quad \frac{}{\emptyset \vdash^\Sigma \rho \triangleright \epsilon} \text{ (cont-empty)}$$

We write $\Gamma \stackrel{\Sigma}{\vdash} \vec{\rho} \triangleright \Sigma$ to mean that the continuation on top of the stack expects to have values of types $\vec{\rho}$ returned to it and requires the capabilities given in Γ . Since we know that the evaluation state, e , will be fully evaluated before executing the continuations in Σ we use “;” to sequence Γ_1 and Γ_2 in the consequent of rule *cont-step*.

We are now in a position to type-check molecules, and hence to type check the multiset of molecules which constitutes a program state.

Definition 2 We write, $\Gamma \vdash^m M$, to mean that the molecule M is well typed, and requires the capabilities in Γ . This relation is defined by the following rules, one for each form of molecule (except heaps, which are treated specially):

$$\frac{\Gamma_1 \vdash^e e : \vec{\rho} \quad \Gamma_2 \stackrel{\Sigma}{\vdash} \vec{\rho} \triangleright \Sigma}{\Gamma_1; \Gamma_2 \vdash^m \langle e, \Sigma \rangle} \text{ (thread)}$$

$$\frac{}{\{\alpha_1 : \text{packet}, \dots, \alpha_i : \text{packet}\} \vdash^m \langle \langle \alpha_1 \bullet \dots \bullet \alpha_i \rangle \rangle_q} \text{ (queue)}$$

$$\frac{\{\vec{x} : \vec{\rho}\} \vdash^e e : \vec{\sigma}}{\emptyset \vdash^m \llbracket \{\vec{x}\} e \rrbracket_f} \text{ (fun) provided } \mathcal{F}(f) = \vec{\rho} \rightarrow \vec{\sigma}$$

The *(thread)* rule checks that the type of the expression state, e , is the type that the stack, Σ , expects to be returned to it. The *(queue)* rule states that a queue owns of all of the packet references contained in the queue. The *(fun)* rule does not require ownership of any packets. It merely checks that its body expression is well typed under the rules given in Figure 3. Recall that function molecules do not themselves perform computation and cannot have any free variables.

We now define our *WellFormed*(\cdot) property over states of the operational semantics:

Definition 3 (WellFormed)

$$\frac{\Gamma_1 \vdash^m M_1 \quad \dots \quad \Gamma_n \vdash^m M_n \quad \Gamma_1 + \dots + \Gamma_n = \{\alpha : \text{packet} \mid \alpha \in \text{dom } H\}}{\text{WellFormed}(\langle H \rangle \mid M_1 \mid \dots \mid M_n)} \text{ (ok)}$$

If *WellFormed*(W) holds, then every molecule in the program types, and every packet in the heap is referenced by exactly one such molecule.

A.1.1 Relating *WellFormed* to *UniqueOwn*

From the rules defining \vdash^m we can see that, if a molecule M types with an environment Γ , then the packet references contained in Γ are exactly those packet references contained in M . That is:

$$\Gamma \vdash^m M \Rightarrow ((\exists \rho. (\alpha \mapsto \rho) \in \Gamma) \iff \alpha \in \text{pr}(M))$$

This follows from the fact that any packet references α in M must be free (there are no binding constructs for packet references, only for the variables that hold them). We can thus observe from the type rules that any packet references that appear in M must

find their way into the top level type environment, Γ . We can similarly see that no packet reference can appear in Γ unless it is mentioned somewhere in M .

Given this, and the definition of ‘+’, we see that *WellFormed* is indeed a stronger condition than *UniqueOwn*. That is:

$$\text{WellFormed}(W) \Rightarrow \text{UniqueOwn}(W)$$

A.2 Lemmas

In this section we prove a few lemmas that turn out to be useful for our proof.

A.2.1 Lemma: Substitution

$$\begin{aligned} \Gamma_1 \vdash^v v : \rho \quad \wedge \quad \Gamma_2, x : \rho \vdash^e e : \vec{\sigma} \\ \Rightarrow \\ \Gamma_1 + \Gamma_2 \vdash^e e\{v/x\} : \vec{\sigma} \\ \text{if } \text{dom}(\Gamma_1) \cap \text{dom}(\Gamma_2, x : \rho) = \emptyset \end{aligned}$$

If an expression e expects a variable x with type ρ , and a value v has type ρ , then we can substitute v for x in e without changing e 's type. The capabilities required by the new expression will be e 's own capabilities Γ_2 and the capabilities Γ_1 required in order to evaluate v .

The proof of this property turns out to be straightforward. v must either be a variable x' , a packet reference α or an integer or boolean constant. We consider each of these cases separately:

- If v is a variable x' , then $\Gamma_1 = \{x' : \rho\}$. We know that x' cannot appear in Γ_2 , thus $\Gamma_2 + \Gamma_1$ is simply $\Gamma_2, x' : \rho$. The effect is simply to substitute one variable name for another. The type rules are insensitive to name choice and so the consequent holds trivially.
- The case where v is a packet reference, α , is the same as the case for a variable. The type rules treat packet references and variables identically and so replacing a variable with a packet reference cannot have an effect on the type. It can be helpful to think of a packet reference as being a variable that is bound by the heap, rather than by a function argument or let expression.
- If v is an integer or boolean constant then $\Gamma_1 = \emptyset$ and ρ is non-linear. From observing the rules for ‘ \vdash^v ’ we can see that substituting a constant for an appropriately typed variable leaves the type unchanged, and merely removes x from the environment. We can also observe that, if x does not appear in an expression, then that expression cannot require x in its environment. This is exactly what we want.

