

Broadcast psi-calculi with an application to wireless protocols

Johannes Borgström · Shuqin Huang ·
Magnus Johansson · Palle Raabjerg · Björn Victor ·
Johannes Åman Pohjola · Joachim Parrow

Received: 16 March 2012 / Revised: 21 July 2013 / Accepted: 15 August 2013 / Published online: 8 November 2013
© Springer-Verlag Berlin Heidelberg 2013

Abstract Psi-calculi is a parametric framework for the extensions of pi-calculus, with arbitrary data structures and logical assertions for facts about data. In this paper we add primitives for broadcast communication in order to model wireless protocols. The additions preserve the purity of the psi-calculi semantics, and we formally prove the standard congruence and structural properties of bisimilarity. We demonstrate the expressive power of broadcast psi-calculi by modelling the wireless ad hoc routing protocol LUNAR and verifying a basic reachability property.

Keywords Psi-calculus · Broadcast communication · Bisimulation · Ad hoc routing protocol

1 Introduction

Psi-calculi is a parametric framework for the extensions of pi-calculus, with arbitrary data structures and logical assertions for facts about data. In earlier papers we have shown how psi-calculi can capture the same phenomena as other proposed extensions of the pi-calculus such as the applied pi-calculus, the spi-calculus, the fusion calculus, the concurrent constraint pi-calculus, and calculi with polyadic communication channels or pattern matching. Psi-calculi

can be even even more general, for example by allowing structured channels, higher-order formalisms such as the lambda calculus for data structures, and predicate logic for assertions [5].

In psi-calculi (described in Sect. 2), the purity of the semantics is on par with the original pi-calculus, the generality and expressiveness exceeds many earlier extensions of the pi-calculus, and the meta-theory is proved correct once and for all using the interactive theorem prover Isabelle/Nominal [34]. The communication paradigm in psi-calculi is binary: for each event, there is one sender and one receiver, just as in the pi-calculus. In several areas, e.g. wireless communications and hardware data buses, a natural paradigm is broadcast, where one transmission can be received by several processes. Broadcast communication cannot be uniformly encoded in the pi-calculus [7].

In this paper we extend the psi-calculi framework with primitives for synchronous unreliable broadcast. These require new operational actions and rules, and new connectivity predicates. In Sect. 3.1, we formally prove the congruence properties of bisimilarity and the soundness of structural equivalence laws using the Isabelle/Nominal theorem prover.

The connectivity predicates allow us to model systems with limited reachability, for instance where a transmitter only reaches nodes within a certain range, and systems with changing reachability, for instance due to physical mobility of nodes. In Sect. 4, we present a technique for treating different generations of connectivity information. Broadcast channels can be globally visible or have limited scope. Scoped channels can be protected from externally imposed connectivity changes, while permitting connectivity changes by processes within the scope of the channel. One of our main contributions is precise requirements that the connectivity predicates must satisfy, in order to model scoped broadcasts

Communicated by Dr. Gerardo Schneider, Gilles Barthe, and Alberto Pardo.

J. Borgström · M. Johansson · P. Raabjerg · B. Victor (✉) ·
J. Åman Pohjola · J. Parrow
Department of Information Technology, Uppsala University,
Uppsala, Sweden
e-mail: bjorn.victor@it.uu.se

S. Huang
Peking University, Beijing, China

with dynamic connectivity, while still satisfying the meta-theoretical results of Sect. 3.1.

We demonstrate the expressive power of the resulting framework in Sect. 5, where we provide a model of the LUNAR protocol for routing in ad hoc wireless networks [33]. The model follows the specification closely and demonstrates several features of the psi-calculi framework: both unicast and broadcast communication, application-specific data structures and logics, classic unstructured channels as well as pairs corresponding to MAC address and port selector. Our model is significantly more succinct than earlier work [35,36] (ca 30 vs. 250 lines). We show an expected basic reachability property of the model: if two network nodes, a sender and a receiver, are both in range of a third node, but not within range of each other, the LUNAR protocol can find a route and transparently handle the delivery of a packet from the sender to the receiver.

We discuss related work on process calculi for wireless broadcast in Sect. 6 and conclude and present ideas for future work in Sect. 7.

This paper is an extended version of [6] that adds clarifications, proofs, and elaborated examples of dynamic topology management.

2 Psi-calculi

This section is a brief recapitulation of psi-calculi; for an extensive treatment including more motivations and examples, see [4,5], from which some examples and explanations below are taken.

We assume a countably infinite set of atomic names \mathcal{N} ranged over by a, b, \dots, z . Intuitively, names will represent the symbols that can be scoped and also represent symbols acting as variables in the sense that they can be subject to substitution. As a general framework for terms and other data containing names, we work in the formalism of *nominal sets* [8,24]. A nominal set is an ordinary set equipped with a formal notion of what it means for a name a to occur in an element X of the set, written $a \in n(X)$ (often pronounced as “ a is in the support of X ”). We write $a\#X$, pronounced “ a is fresh for X ”, for $a \notin n(X)$, and if A is a finite set of names, we write $A\#X$ to mean $\forall a \in A. a\#X$. We require all elements to have finite support, i.e. $n(X)$ is finite for all X . In the following, \tilde{a} means a finite sequence of names, a_1, \dots, a_n . The empty sequence is written ϵ , and the concatenation of \tilde{a} and \tilde{b} is written $\tilde{a}\tilde{b}$. When occurring as an operand of a set operator, \tilde{a} means the corresponding set of names $\{a_1, \dots, a_n\}$. We also use sequences of other nominal sets in the same way. For names, we write $(\tilde{a} \tilde{b})$ for the *name swapping* that swaps each element of \tilde{a} with the corresponding element of \tilde{b} ; here it is implicit that \tilde{a} and \tilde{b} have the same length and that the names in \tilde{a} (respectively, \tilde{b}) are pairwise distinct. A

function f is *equivariant* if $(a b) \cdot f(X) = f((a b) \cdot X)$ holds for all X , and similarly for functions and relations of any arity. Intuitively, equivariance means that all names are treated equally.