By applying this rule repeatedly we see that the following, more general, property holds:

$$\begin{aligned} \Gamma_1 \vdash^v v_1 : \rho_1 \quad \wedge \quad \dots \quad \wedge \quad \Gamma_j \vdash^v v_j : \rho_j \quad \wedge \quad \Gamma, x_1 : \rho_1, \dots, x_j : \rho_j \vdash^e e : \vec{\sigma} \\ \Rightarrow \\ \Gamma + \Gamma_1 + \dots + \Gamma_j \vdash^e e\{v_1/x_1, \dots, v_j/x_j\} : \vec{\sigma} \\ \text{if } \text{dom}(\Gamma_1 + \dots + \Gamma_j) \cap \text{dom}(\Gamma, x_1 : \rho_1, \dots, x_j : \rho_j) = \emptyset \end{aligned}$$

A.2.2 Lemma: associativity of ;

$$(\Gamma_1; \Gamma_2); \Gamma_3 = \Gamma_1; (\Gamma_2; \Gamma_3)$$

Proof by inspection of the rules defining ‘;’.

A.2.3 Lemma: commutativity of +

$$\Gamma_1 + \Gamma_2 = \Gamma_2 + \Gamma_1$$

Proof by inspection of the rules defining ‘+’.

A.3 Preservation of WellFormedness

Having strengthened *UniqueOwn* to *WellFormed* we can similarly strengthen preservation of unique ownership to preservation of wellformedness. We thus seek to prove the following theorem:

$$\stackrel{P}{\vdash} P \wedge \text{init}(P) \rightarrow^* W \Rightarrow \text{WellFormed}(W)$$

If the program P is well typed, then any state W reachable from the initial state $\text{init}(P)$ through evaluation transitions must be well formed (and thus satisfy unique ownership).

The proof of this theorem proceeds by induction over the number of evaluation steps between $\text{init}(P)$ and W .

The base case (zero steps) is easy. Recall from Section 4 that, in the initial state $\text{init}(P)$, the heap is empty, all queues are empty, all stacks are empty, and the bodies of the threads are the bodies of the (closed and well typed) bodies of the *main* functions. It thus follows trivially that all such initial states must be well formed.

The inductive step is somewhat harder. We must show that, if a state W is well formed, and the semantics can transform the state W into a new state W' then W' must also be well formed. That is:

$$\text{WellFormed}(W) \wedge W \rightarrow W' \Rightarrow \text{WellFormed}(W')$$

In the following subsections we prove that this holds for all rules defining \rightarrow . For each rule, we assume well-formedness for the antecedent and prove well-formedness for the consequent.

A.3.1 Rule (let):

$$\vec{\Delta} \mid [\{\vec{x}\}e_f]_f \mid (\text{let } \vec{y} = f(\vec{v}) \text{ in } e_b, \Sigma) \rightarrow \vec{\Delta} \mid [\{\vec{x}\}e_f]_f \mid (e_f\{\vec{v}/\vec{x}\}, \{\vec{y}\}e_b \bullet \Sigma)$$

Referring to the property we wish to prove, we see that:

$$\begin{aligned} W &= \vec{\Delta} \mid [\{\vec{x}\}e_f]_f \mid (\text{let } \vec{y} = f(\vec{v}) \text{ in } e_b, \Sigma) \\ W' &= \vec{\Delta} \mid [\{\vec{x}\}e_f]_f \mid (e_f\{\vec{v}/\vec{x}\}, \{\vec{y}\}e_b \bullet \Sigma) \end{aligned}$$

If $WellFormed(W)$ holds, then it must be the result of an application of rule (*ok*). In which case, the following must hold:

$$\vec{\Delta} = \langle H \rangle \mid \Delta_1 \mid \dots \mid \Delta_n \quad (1)$$

$$\Gamma_{d_1} \vdash^m \Delta_1 \quad \dots \quad \Gamma_{d_n} \vdash^m \Delta_n \quad (2)$$

$$\emptyset \vdash^m \llbracket \{\vec{x}\}e_f \rrbracket_f \quad (3)$$

$$\Gamma \vdash^m (\mathbf{let} \vec{y} = f(\vec{v}) \mathbf{in} e_b, \Sigma) \quad (4)$$

$$\Gamma + \Gamma_{d_1} + \dots + \Gamma_{d_n} = \{\alpha : packet \mid \alpha \in dom(H)\} \quad (5)$$

If (4) holds, then it must be the result of an application of rule (*thread*) in which case the following must hold:

$$\Gamma = \Gamma_1; \Gamma_2 \quad (6)$$

$$\Gamma_1 \vdash^e (\mathbf{let} \vec{y} = f(\vec{v}) \mathbf{in} e_b) : \vec{\tau} \quad (7)$$

$$\Gamma_2 \vdash^\Sigma \vec{\tau} \triangleright \Sigma \quad (8)$$

If (7) holds, then it must be the result of an application of rule (*apseq*)⁷, in which case the following must hold:

$$\Gamma_1 = (\Gamma_{a_1} + \dots + \Gamma_{a_j}); \Gamma_b \quad (9)$$

$$\mathcal{F}(f) = \rho_1 \times \dots \times \rho_j \rightarrow \sigma_1 \times \dots \times \sigma_i \quad (10)$$

$$\Gamma_{a_1} \vdash^v v_1 : \rho_1 \quad \dots \quad \Gamma_{a_j} \vdash^v v_j : \rho_j \quad (11)$$

$$\Gamma_b, y_1 : \sigma_1, \dots, y_i : \sigma_i \vdash^e e_b : \vec{\tau} \quad (12)$$

If (3) holds, then it must be the result of an application of rule (*fun*). Given that we know (10), we can deduce that:

$$x_1 : \rho_1, \dots, x_j : \rho_j \vdash^e e_f : \sigma_1 \times \dots \times \sigma_i \quad (13)$$

From (5), (6), (9) and the definition of ‘;’ we can deduce that $dom(\Gamma_{a_1} + \dots + \Gamma_{a_j})$ contains only packet references (it cannot contain variables). Thus:

$$dom(\Gamma_{a_1} + \dots + \Gamma_{a_j}) \cap dom(x_1 : \rho_1, \dots, x_j : \rho_j) = \emptyset \quad (14)$$

By (13), (11), (14) and Lemma (A.2.1) we can deduce that:

$$\Gamma_{a_1} + \dots + \Gamma_{a_j} \vdash^e e_f \{\vec{v}/\vec{x}\} : \tau_{r_1} \times \dots \times \tau_{r_i} \quad (15)$$

Given (12) and (8), we can apply rule (*cont-step*), giving:

$$\Gamma_b; \Gamma_2 \vdash^\Sigma \vec{\sigma} \triangleright \{\vec{y}\}e_b \bullet \Sigma \quad (16)$$

From (6), (9) and Lemma (A.2.2) we can deduce that:

$$\Gamma = \Gamma_{a_1} + \dots + \Gamma_{a_j}; (\Gamma_b; \Gamma_2) \quad (17)$$

⁷It cannot be an application of rule (*weak*) because we know from (5) and (6) that all bindings in Γ_1 must be for type *packet* or *!packet*.

Given (15), (16) and (17) we can apply rule (*proc*), giving:

$$\Gamma \vdash^m (e_f\{\vec{v}/\vec{x}\}, \{\vec{y}\}e_b \bullet \Sigma) \quad (18)$$

Given (1), (2), (3), (18), and (5) we can apply rule (*ok*), giving:

$$WellFormed(\vec{\Delta} \mid [\{\vec{x}\}e_f]_f \mid (e_f\{\vec{v}/\vec{x}\}, \{\vec{y}\}e_b \bullet \Sigma)) \quad (19)$$

which is what we wanted.

A.3.2 Rule (*return*):

$$\vec{\Delta} \mid (\mathbf{return} \vec{v}, \{\vec{x}\}e \bullet \Sigma) \rightarrow \vec{\Delta} \mid (e\{\vec{v}/\vec{x}\}, \Sigma)$$

Referring to the property we wish to prove, we see that:

$$\begin{aligned} W &= \vec{\Delta} \mid (\mathbf{return} \vec{v}, \{\vec{x}\}e \bullet \Sigma) \\ W' &= \vec{\Delta} \mid (e\{\vec{v}/\vec{x}\}, \Sigma) \end{aligned}$$

If $WellFormed(W)$ holds, then it must be the result of an application of rule (*ok*), in which case the following must also hold:

$$\vec{\Delta} = \langle H \rangle \mid \Delta_1 \mid \dots \mid \Delta_n \quad (1)$$

$$\Gamma_{d_1} \vdash^m \Delta_1 \quad \dots \quad \Gamma_{d_n} \vdash^m \Delta_n \quad (2)$$

$$\Gamma \vdash^m (\mathbf{return} \vec{v}, \{\vec{x}\}e \bullet \Sigma) \quad (3)$$

$$\Gamma + \Gamma_{d_1} + \dots + \Gamma_{d_n} = \{\alpha : packet \mid \alpha \in dom(H)\} \quad (4)$$

If (3) holds then it must be the result of an application of rule (*proc*), in which case the following must hold:

$$\Gamma = \Gamma_1; \Gamma_2 \quad (5)$$

$$\Gamma_1 \vdash^e \mathbf{return} \vec{v} : \vec{\tau} \quad (6)$$

$$\Gamma_2 \vdash^\Sigma \vec{\tau} \triangleright \{\vec{x}\}e \bullet \Sigma \quad (7)$$

If (6) holds, then it must be the result of an application of rule (*ret*)⁸, in which case:

$$\Gamma_1 = \Gamma_{v_1} + \dots + \Gamma_{v_i} \quad (8)$$

$$\Gamma_{v_1} \vdash^v v_1 : \tau_1 \quad \dots \quad \Gamma_{v_i} \vdash^v v_i :_e \tau_i \quad (9)$$

$$\text{no } \tau_1 \quad \dots \quad \tau_i \text{ is !}packet \quad (10)$$

Similarly, if (7) holds, then it must be the result of an application of rule (*cont-step*), in which case:

$$\Gamma_2 = \Gamma_e; \Gamma_s \quad (11)$$

$$\Gamma_e, \vec{x} : \vec{\tau} \vdash^e e : \sigma \quad (12)$$

$$\Gamma_s \vdash^\Sigma \sigma \triangleright \Sigma \quad (13)$$

⁸Again, we know from (4) and (5) that Γ_1 cannot contain non-linear types and thus the (*weaken*) rule cannot be applied.

From (4), (5), (8), and (10) we can see that Γ_1 can only contain packet reference bindings, and all such bindings have type *packet*. Hence, by (5) and the definition of ‘;’ we can observe that Γ_1 and Γ_2 have disjoint domains. Hence, by (11), (8) we can observe that:

$$\text{dom}(\Gamma_{v_1} + \dots + \Gamma_{v_i}) \cap \text{dom}(\Gamma_e, \vec{x} : \vec{\tau}) = \emptyset \quad (14)$$

From (9), (12), (14) and Lemma (A.2.1) we can deduce that:

$$\Gamma_e + \Gamma_{v_1} + \dots + \Gamma_{v_i} \vdash^e e\{\vec{v}/\vec{x}\} : \sigma \quad (15)$$

We know that none of the environments Γ_v can contain *!packet* types ((10)). It is thus possible to rewrite (15) as:

$$\Gamma_{v_1} + \dots + \Gamma_{v_i}; \Gamma_e \vdash^e e\{\vec{v}/\vec{x}\} : \sigma \quad (16)$$

Similarly, we can observe from (5), (8), (11) and associativity of environment sequencing (Lemma A.2.2) that:

$$\Gamma = (\Gamma_{v_1} + \dots + \Gamma_{v_i}; \Gamma_e); \Gamma_s \quad (17)$$

Using (16), (13) and (17) as assumptions, we can apply rule (*proc*) to give:

$$\Gamma \vdash^m \langle e\{\vec{v}/\vec{x}\}, \Sigma \rangle \quad (18)$$

Combining (18) with (1), (2) and (4), we can apply rule (*ok*) giving:

$$\text{WellFormed}(\vec{\Delta} \mid \langle e\{\vec{v}/\vec{x}\}, \Sigma \rangle) \quad (19)$$

as required.