A *nominal data type* is a nominal set together with a set of functions on it. In particular, we shall consider substitution functions that substitute elements for names. If X is an element of a data type, \tilde{a} is a sequence of names without duplicates and \tilde{Y} is an equally long sequence of elements of possibly another data type, the *substitution* $X[\tilde{a} := \tilde{Y}]$ is an element of the same data type as X . Substitution is required to satisfy a law akin to alpha-conversion: if $\tilde{b}\#X, \tilde{a}$ then $X[\tilde{a} := \tilde{T}] = ((\tilde{b} \tilde{a}) \cdot X)[\tilde{b} := \tilde{T}]$. Intuitively, this ensures that substitutions for bound names yield the same result no matter which alpha-equivalent version is used.

We use nominal data types in order to obtain a general framework, allowing many different instantiations. Our only requirements are on the notions of support, name swapping, and substitution. Thus, we can handle data types that are not inductively defined, such as equivalence classes or sets defined by comprehension or co-induction. Examples include higher-order data types such as the lambda calculus. As long as the term language satisfies the axioms of a nominal data type, it can be used in our framework. Similarly, the notions of conditions, i.e. the tests on data that agents can perform during their execution, and assertions, i.e. the facts that can be used to resolve conditions, are formulated as nominal data types. This means that logics with binders and even higher-order logics can be used. Moreover, alpha-variants of terms can be formally equated by taking the quotient of terms under alpha equality, thereby facilitating the formalism and proofs.

A psi-calculus is defined by instantiating three nominal data types and four operators:

Definition 1 (*Psi-calculus parameters*) A psi-calculus requires the three (not necessarily disjoint) nominal data types: the (data) terms \mathbf{T} , ranged over by M, N , the conditions \mathbf{C} , ranged over by φ , the assertions \mathbf{A} , ranged over by Ψ , and the four equivariant operators:

- $\leftrightarrow : \mathbf{T} \times \mathbf{T} \rightarrow \mathbf{C}$ Channel Equivalence
- $\otimes : \mathbf{A} \times \mathbf{A} \rightarrow \mathbf{A}$ Composition
- $\mathbf{1} : \mathbf{A}$ Unit
- $\vdash \subseteq \mathbf{A} \times \mathbf{C}$ Entailment

and substitution functions $[\tilde{a} := \tilde{M}]$, substituting terms for names, on each of \mathbf{T} , \mathbf{C} , and \mathbf{A} , where the substitution function on \mathbf{T} , in addition to the alpha-conversion-like law above, satisfies the following name preservation law: if $\tilde{a} \subseteq n(M)$ and $b \in n(\tilde{N})$, then $b \in n(M[\tilde{a} := \tilde{N}])$.

The binary functions above will be written in infix. Thus, if M and N are terms, then $M \leftrightarrow N$ is a condition, pronounced “ M and N are channel equivalent”, and if Ψ and Ψ' are

assertions, then so is $\Psi \otimes \Psi'$. Also, we write $\Psi \vdash \varphi$, “ Ψ entails φ ”, for $(\Psi, \varphi) \in \vdash$.

As an example, we can choose data terms inductively generated by some signature, assertions, and conditions to be elements of a first-order logic with equality over these terms, entailment to be logical implication, \otimes to be conjunction and $\mathbf{1}$ to be TRUE. We call this example instance **euf**.

We say that two assertions are equivalent, written $\Psi \simeq \Psi'$, if they entail the same conditions, i.e. for all φ we have that $\Psi \vdash \varphi \Leftrightarrow \Psi' \vdash \varphi$. We impose certain straightforward requisites on the sets and operators. In brief, channel equivalence must be symmetric and transitive (but not necessarily reflexive), \otimes must be compositional with regard to \simeq , and the assertions with $(\otimes, \mathbf{1})$ form an abelian monoid modulo \simeq . In the **euf** instance, we can let channel equivalence be term equality: symmetry and reflexivity clearly hold, logical conjunction does form an abelian monoid with TRUE as unit, and compositionality of assertion composition follows from the tautology $(\Psi \Rightarrow (\Psi_1 \Leftrightarrow \Psi_2)) \Rightarrow (\Psi \Rightarrow (\Psi' \wedge \Psi_1 \Leftrightarrow \Psi' \wedge \Psi_2))$. For details see [5].

A *frame* F can intuitively be thought of as an assertion with local names: it is of the form $(v\tilde{b})\Psi$ where \tilde{b} is a sequence of names that bind into the assertion Ψ . We use F, G to range over frames. We overload Ψ to also mean the frame $(v\epsilon)\Psi$ and \otimes to composition on frames defined by $(v\tilde{b}_1)\Psi_1 \otimes (v\tilde{b}_2)\Psi_2 = (v\tilde{b}_1\tilde{b}_2)(\Psi_1 \otimes \Psi_2)$ where $\tilde{b}_1\#\tilde{b}_2, \Psi_2$ and vice versa. We write $\Psi \otimes F$ to mean $(v\epsilon)\Psi \otimes F$, and $(v\tilde{c})((v\tilde{b})\Psi)$ for $(v\tilde{c}\tilde{b})\Psi$.

Alpha-equivalent frames are identified. We define $F \vdash \varphi$ to mean that there exists an alpha variant $(v\tilde{b})\Psi$ of F such that $\tilde{b}\#\varphi$ and $\Psi \vdash \varphi$. We also define $F \simeq G$ to mean that for all φ it holds that $F \vdash \varphi$ iff $G \vdash \varphi$. Intuitively, a condition is entailed by a frame if it is entailed by the assertion and does not contain any names bound by the frame, and two frames are equivalent if they entail the same conditions.

In the **euf** example, assume that the term $\text{enc}(M, k)$ represents the encoding of message M with key k , and let Ψ be the assertion $C = \text{enc}(M, k)$, stating that the ciphertext C is the result of encoding M by k . If an agent contains this assertion, the environment of the agent will be able to use it to resolve tests on the data. In particular, it may infer that $C = \text{enc}(M, k)$, i.e. it can test whether this C is the encryption of M . Access to the key k can be restricted by enclosing it in a scope: if the environment instead has access to the assertion $(vk)\Psi$, it can *not* infer that C is the encoding of M (assuming conditions only contain equality tests on terms, and no quantifiers). For more discussion, see [5].

Definition 2 (*Psi-calculus agents*) Given valid psi-calculus parameters as in Definition 1, the psi-calculus *agents*, ranged over by P, Q, \dots , are of the following forms.

$\mathbf{0}$	Nil
$\overline{MN}.P$	Output
$\underline{M}(\lambda\tilde{x})N.P$	Input
case $\varphi_1 : P_1 [] \dots [] \varphi_n : P_n$	Case
$(\nu a)P$	Restriction
$P \mid Q$	Parallel
$!P$	Replication
$(\downarrow\Psi)$	Assertion

Restriction binds a in P , and Input binds \tilde{x} in both N and P . We identify alpha-equivalent agents. An assertion is *guarded* if it is a subterm of an Input or Output. An agent is *assertion guarded* if it contains no unguarded assertions. An agent is *well formed* if in $\underline{M}(\lambda\tilde{x})N.P$ it holds that $\tilde{x} \subseteq n(N)$ is a sequence without duplicates, that in a replication $!P$ the agent P is assertion guarded, and that in **case** $\varphi_1 : P_1 [] \dots [] \varphi_n : P_n$ the agents P_i are assertion guarded.

In the Output and Input forms, M is called the subject and N the object. Output and Input are similar to those in the pi-calculus, but arbitrary terms can function as both subjects and objects. In the input $\underline{M}(\lambda\tilde{x})N.P$, the intuition is that the pattern $(\lambda\tilde{x})N$ can match any term obtained by instantiating \tilde{x} , e.g. $\underline{M}(\lambda x, y)f(x, y).P$ can only communicate with an output $\overline{M}f(N_1, N_2)$ for some data terms N_1, N_2 . This can be thought of as a generalisation of the polyadic pi-calculus where the patterns are just tuples of names. Another significant extension is that we allow arbitrary data terms also as communication channels. Thus, it is possible to include functions that create channels.

The **case** construct behaves as one of the P_i for which the corresponding φ_i is true. The agent **case** $\varphi_1 : P_1 [] \dots [] \varphi_n : P_n$ is sometimes abbreviated as **case** $\tilde{\varphi} : \tilde{P}$, or if $n = 1$ as **if** φ_1 **then** P_1 . Input subjects are underlined to facilitate parsing of complicated expressions; in simple cases we often omit the underline. We sometimes write $\underline{M}(x).P$ for $\underline{M}(\lambda x)x.P$.

One of the simplest examples of a psi-calculus is the pi-calculus [21] that can be represented using names as the only data terms, $\mathbf{1}$ as the only assertion, and equality tests on names as conditions. Channel equivalence \Leftrightarrow is also equality on names. Substitution is the standard syntactic replacement of names for names. Choice in the pi-calculus can be represented using the **case** statement: $P + Q$ corresponds to $(\nu a)(\mathbf{case} a = a : P [] a = a : Q)$, where $a\#P, Q$, and the pi-calculus match construct $[a = b]P$ corresponds to **if** $a = b$ **then** P . The formal correspondence between this psi-calculus instance and the original pi-calculus is proved in [5].

As indicated in the encryption example above, the conditions tested in a process are affected by the assertions of parallel processes. For example, in $P \mid Q$, the assertions of P can affect the conditions tested in Q , and thereby its transi-

tions. We introduce the *frame of an agent* as the combination of its top-level assertions, retaining all the binders: this is precisely what can affect a parallel agent. The *frame* $\mathcal{F}(P)$ of an agent P is defined inductively as follows:

$$\begin{aligned} \mathcal{F}(\underline{M}(\lambda\tilde{x})N . P) &= \mathcal{F}(\overline{M}N . P) \\ &= \mathcal{F}(\mathbf{0}) = \mathcal{F}(\mathbf{case} \tilde{\varphi} : \tilde{P}) = \mathcal{F}(!P) = \mathbf{1} \\ \mathcal{F}(\langle\!\langle\Psi\rangle\!\rangle) &= (\nu\epsilon)\Psi \\ \mathcal{F}(P \mid Q) &= \mathcal{F}(P) \otimes \mathcal{F}(Q) \\ \mathcal{F}((\nu b)P) &= (\nu b)\mathcal{F}(P) \end{aligned}$$

For a simple example, if $a\#\Psi_1$:

$$\mathcal{F}(\langle\!\langle\Psi_1\rangle\!\rangle \mid (\nu a)(\langle\!\langle\Psi_2\rangle\!\rangle \mid \overline{M}N . \langle\!\langle\Psi_3\rangle\!\rangle)) = (\nu a)(\Psi_1 \otimes \Psi_2)$$

Here, Ψ_3 occurs under a prefix and is therefore not included in the frame.