A.3.3 Rules *if*₁ and *if*₂

We will give a proof for *if*₁. The proof for *if*₂ is virtually identical.

$$\vec{\Delta} \mid \langle \text{if } true \text{ then } e_1 \text{ else } e_2, \Sigma \rangle \rightarrow \vec{\Delta} \mid \langle e_1, \Sigma \rangle$$

Referring to the property we wish to prove, we see that:

$$\begin{aligned} W &= \vec{\Delta} \mid \langle \text{if } true \text{ then } e_1 \text{ else } e_2, \Sigma \rangle \\ W' &= \vec{\Delta} \mid \langle e_1, \Sigma \rangle \end{aligned}$$

If *WellFormed*(*W*) holds, then it must be the result of an application of rule (*ok*), in which case the following must also hold:

$$\vec{\Delta} = \langle H \rangle \mid M_1 \mid \dots \mid M_n \quad (1)$$

$$\Gamma_{d_1} \vdash^m M_1 \quad \dots \quad \Gamma_{d_n} \vdash^m M_n \quad (2)$$

$$\Gamma \vdash^m \langle \text{if } true \text{ then } e_1 \text{ else } e_2, \Sigma \rangle \quad (3)$$

$$\Gamma + \Gamma_{d_1} + \dots + \Gamma_{d_n} = \{\alpha : \text{packet} \mid \alpha \in \text{dom}(H)\} \quad (4)$$

If (3) holds, then it must be the result of an application of rule (*proc*), in which case the following must hold:

$$\Gamma = \Gamma_1; \Gamma_2 \quad (5)$$

$$\Gamma_1 \vdash^e \text{if } true \text{ then } e_1 \text{ else } e_2 : \vec{\sigma} \quad (6)$$

$$\Gamma_2 \vdash^\Sigma \vec{\sigma} \triangleright \Sigma \quad (7)$$

If (6) holds, then it must be the result of an application of rule (*if*) in which case:

$$\Gamma_1 = \Gamma_3; \Gamma_4 \quad (8)$$

$$\Gamma_3 \vdash^v true : bool \quad (9)$$

$$\Gamma_4 \vdash^e e_1 : \vec{\sigma} \quad (10)$$

$$\Gamma_4 \vdash^e e_2 : \vec{\sigma} \quad (11)$$

(9) can only be the result of an application of rule (*bool*), in which case:

$$\Gamma_3 = \emptyset \quad (12)$$

Hence, from (8) and the definition of ‘;’, we can see that:

$$\Gamma_1 = \Gamma_4 \quad (13)$$

Using (10), (7) and (5) as assumptions, and noting (13) we can apply rule (*proc*) to give:

$$\Gamma \vdash^m (\!| e_1, \Sigma \!|) \quad (14)$$

Combining (14) with (1), (2) and (4), we can apply rule (*ok*) giving:

$$WellFormed(\vec{\Delta} \mid (\!| e_1, \Sigma \!|)) \quad (15)$$

as desired.

A.3.4 Rule *new*

$$\vec{\Delta} \mid \langle H \rangle \mid (\!| \text{let } x = \text{new } i \text{ in } e, \Sigma \!|) \rightarrow \vec{\Delta} \mid \langle H[\alpha \mapsto \text{empty}_i] \rangle \mid (\!| e\{\alpha/x\}, \Sigma \!|) \quad \text{where } \alpha \text{ is fresh}$$

Referring to the property we wish to prove, we see that:

$$\begin{aligned} W &= \vec{\Delta} \mid \langle H \rangle \mid (\!| \text{let } x = \text{new } i \text{ in } e, \Sigma \!|) \\ W' &= \vec{\Delta} \mid \langle H[\alpha \mapsto \text{empty}_i] \rangle \mid (\!| e\{\alpha/x\}, \Sigma \!|) \end{aligned}$$

If $WellFormed(W)$ holds, then it must be the result of an application of rule (*ok*), in which case the following must also hold:

$$\vec{\Delta} = M_1 \mid \dots \mid M_n \quad (1)$$

$$\Gamma_{d_1} \vdash^m M_1 \quad \dots \quad \Gamma_{d_n} \vdash^m M_n \quad (2)$$

$$\Gamma \vdash^m (\!| \text{let } x = \text{new } i \text{ in } e, \Sigma \!|) \quad (3)$$

$$\Gamma + \Gamma_{d_1} + \dots + \Gamma_{d_n} = \{ \alpha' : \text{packet} \mid \alpha' \in \text{dom}(H) \} \quad (4)$$

If (3) holds, then it must be the result of an application of rule (*proc*), in which case the following must hold:

$$\Gamma = \Gamma_1; \Gamma_2 \quad (5)$$

$$\Gamma_1 \vdash^e \text{let } x = \text{new } i \text{ in } e : \vec{\tau} \quad (6)$$

$$\Gamma_2 \vdash^\Sigma \vec{\tau} \triangleright \Sigma \quad (7)$$

If (6) holds then it must be the result of an application of rule (*apseq*), in which case the following must hold:

$$\mathcal{F}(\text{new}) = \text{int} \rightarrow \text{packet} \quad (8)$$

$$\emptyset \vdash^v i : \text{int} \quad (9)$$

$$\Gamma_1, x : \text{packet} \vdash^e e : \vec{\tau} \quad (10)$$

By (10), Lemma (A.2.1), we can deduce the following for fresh α :

$$\Gamma_1 + \{\alpha : \text{packet}\} \vdash^e e\{\alpha/x\} : \vec{\tau} \quad (11)$$

Given (11) and (7) we can apply rule (*proc*), giving:

$$\Gamma_1 + \{\alpha : \text{packet}\}; \Gamma_2 \vdash^m \langle e\{\alpha/x\}, \Sigma \rangle \quad (12)$$

By Lemma (A.2.3) and (5), this can be rewritten as:

$$\{\alpha : \text{packet}\} + \Gamma \vdash^m \langle e\{\alpha/x\}, \Sigma \rangle \quad (13)$$

Given (4) and the freshness of α , we can deduce that:

$$\{\alpha : \text{packet}\} + \Gamma + \Gamma_{d_1} + \dots + \Gamma_{d_n} = \{\alpha' : \text{packet} \mid \alpha' \in \text{dom}(H[\alpha \mapsto \text{empty}_i])\} \quad (14)$$

Given (1), (2), (13), and (14) we can apply rule (*ok*), giving:

$$\text{WellFormed}(\vec{\Delta} \mid \langle H[\alpha' \mapsto \text{empty}_i] \rangle \mid \langle e\{\alpha'/x\}, \Sigma \rangle) \quad (15)$$

which is what we wanted.

A.3.5 Rule *kill*

$$\vec{\Delta} \mid \langle H \rangle \mid \langle \text{let } () = \text{kill } \alpha \text{ in } e, \Sigma \rangle \rightarrow \vec{\Delta} \mid \langle H \setminus \alpha \rangle \mid \langle e, \Sigma \rangle$$

Referring to the property we wish to prove, we see that:

$$W = \vec{\Delta} \mid \langle H \rangle \mid \langle \text{let } () = \text{kill } \alpha \text{ in } e, \Sigma \rangle$$

$$W' = \vec{\Delta} \mid \langle H \setminus \alpha \rangle \mid \langle e, \Sigma \rangle$$

If $\text{WellFormed}(W)$ holds, then it must be the result of an application of rule (*ok*), in which case the following must also hold:

$$\vec{\Delta} = M_1 \mid \dots \mid M_n \quad (1)$$

$$\Gamma_{d_1} \vdash^m M_1 \quad \dots \quad \Gamma_{d_n} \vdash^m M_n \quad (2)$$

$$\Gamma \vdash^m \langle \text{let } () = \text{kill } \alpha \text{ in } e, \Sigma \rangle \quad (3)$$

$$\Gamma + \Gamma_{d_1} + \dots + \Gamma_{d_n} = \{\alpha' : \text{packet} \mid \alpha' \in \text{dom}(H)\} \quad (4)$$

If **(3)** holds, then it must be the result of an application of rule *(proc)*, in which case the following must hold:

$$\Gamma = \Gamma_1; \Gamma_2 \quad (5)$$

$$\Gamma_1 \vdash^e \text{let } () = \text{kill } \alpha \text{ in } e : \vec{\tau} \quad (6)$$

$$\Gamma_2 \vdash^\Sigma \vec{\tau} \triangleright \Sigma \quad (7)$$

If **(6)** holds then it must be the result of an application of rule *(apseq)*, in which case the following must hold:

$$\Gamma_1 = \{\alpha : \text{packet}\}; \Gamma_3 \quad (8)$$

$$\mathcal{F}(\text{kill}) = \text{packet} \rightarrow () \quad (9)$$

$$\{\alpha : \text{packet}\} \vdash^v \alpha : \text{packet} \quad (10)$$

$$\Gamma_3 \vdash^e e : \vec{\tau} \quad (11)$$

Given **(11)** and **(7)** we can apply rule *(proc)*, giving:

$$\Gamma_3; \Gamma_2 \vdash^m \langle e, \Sigma \rangle \quad (12)$$

Given **(5)**, **(8)** and Lemma **(A.2.2)** we can deduce that:

$$\Gamma = \{\alpha : \text{packet}\}; (\Gamma_3; \Gamma_2) \quad (13)$$

We can observe that $\{\alpha : \text{packet}\}$ contains no *!packet* types. It is thus possible to rewrite **(13)** as:

$$\Gamma = \{\alpha : \text{packet}\} + (\Gamma_3; \Gamma_2) \quad (14)$$

Given **(4)** and **(14)** we can deduce that:

$$(\Gamma_3; \Gamma_2) + \Gamma_{d_1} + \dots + \Gamma_{d_n} = \{\alpha' : \text{packet} \mid \alpha' \in \text{dom}(H \setminus \alpha)\} \quad (15)$$

Using **(1)**, **(2)**, **(12)** and **(15)**, we can apply rule *(ok)*, giving:

$$\text{WellFormed}(\vec{\Delta} \mid \langle H \setminus \alpha \rangle \mid \langle e, \Sigma \rangle) \quad (16)$$

as required.