The *actions* ranged over by α, β are of the following three kinds:

Output $\overline{M}(\nu\tilde{a})N$ where $\alpha \subseteq n(N)$, Input $\underline{M}N$, and Silent τ . Here, we refer to M as the *subject* and N as the *object*. We define $\text{bn}(\overline{M}(\nu\tilde{a})N) = \tilde{a}$, and $\text{bn}(\alpha) = \emptyset$ if α is an input or τ . We also define $n(\tau) = \emptyset$ and $n(\alpha) = n(M) \cup n(N)$ for the input and output actions. As in the pi-calculus, the output $\overline{M}(\nu\tilde{a})N$ represents an action sending N along M and opening the scopes of the names \tilde{a} . Note in particular that the support of this action includes \tilde{a} . Thus, $\overline{M}(\nu a)a$ and $\overline{M}(\nu b)b$ are different actions.

Definition 3 (Transitions) A *transition* is written $\Psi \triangleright P \xrightarrow{\alpha} P'$, meaning that in the environment Ψ the well-formed agent P can do an α to become P' . The transitions are defined

inductively in Table 1. We write $P \xrightarrow{\alpha} P'$ without an assertion to mean $\mathbf{1} \triangleright P \xrightarrow{\alpha} P'$.

Agents, frames, and transitions are identified by alpha equivalence. In a transition the names in $\text{bn}(\alpha)$ bind into both the action object and the derivative; therefore, $\text{bn}(\alpha)$ is in the support of α but not in the support of the transition. This means that the bound names can be chosen fresh, substituting each occurrence in both the object and the derivative.

The environmental assertions $\Psi \triangleright \dots$ in Table 1 express the effect that the environment has on the agent: enabling conditions in CASE, giving rise to action subjects in IN and OUT, and enabling interactions in COM. The environment Ψ increases towards the leaves of the derivation only in the rules for the parallel operator, where an agent is part of the environment for another agent. If all environmental assertions are erased and channel equivalence is replaced by identity, we get the standard laws of the pi-calculus enriched with data structures.

For a simple example of a transition, suppose for an assertion Ψ and condition φ that $\Psi \vdash \varphi$. Assume that

$$\forall \Psi'. \Psi' \triangleright Q \xrightarrow{\alpha} Q'$$

i.e. Q has an action α regardless of the environment. Then by the CASE rule we get

$$\Psi \triangleright \mathbf{if} \varphi \mathbf{then} Q \xrightarrow{\alpha} Q'$$

i.e. **if then** Q has the same transition if the environment is Ψ . Since $\mathcal{F}(\langle\!\langle\Psi\rangle\!\rangle) = \Psi$ and $\Psi \otimes \mathbf{1} = \Psi$, if $\text{bn}(\alpha)\#\Psi$ we get by PAR that

$$\mathbf{1} \triangleright \langle\!\langle\Psi\rangle\!\rangle \mid \mathbf{if} \varphi \mathbf{then} Q \xrightarrow{\alpha} \langle\!\langle\Psi\rangle\!\rangle \mid Q'$$

Table 1 Structured operational semantics

	$\text{IN} \frac{\Psi \vdash K \leftrightarrow M}{\Psi \triangleright \underline{M}(\lambda\tilde{y})N . P \xrightarrow{\underline{K}N[\tilde{y}:=\tilde{L}]} P[\tilde{y}:=\tilde{L}]}$	$\text{OUT} \frac{\Psi \vdash M \leftrightarrow K}{\Psi \triangleright \overline{M}N . P \xrightarrow{\overline{K}N} P}$
	$\text{CASE} \frac{\Psi \triangleright P_i \xrightarrow{\alpha} P' \quad \Psi \vdash \varphi_i}{\Psi \triangleright \mathbf{case} \tilde{\varphi} : \tilde{P} \xrightarrow{\alpha} P'}$	
	$\text{COM} \frac{\Psi \otimes \Psi_P \otimes \Psi_Q \vdash M \leftrightarrow K \quad \Psi_Q \otimes \Psi \triangleright P \xrightarrow{\overline{M}(\nu\tilde{a})N} P' \quad \Psi_P \otimes \Psi \triangleright Q \xrightarrow{\underline{K}N} Q' \quad \tilde{a}\#Q}{\Psi \triangleright P \mid Q \xrightarrow{\tau} (\nu\tilde{a})(P' \mid Q')}$	
	$\text{PAR} \frac{\Psi_Q \otimes \Psi \triangleright P \xrightarrow{\alpha} P' \quad \text{bn}(\alpha)\#Q}{\Psi \triangleright P \mid Q \xrightarrow{\alpha} P' \mid Q}$	
	$\text{SCOPE} \frac{\Psi \triangleright P \xrightarrow{\alpha} P' \quad b\#\alpha, \Psi}{\Psi \triangleright (\nu b)P \xrightarrow{\alpha} (\nu b)P'}$	
	$\text{OPEN} \frac{\Psi \triangleright P \xrightarrow{\overline{M}(\nu\tilde{a})N} P' \quad b\#\tilde{a}, \Psi, M \quad b \in n(N)}{\Psi \triangleright (\nu b)P \xrightarrow{\overline{M}(\nu\tilde{a} \cup \{b\})N} P'}$	
	$\text{REP} \frac{\Psi \triangleright P \mid !P \xrightarrow{\alpha} P'}{\Psi \triangleright !P \xrightarrow{\alpha} P'}$	

Symmetric versions of COM and PAR are elided in the rule COM we assume that $\mathcal{F}(P) = (\nu\tilde{b}_P)\Psi_P$ and $\mathcal{F}(Q) = (\nu\tilde{b}_Q)\Psi_Q$ where \tilde{b}_P is fresh for all of Ψ, \tilde{b}_Q, Q, M , and P and that \tilde{b}_Q is similarly fresh. In the rule PAR we assume that $\mathcal{F}(Q) = (\nu\tilde{b}_Q)\Psi_Q$ where \tilde{b}_Q is fresh for Ψ, P , and α . In OPEN the expression $\tilde{a} \cup \{b\}$ means the sequence \tilde{a} with b inserted anywhere

The notion of strong bisimulation is used to formalise the intuition that two agents “behave in the same way”.