A.3.6 Rule *read*

$$\vec{\Delta} \mid \langle H \rangle \mid \langle \text{let } x = \text{read}(\alpha, i) \text{ in } e, \Sigma \rangle \rightarrow \vec{\Delta} \mid \langle H \rangle \mid \langle e\{H(\alpha)(i)/x\}, \Sigma \rangle$$

Referring to the property we wish to prove, we see that:

$$\begin{aligned} W &= \vec{\Delta} \mid \langle H \rangle \mid \langle \text{let } x = \text{read}(\alpha, i) \text{ in } e, \Sigma \rangle \\ W' &= \vec{\Delta} \mid \langle H \rangle \mid \langle e\{H(\alpha)(i)/x\}, \Sigma \rangle \end{aligned}$$

If $WellFormed(W)$ holds, then it must be the result of an application of rule (*ok*), in which case the following must also hold:

$$\vec{\Delta} = M_1 \mid \dots \mid M_n \quad (1)$$

$$\Gamma_{d_1} \vdash^m M_1 \dots \Gamma_{d_n} \vdash^m M_n \quad (2)$$

$$\Gamma \vdash^m (\text{let } x = \text{read}(\alpha, i) \text{ in } e, \Sigma) \quad (3)$$

$$\Gamma + \Gamma_{d_1} + \dots + \Gamma_{d_n} = \{\alpha' : \text{packet} \mid \alpha' \in \text{dom}(H)\} \quad (4)$$

If (3) holds, then it must be the result of an application of rule (*proc*), in which case the following must hold:

$$\Gamma = \Gamma_1; \Gamma_2 \quad (5)$$

$$\Gamma_1 \vdash^e \text{let } x = \text{read}(\alpha, i) \text{ in } e : \vec{\tau} \quad (6)$$

$$\Gamma_2 \vdash^\Sigma \vec{\tau} \triangleright \Sigma \quad (7)$$

If (6) holds then it must be the result of an application of rule (*apseq*), in which case the following must hold:

$$\mathcal{F}(\text{read}) = \text{packet} \times \text{int} \rightarrow \text{int} \quad (8)$$

$$\Gamma_1, x : \text{int} \vdash^e e : \vec{\tau} \quad (9)$$

By (9) and Lemma (A.2.1) we can deduce that:

$$\Gamma_1 \vdash^e e\{H(\alpha)(i)/x\} : \vec{\tau} \quad (10)$$

Given (10), (7) and (5), we can apply rule (*proc*), giving:

$$\Gamma \vdash^m (e\{H(\alpha)(i)/x\}, \Sigma) \quad (11)$$

Given (1), (2), (11), and (4), we can apply rule (*ok*), giving:

$$WellFormed(\vec{\Delta} \mid \langle H \rangle \mid (e\{H(\alpha)(i)/x\}, \Sigma)) \quad (12)$$

as required.

A.3.7 Rule update

$$\vec{\Delta} \mid \langle H \rangle \mid (\text{let } () = \text{update}(\alpha, i, v) \text{ in } e, \Sigma) \rightarrow \vec{\Delta} \mid \langle H[\alpha \mapsto H(\alpha)[i \mapsto v]] \rangle \mid (e, \Sigma)$$

Referring to the property we wish to prove, we see that:

$$W = \vec{\Delta} \mid \langle H \rangle \mid (\text{let } () = \text{update}(\alpha, i, v) \text{ in } e, \Sigma)$$

$$W' = \vec{\Delta} \mid \langle H[\alpha \mapsto H(\alpha)[i \mapsto v]] \rangle \mid (e, \Sigma)$$

If $WellFormed(W)$ holds, then it must be the result of an application of rule (*ok*), in which case the following must also hold:

$$\vec{\Delta} = M_1 \mid \dots \mid M_n \quad (1)$$

$$\Gamma_{d_1} \vdash^m M_1 \dots \Gamma_{d_n} \vdash^m M_n \quad (2)$$

$$\Gamma \vdash^m (\text{let } () = \text{update}(\alpha, i, v) \text{ in } e, \Sigma) \quad (3)$$

$$\Gamma + \Gamma_{d_1} + \dots + \Gamma_{d_n} = \{\alpha' : \text{packet} \mid \alpha' \in \text{dom}(H)\} \quad (4)$$

If **(3)** holds, then it must be the result of an application of rule (*proc*), in which case the following must hold:

$$\Gamma = \Gamma_1; \Gamma_2 \quad (5)$$

$$\Gamma_1 \vdash^e \text{let } () = \text{update}(\alpha, i, v) \text{ in } e : \vec{\tau} \quad (6)$$

$$\Gamma_2 \vdash^\Sigma \vec{\tau} \triangleright \Sigma \quad (7)$$

If **(6)** holds then it must be the result of an application of rule (*apseq*), in which case the following must hold:

$$\mathcal{F}(\text{update}) = \text{packet} \times \text{int} \times \text{int} \rightarrow () \quad (8)$$

$$\Gamma_1 \vdash^e e : \vec{\tau} \quad (9)$$

Given **(9)**, **(7)** and **(5)**, we can apply rule (*proc*), giving:

$$\Gamma \vdash^m \langle e, \Sigma \rangle \quad (10)$$

Given **(1)**, **(2)**, **(10)**, and **(4)**, we can apply rule (*ok*), giving:

$$\langle H[\alpha \mapsto H(\alpha)[i \mapsto v]] \rangle \mid \langle e, \Sigma \rangle \quad (11)$$

as required.

A.3.8 Rule *deq*

$$\vec{\Delta} \mid \langle \langle \alpha \bullet Q \rangle \rangle_q \mid \langle \text{let } x = \text{q.deq}() \text{ in } e, \Sigma \rangle \rightarrow \vec{\Delta} \mid \langle \langle Q \rangle \rangle_q \mid \langle e\{\alpha/x\}, \Sigma \rangle$$

Referring to the property we wish to prove, we see that:

$$W = \vec{\Delta} \mid \langle \langle \alpha \bullet Q \rangle \rangle_q \mid \langle \text{let } x = \text{q.deq}() \text{ in } e, \Sigma \rangle$$

$$W' = \vec{\Delta} \mid \langle \langle Q \rangle \rangle_q \mid \langle e\{\alpha/x\}, \Sigma \rangle$$

If *WellFormed*(*W*) holds, then it must be the result of an application of rule (*ok*), in which case the following must also hold:

$$\vec{\Delta} = \langle H \rangle \mid M_1 \mid \dots \mid M_n \quad (1)$$

$$\Gamma_{d_1} \vdash^m M_1 \quad \dots \quad \Gamma_{d_n} \vdash^m M_n \quad (2)$$

$$\Gamma \vdash^m \langle \text{let } x = \text{q.deq}() \text{ in } e, \Sigma \rangle \quad (3)$$

$$\Gamma_q \vdash^m \langle \langle \alpha \bullet Q \rangle \rangle_q \quad (4)$$

$$\Gamma + \Gamma_q + \Gamma_{d_1} + \dots + \Gamma_{d_n} = \{\alpha' : \text{packet} \mid \alpha' \in \text{dom}(H)\} \quad (5)$$

If **(3)** holds, then it must be the result of an application of rule (*proc*), in which case the following must hold:

$$\Gamma = \Gamma_1; \Gamma_2 \quad (6)$$

$$\Gamma_1 \vdash^e \text{let } x = \text{q.deq}() \text{ in } e : \vec{\tau} \quad (7)$$

$$\Gamma_2 \vdash^\Sigma \vec{\tau} \triangleright \Sigma \quad (8)$$

If (7) holds then it must be the result of an application of rule (*apseq*), in which case the following must hold:

$$\mathcal{F}(q.deq) = () \rightarrow packet \quad (9)$$

$$\Gamma_1, x : packet \vdash^e e : \vec{\tau} \quad (10)$$

From (4), (5) and (6) we can deduce that:

$$\alpha \notin \text{dom}(\Gamma_1) \quad (11)$$

By (10), (11) and Lemma (A.2.1) we can deduce that:

$$\Gamma_1 + \{\alpha : packet\} \vdash^e e\{\alpha/x\} : \vec{\tau} \quad (12)$$

Given (12) and (8) we can apply rule (*proc*), giving:

$$\Gamma_1 + \{\alpha : packet\}; \Gamma_2 \vdash^m \langle\!\langle e\{\alpha/x\}, \Sigma \rangle\!\rangle \quad (13)$$

By Lemma (A.2.3) and (6), this can be rewritten as:

$$\{\alpha : packet\} + \Gamma \vdash^m \langle\!\langle e\{\alpha/x\}, \Sigma \rangle\!\rangle \quad (14)$$

If (4) holds, then it must be the result of an application of rule (*queue*). We can thus deduce that:

$$\Gamma_q = \{\alpha : packet\} + \Gamma'_q \quad (15)$$

$$\Gamma'_q \vdash^m \langle\!\langle Q \rangle\!\rangle_q \quad (16)$$

From (5), (15) and the associativity of +, we can deduce that:

$$\{\alpha : packet\} + \Gamma + \Gamma'_q + \Gamma_{d_1} + \dots + \Gamma_{d_n} = \{\alpha' : packet \mid \alpha' \in \text{dom}(H)\} \quad (17)$$

Given (1), (2), (14), (16), and (17), we can apply rule (*ok*), giving:

$$\text{WellFormed}(\vec{\Delta} \mid \langle\!\langle Q \rangle\!\rangle_q \mid \langle\!\langle e\{\alpha/x\}, \Sigma \rangle\!\rangle) \quad (18)$$

as required.

A.3.9 Rule *enq*

$$\vec{\Delta} \mid \langle\!\langle Q \rangle\!\rangle_q \mid \langle\!\langle \text{let } () = \text{q.enq}(\alpha) \text{ in } e, \Sigma \rangle\!\rangle \rightarrow \vec{\Delta} \mid \langle\!\langle Q \bullet \alpha \rangle\!\rangle_q \mid \langle\!\langle e, \Sigma \rangle\!\rangle$$

Referring to the property we wish to prove, we see that:

$$\begin{aligned} W &= \vec{\Delta} \mid \langle\!\langle Q \rangle\!\rangle_q \mid \langle\!\langle \text{let } () = \text{q.enq}(\alpha) \text{ in } e, \Sigma \rangle\!\rangle \\ W' &= \vec{\Delta} \mid \langle\!\langle Q \bullet \alpha \rangle\!\rangle_q \mid \langle\!\langle e, \Sigma \rangle\!\rangle \end{aligned}$$

If $WellFormed(W)$ holds, then it must be the result of an application of rule (ok) , in which case the following must also hold:

$$\vec{\Delta} = \langle H \rangle \mid M_1 \mid \dots \mid M_n \quad (1)$$

$$\Gamma_{d_1} \vdash^m M_1 \dots \Gamma_{d_n} \vdash^m M_n \quad (2)$$

$$\Gamma \vdash^m (\text{let } () = \text{q.enq}(\alpha) \text{ in } e, \Sigma) \quad (3)$$

$$\Gamma_q \vdash^m \langle\langle Q \rangle\rangle_q \quad (4)$$

$$\Gamma + \Gamma_q + \Gamma_{d_1} + \dots + \Gamma_{d_n} = \{\alpha' : \text{packet} \mid \alpha' \in \text{dom}(H)\} \quad (5)$$