Definition 4 (*Strong bisimulation*) A strong bisimulation \mathcal{R} is a ternary relation on assertions and pairs of agents such that $\mathcal{R}(\Psi, P, Q)$ implies

1. Static equivalence: $\Psi \otimes \mathcal{F}(P) \simeq \Psi \otimes \mathcal{F}(Q)$; and
2. Symmetry: $\mathcal{R}(\Psi, Q, P)$; and
3. Extension of arbitrary assertion: $\forall \Psi'. \mathcal{R}(\Psi \otimes \Psi', P, Q)$; and
4. Simulation: for all α, P' such that $\Psi \triangleright P \xrightarrow{\alpha} P'$ and $\text{bn}(\alpha)\#\Psi, Q$, there exists Q' such that $\Psi \triangleright Q \xrightarrow{\alpha} Q'$ and $\mathcal{R}(\Psi, P', Q')$.

We define $P \sim_{\Psi} Q$ to mean that there exists a bisimulation \mathcal{R} such that $\mathcal{R}(\Psi, P, Q)$ and write \sim for \sim_1 .

Strong bisimulation is a congruence in the usual sense: it is preserved by all operators except input prefix and satisfies the expected algebraic laws such as scope extension $P \mid (\nu a)Q \sim (\nu a)(P \mid Q)$ if $a\#P$. For details see [4,5]. Note that these meta-theoretical results have been proven to hold for all psi-calculus instances using the interactive theorem prover Isabelle/Nominal [34].

Psi-calculi can capture the same phenomena as a wide range of previously proposed individual extensions of the pi-calculus. Examples in [4,5] range from foundational calculi such as polyadic pi-calculus, polyadic synchronisation pi-calculus, fusion calculus, and concurrent constraint calculi, to applied calculi for cryptography and systems with frequency hopping communication protocols. Each previous pi-calculus extension in the literature has needed new proofs of basic results such as scope extension and bisimulation congruence. Instead, formulated as psi-calculus instances, all the meta-theory of psi-calculi is automatically inherited.

3 Broadcast psi-calculi

In this section we extend the unicast psi-calculi of the previous section with a communication paradigm for synchronous unreliable non-blocking broadcast (suitable for modelling wireless communication). We introduce the notion of a *broadcast channel* as an abstraction of relevant properties of the transmission, such as frequency, sender location, and signal strength. Formally, a broadcast channel is just a term. We assume the so-called *connectivity predicates* that regulate which prefix subjects can send on or receive from which broadcast channels. These predicates may depend on assertions and therefore change as an agent evolves.

As an example, assume that the connectivity information Ψ allows the sender M_0 to send on the broadcast channel K , and receivers M_1 and M_2 to listen on K . We would then have

the following transition:

$$\Psi \triangleright \overline{M_0} N.P \mid \underline{M_1}(x).Q \mid \underline{M_2}(y).R \xrightarrow{!K N} P \mid Q[x := N] \mid R[y := N].$$

Here, in one action two processes both receive the N sent along K , and moreover, the action label retains the broadcast output action $!K N$, meaning that in a larger context even more processes could receive N .

Formally, we assume a psi-calculus with the following extra predicates:

Definition 5 (*Extra predicates for broadcast*)

$$\dot{\prec}: \mathbf{T} \times \mathbf{T} \rightarrow \mathbf{C} \quad \text{Output Connectivity}$$

$$\dot{\succ}: \mathbf{T} \times \mathbf{T} \rightarrow \mathbf{C} \quad \text{Input Connectivity}$$

The first predicate, $M \dot{\prec} K$, is pronounced “ M is out-connected to K ” and means that an output prefix $\overline{M} N$ can result in a broadcast on channel K . The second, $K \dot{\succ} M$, is pronounced “ M is in-connected to K ” and means that an input prefix $\underline{M}(\lambda\tilde{x})N$ can receive broadcast messages from channel K . As usual in broadcast calculi, the receivers need to use the same broadcast channel as the sender in order to receive a message.

As an example, we can model lookup in a routing table: if the term tab is a list of pairs of identifiers and channels, we can let $\Psi \vdash \text{lookup}(tab, id) \dot{\prec} ch$ be true iff (id, ch) appears in the routing table tab . We can also model connectivity: if Ψ contains connectivity information between channels ch and receivers n , we may let $\Psi \vdash ch \dot{\succ} rcv(n, ch)$ be true if n is connected to ch according to Ψ .

In contrast to unicast connectivity, we do not require broadcast connectedness to be symmetric or transitive, so in particular $M \dot{\prec} K$ might not be equivalent to $K \dot{\succ} M$. Instead, for technical reasons related to scope extension (cf. Example 13), broadcast channels must have no greater support than the input and output prefixes that send and receive on them.

Definition 6 (*Requirements for broadcast*)

1. $\Psi \vdash M \dot{\prec} K \implies n(M) \supseteq n(K)$
2. $\Psi \vdash K \dot{\succ} M \implies n(K) \subseteq n(M)$

Definition 7 (*Transitions of broadcast psi*) To the actions of psi-calculi we add broadcast input, written $?K N$ for a reception of N on K , and broadcast output, written $!K(\nu\tilde{a})N$ for a broadcast of N on K , with names \tilde{a} fresh in K . As before, we omit $(\nu\tilde{a})$ when \tilde{a} is empty, and in examples we omit N when it is not relevant. The transitions of well-formed agents are defined inductively in Tables 1 and 2, where we let α range over both unicast and broadcast actions.

The rule BROUT allows transmission on a broadcast channel K that the subject M of an output prefix is out-connected

Table 2 Operational broadcast semantics

$\text{BROUT} \frac{\Psi \vdash M \dot{<} K}{\Psi \triangleright \overline{M} N . P \xrightarrow{!K N} P}$ $\text{BRMERGE} \frac{\Psi_Q \otimes \Psi \triangleright P \xrightarrow{?K N} P' \quad \Psi_P \otimes \Psi \triangleright Q \xrightarrow{?K N} Q'}{\Psi \triangleright P \mid Q \xrightarrow{?K N} P' \mid Q'}$ $\text{BRCOM} \frac{\Psi_Q \otimes \Psi \triangleright P \xrightarrow{!K (v\tilde{a})N} P' \quad \Psi_P \otimes \Psi \triangleright Q \xrightarrow{?K N} Q' \quad \tilde{a} \# Q}{\Psi \triangleright P \mid Q \xrightarrow{!K (v\tilde{a})N} P' \mid Q'}$ $\text{BROPEN} \frac{\Psi \triangleright P \xrightarrow{!K (v\tilde{a})N} P' \quad \Psi \triangleright (vb) P \xrightarrow{!K (v\tilde{a} \cup (b))N} P'}{\Psi \triangleright (vb) P \xrightarrow{!K (v\tilde{a})N} P' \quad b \# \tilde{a}, \Psi, K \quad b \in n(N)}$ $\text{BRCLOSE} \frac{\Psi \triangleright P \xrightarrow{!K (v\tilde{a})N} P' \quad b \in n(K) \quad b \# \Psi}{\Psi \triangleright (vb) P \xrightarrow{\tau} (vb)(v\tilde{a})P'}$	$\text{BRIN} \frac{\Psi \vdash K \dot{>} M}{\Psi \triangleright \underline{M}(\lambda.\tilde{y})N . P \xrightarrow{?K N[\tilde{y}:=\tilde{L}]} P[\tilde{y}:=\tilde{L}]}$
--	--

A symmetric version of BRCOM is elided. In rules BRCOM and BRMERGE, we assume that $\mathcal{F}(P) = (v\tilde{b}_P)\Psi_P$ and $\mathcal{F}(Q) = (v\tilde{b}_Q)\Psi_Q$ where \tilde{b}_P is fresh for P, \tilde{b}_Q, Q, K , and Ψ and that \tilde{b}_Q is fresh for Q, \tilde{b}_P, P, K , and Ψ

to. Similarly, the rule BRIN allows input from a broadcast channel K that the subject M of an input prefix is in-connected to. The environmental assertion Ψ determines whether a prefix is connected to a broadcast channel and thus gives rise to a broadcast in BRIN and BROUT. In the same way, it determines whether a prefix is channel equivalent to something else and thus gives rise to a unicast in IN and OUT. The same prefix could theoretically be used for both kinds of communication, although it may be unusual to find situations where that would be useful.

When two parallel processes both receive a broadcast on the same channel, the rule BRMERGE combines the two actions. This rule is necessary to ensure the associativity of parallel composition. After a broadcast communication using BRCOM, the resulting action is the original transmission. This is different from the unicast COM rule, where a communication yields an internal action τ . The BROPEN rule allows broadcast communication of data containing scoped names. Rule BRCLOSE states that a broadcast transmission does not reach beyond its scope. This allows for broadcasting on restricted channels. Dually, the SCOPE rule (of Table 1) ensures that broadcast receivers on restricted channels cannot proceed unless a message is sent. The PAR rule allows for broadcasts to bypass a process, as in most other broadcast calculi for wireless systems.

3.1 Meta-theory

We have developed a meta-theory for broadcast psi-calculi. Theorems 8, 10, and 11 give us assurance that any broadcast psi-calculus has a compositional labelled bisimilarity that respects important structural laws. The proofs of these results are mostly straightforward extensions of the corresponding proofs for standard (unicast) psi-calculi [3, 14], where some technical lemmas can be simplified because of

the requirement of syntactic equality of channels in rules BRCOM and BRMERGE. Most of the added complications are caused by the fact that the BRCOM rule defers the closing of the communication to BRCLOSE; cf. Lemma 12. The proofs [28] are formally verified in the interactive theorem prover Isabelle/Nominal. The full formalisation of broadcast psi-calculi amounts to ca 33,000 lines of Isabelle code, of which about 21,000 lines are reused from our earlier work [5].

In the following, we restrict attention to well-formed agents.