If **(3)** holds, then it must be the result of an application of rule $(proc)$, in which case the following must hold:

$$\Gamma = \Gamma_1; \Gamma_2 \quad (6)$$

$$\Gamma_1 \vdash^e \text{let } () = \text{q.enq}(\alpha) \text{ in } e : \vec{\tau} \quad (7)$$

$$\Gamma_2 \vdash^\Sigma \vec{\tau} \triangleright \Sigma \quad (8)$$

If **(7)** holds then it must be the result of an application of rule $(apseq)$, in which case the following must hold:

$$\Gamma_1 = \{\alpha : \text{packet}\}; \Gamma'_1 \quad (9)$$

$$\mathcal{F}(q.enq) = \text{packet} \rightarrow () \quad (10)$$

$$\{\alpha : \text{packet}\} \vdash^v \alpha : \text{packet} \quad (11)$$

$$\Gamma'_1 \vdash^e e : \vec{\tau} \quad (12)$$

Given **(12)** and **(8)** we can apply rule $(proc)$, giving:

$$\Gamma'_1; \Gamma_2 \vdash^m (e, \Sigma) \quad (13)$$

Given **(6)**, **(9)**, Lemma **(A.2.2)** and the lack of $!packet$ types in $\{\alpha : \text{packet}\}$ we can deduce that:

$$\Gamma = \{\alpha : \text{packet}\} + (\Gamma'_1; \Gamma_2) \quad (14)$$

Given **(4)** and the definition of rule $(queue)$ we can deduce that:

$$\{\alpha : \text{packet}\} + \Gamma_q \vdash^m \langle\langle q \bullet Q \rangle\rangle_q \quad (15)$$

$$(16)$$

Given **(5)**, **(14)**, **(15)** and the associativity of $+$ we can deduce that:

$$(\Gamma'_1; \Gamma_2) + (\{\alpha : \text{packet}\} + \Gamma_q) + \Gamma_{d_1} + \dots + \Gamma_{d_n} = \{\alpha' : \text{packet} \mid \alpha' \in \text{dom}(H)\} \quad (17)$$

Using **(1)**, **(2)**, **(13)** and **(17)**, we can apply rule (ok) , giving:

$$WellFormed(\vec{\Delta} \mid \langle\langle Q \bullet \alpha \rangle\rangle_q \mid (e, \Sigma)) \quad (18)$$

This completes the proof of the preservation of well-formedness.

References

- [1] ALLEN, J *et al.* PowerNP network processor: Hardware, software and applications. *IBM Journal of research and development* 47, 2–3 (3003), 177–194.
- [2] BERRY, G., AND BOUDOL, G. The chemical abstract machine. *Theoretical Computer Science* 96 (1992), 217–248.
- [3] BOYAPATI, C., LEE, R., AND RINARD, M. Ownership types for safe programming: Preventing data races and deadlocks. In *Proceedings of the 17th Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA-02)* (Seattle, WA, November 2002), ACM Press.
- [4] CLARKE, D., AND WRIGSTAD, T. External uniqueness is unique enough. In *Proceedings of the 17th European Conference for Object-Oriented Programming (ECOOP)* (July 2003), vol. 2743 of *LNCS*, Springer-Verlag.
- [5] CLARKE, D. G., POTTER, J. M., AND NOBLE, J. Ownership types for flexible alias protection. In *Proceedings of the 13th Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA-98)* (New York, NY, October 1998), vol. 33:10 of *ACM SIGPLAN Notices*, ACM Press, pp. 48–64.
- [6] FLANAGAN, C., SABRY, A., DUBA, B. F., AND FELLEISEN, M. The essence of compiling with continuations. In *Proceedings ACM SIGPLAN 1993 Conf. on Programming Language Design and Implementation, PLDI'93, Albuquerque, NM, USA, 23–25 June 1993*, vol. 28(6). ACM Press, New York, NY, 1993, pp. 237–247.
- [7] GROSSMAN, D., MORRISETT, G., JIM, T., HICKS, M., WANG, Y., AND CHENEY, J. Region-based memory management in Cyclone. In *Proceedings of the 2002 Conference on Programming Language Design and Implementation (PLDI)*, ACM press.
- [8] HILT HIGH LEVEL SOFTWARE TOOLS B.V. Concurrent clean language report. Available from: <ftp://ftp.cs.kun.nl/pub/Clean/Clean20/doc/CleanRep2.0.pdf>.
- [9] INTEL CORPORATION. Intel IXDP2400 & IXDP2800 advanced development platforms product brief. Available from: <http://www.intel.com/design/network/prodbrf/ixdp2x00.htm>.
- [10] INTEL CORPORATION. Intel IXP2400 Network Processor: Flexible, high-performance solution for access and edge applications. Available from: <http://www.intel.com/design/network/papers/ixp2400.htm>.
- [11] JOHNSON, E. J., AND KUNZE, A. *IXP-1200 Programming: The Microengine Coding Guide for the Intel IXP1200 Network Processor Family*. Intel Press, 2002.
- [12] KOBAYASHI, N. Quasi-linear types. In *Proceedings of The 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 99)* (San Antonio, TX, 1999), pp. 29–42.

- [13] POSTEL, J. Internet Protocol RFC 791. Internet Engineering Task Force, September 1981. Available from: <http://www.ietf.org/rfc.html>.
- [14] POSTEL, J. Internet Control Message Protocol, RFC 792. Internet Engineering Task Force, September 1981. Available from: <http://www.ietf.org/rfc.html>.
- [15] STROUSTRUP, B. *The C++ programming language (third edition)*. Addison-Wesley.
- [16] TOFTE, M., AND TALPIN, J.-P. Region-based memory management. *Information and Computation* (1997).
- [17] TURNER, D. N., WADLER, P., AND MOSSIN, C. Once upon a type. In *Functional Programming Languages and Computer Architecture* (1995), ACM Press.
- [18] WALKER, D., AND WATKINS, K. On regions and linear types. In *International Conference on Functional Programming* (2001), ACM press, pp. 181–192.
- [19] YAVATKAR, R., AND H. VIN (EDS.). *IEEE Network Magazine. Special issue on Network Processors: Architecture, Tools, and Applications 17, 4* (July 2003).