Theorem 8 (Congruence properties of strong bisimulation) *For all Ψ :*

$$\begin{aligned}
 P \dot{\sim}_\Psi Q &\implies P \mid R \dot{\sim}_\Psi Q \mid R \\
 P \dot{\sim}_\Psi Q &\implies (va)P \dot{\sim}_\Psi (va)Q \quad \text{if } a \# \Psi \\
 P \dot{\sim}_\Psi Q &\implies !P \dot{\sim}_\Psi !Q \quad \text{if } P, Q \text{ assertion guarded} \\
 \forall i. P_i \dot{\sim}_\Psi Q_i &\implies \text{case } \tilde{\varphi} : \tilde{P} \dot{\sim}_\Psi \text{case } \tilde{\varphi} : \tilde{Q} \\
 P \dot{\sim}_\Psi Q &\implies \overline{M} N . P \dot{\sim}_\Psi \overline{M} N . Q \\
 (\forall \tilde{L}. P[\tilde{x} := \tilde{L}] \dot{\sim}_\Psi Q[\tilde{x} := \tilde{L}]) &\implies \underline{M}(\lambda.\tilde{x})N . P \dot{\sim}_\Psi \underline{M}(\lambda.\tilde{x})N . Q
 \end{aligned}$$

As usual in channel-passing calculi, bisimulation is not a congruence for input prefix. We can characterise strong bisimulation congruence in the usual way.

Definition 9 (Strong congruence) $P \dot{\sim}_\Psi Q$ iff for all sequences σ of substitutions it holds that $P\sigma \dot{\sim}_\Psi Q\sigma$. We write $P \sim Q$ for $P \sim_1 Q$.

Theorem 10 *Strong congruence \sim_Ψ is a congruence for all Ψ .*

The standard rules of structural equivalence are sound for bisimilarity congruence.

Theorem 11 (Structural equivalence) *Assume that $a\#Q, \tilde{x}, M, N, \tilde{\varphi}$. Then*

$$\begin{array}{l} \text{case } \tilde{\varphi} : \widetilde{(va)P} \sim (va)\text{case } \tilde{\varphi} : \tilde{P} \quad (va)\mathbf{0} \sim \mathbf{0} \\ \underline{M}(\lambda\tilde{x})N . (va)P \sim (va)\underline{M}(\lambda\tilde{x})(N) . P \quad Q \mid (va)P \sim (va)(Q \mid P) \\ \overline{M}N . (va)P \sim (va)\overline{M}N . P \quad (vb)(va)P \sim (va)(vb)P \\ P \mid (Q \mid R) \sim (P \mid Q) \mid R \quad !P \sim P \mid !P \\ P \mid Q \sim Q \mid P \quad P \sim P \mid \mathbf{0} \end{array}$$

When proving Theorem 11, we encountered an unusual complication in the proof of the commutativity of restriction, due to the BRCLOSE rule. Since this rule can insert binder sequences under name restrictions, the simulation proof needs to allow for permutations of sequences of top-level binders. This is the main difference in our meta-theoretical proofs as compared to the original psi-calculi. We write $\tilde{a} \equiv \tilde{b}$ to denote that the sequence \tilde{a} is a rearrangement of \tilde{b} , preserving the number of occurrences of each name.

Lemma 12 *For all Ψ, P, x, y , we have $(vy)(vx)P \sim_{\Psi} (vx)(vy)P$.*

Proof In standard psi-calculi, the proof of this result uses the candidate relation $\mathcal{S}_0 \stackrel{\text{def}}{=} \{(\Psi, (vy)(vx)P, (vx)(vy)P) : x, y\#\Psi\}$. Here, we inductively close this relation under restriction, yielding \mathcal{S} :

$$\mathcal{S} \stackrel{\text{def}}{=} \mathcal{S}_0 \cup \{(\Psi, (va)P, (va)Q) : (\Psi, P, Q) \in \mathcal{S} \wedge a\#\Psi\}.$$

We show that \mathcal{S} is a bisimulation up to transitivity [29] (at every Ψ). That is, we only require the derivatives after a simulation step to be related by \mathcal{S}^* , inductively defined as

$$\mathcal{S}^* \stackrel{\text{def}}{=} \{(\Psi, P, P)\} \cup \{(\Psi, P, R) : \exists Q. (\Psi, P, Q) \in \mathcal{S}^* \wedge (\Psi, Q, R) \in \mathcal{S}\}.$$

We have proven “up to transitivity” to be sound, i.e. every bisimulation up to transitivity is a subset of some ordinary bisimulation.

The interesting part of the proof is in the simulation clause. We here consider only the base case of the definition of \mathcal{S} (i.e. \mathcal{S}_0), where we need to prove that for all α, P' such that $\text{bn}(\alpha)\#\Psi, Q$ and $\Psi \triangleright (vy)(vx)P \xrightarrow{\alpha} P'$ there exists a Q' such that $\Psi \triangleright (vx)(vy)P \xrightarrow{\alpha} Q'$ and $(\Psi, P', Q') \in \mathcal{S}^*$.

We first define a relation \mathcal{R} that safely approximates \mathcal{S}^* (i.e. $\mathcal{R} \subseteq \mathcal{S}^*$) and is easier to work with.

$$\mathcal{R} \stackrel{\text{def}}{=} \{(\Psi, (v\tilde{a})P, (v\tilde{b})P) : \tilde{a}\#\Psi \wedge \tilde{a} \equiv \tilde{b}\}.$$

By induction on the length of \tilde{a} , we get that for all $\tilde{a}, \tilde{b}, \Psi, P$ such that $\tilde{a}\#\Psi$ and $\tilde{a} \equiv \tilde{b}$ we have $(\Psi, (v\tilde{a})P, (v\tilde{b})P) \in \mathcal{S}^*$. From this follows that the relation $\mathcal{R} \subseteq \mathcal{S}^*$; in order to show that the derivatives $(\Psi, P', Q') \in \mathcal{S}^*$ after a simulation step, we instead prove $(\Psi, P', Q') \in \mathcal{R}$.

The simulation proof is by case analysis on the derivations of transitions of $(vy)(vx)P$. We here focus on the following derivation.

$$\text{SCOPE} \frac{\text{BRCLOSE} \frac{\Psi \triangleright P \xrightarrow{\overline{!M}(v\tilde{a})N} P'}{\Psi \triangleright (vx)P \xrightarrow{\tau} (vx)(v\tilde{a})P'} \quad x \in \mathfrak{n}(M), x\#\Psi}{\Psi \triangleright (vy)(vx)P \xrightarrow{\tau} (vy)(vx)(v\tilde{a})P'} \quad y\#\tau, \Psi$$

We assume that $\tilde{a}\#\Psi, P, M, x, y$. There are three cases to consider.

1. $y\#\overline{!M}(v\tilde{a})N$: We have the following transition.

$$\text{BRCLOSE} \frac{\text{SCOPE} \frac{\Psi \triangleright P \xrightarrow{\overline{!M}(v\tilde{a})N} P'}{\Psi \triangleright (vy)P \xrightarrow{\tau} (vy)P'} \quad y\#\overline{!M}(v\tilde{a})N, \Psi}{\Psi \triangleright (vx)(vy)P \xrightarrow{\tau} (vx)(v\tilde{a})(vy)P'} \quad x \in \mathfrak{n}(M), x\#\Psi$$

Since $x, y, \tilde{a}\#\Psi$ and $(x, \tilde{a}, y) \equiv (y, x, \tilde{a})$ we have $(\Psi, (vy)(v\tilde{a})P', (vx)(v\tilde{a})(vy)P') \in \mathcal{R} \subseteq \mathcal{S}^*$.

2. $y \in \mathfrak{n}(\overline{!M}(v\tilde{a})N)$ and $y \in \mathfrak{n}(M)$: We have the following transition.

$$\text{SCOPE} \frac{\text{BRCLOSE} \frac{\Psi \triangleright P \xrightarrow{\overline{!M}(v\tilde{a})N} P'}{\Psi \triangleright (vy)P \xrightarrow{\tau} (vy)(v\tilde{a})P'} \quad y \in \mathfrak{n}(M), y\#\Psi}{\Psi \triangleright (vx)(vy)P \xrightarrow{\tau} (vx)(vy)(v\tilde{a})P'} \quad x\#\tau, \Psi$$

Since $x, y\#\Psi$ and $y, x \equiv x, y$ we have $(\Psi, (vy)(v\tilde{a})P', (vx)(vy)(v\tilde{a})P') \in \mathcal{R} \subseteq \mathcal{S}^*$

3. $y \in \mathfrak{n}(\overline{!M}(v\tilde{a})N)$ and $y\#M$: We then have $y \in \mathfrak{n}(N)$ and derive

$$\text{BRCLOSE} \frac{\text{BROPEN} \frac{\Psi \triangleright P \xrightarrow{\overline{!M}(v\tilde{a})N} P'}{\Psi \triangleright (vy)P \xrightarrow{\tau} (vy)(v\tilde{a})P'} \quad y\#\tilde{a}, \Psi, M, y \in \mathfrak{n}(N)}{\Psi \triangleright (vx)(vy)P \xrightarrow{\tau} (vx)(vy)(v\tilde{a})P'} \quad x \in \mathfrak{n}(M), x\#\Psi$$

Since $x, y\#\Psi$ and $y, x \equiv x, y$ we have $(\Psi, (vy)(vx)(v\tilde{a})P', (vx)(vy)(v\tilde{a})P') \in \mathcal{R} \subseteq \mathcal{S}^*$. \square

The soundness proof for scope extension uses the same ideas as the proof of Lemma 12.

3.2 Motivating the requisites

An apparently simpler way to define broadcast connectivity is to have just one binary connectivity predicate relating input and output prefixes, as \leftrightarrow does for unicast communication. However, such a predicate would need to be transitive and symmetric for Theorem 11 to hold, for the same reasons as in the original psi-calculus (detailed in [5]). In wireless broadcast communication systems, symmetry and transitivity do not necessarily hold, and the requirements would not be reasonable.

A weaker version of condition 2 (respectively 1) of Definition 6 would be to require $n(K) \subseteq n(M, \Psi)$ whenever $\Psi \vdash K \dot{\succ} M$ (respectively $\Psi \vdash M \dot{\prec} K$). However, this leads to structural equivalence not being sound for bisimulation: the scope extension case of Theorem 11 fails, as we see in the following example.

Example 13 We let $\mathbf{A} = \mathcal{P}_{\text{fin}}(\mathcal{N})$ with $\mathbf{1} = \emptyset$ and $\otimes = \cup$. We let $\mathbf{T} = \mathcal{N}$ and $\mathbf{C} = \{a \dot{\leftrightarrow} b, a \dot{\prec} b, a \dot{\succ} b : a, b \in \mathcal{N}\}$. We define \vdash by $\forall \Psi, a, b, \Psi \vdash b \dot{\prec} b$ iff $b \in \Psi$ and $\Psi \vdash b \dot{\succ} a$ iff $b \in \Psi$. Note that this definition of entailment does not satisfy Definition 6, since we may have $\Psi \vdash b \dot{\succ} a$ for some $b \neq a$.

We let $P := (\nu a)(\{a\} \mid \bar{a}.\mathbf{0} \mid \underline{c}.\mathbf{0})$. Here $\mathbf{1} \triangleright P \xrightarrow{\tau} (\nu a)(\{a\} \mid \mathbf{0} \mid \mathbf{0})$. However, P results from scope extension from $Q := (\nu a)(\{a\} \mid \bar{a}.\mathbf{0} \mid \underline{c}.\mathbf{0})$, but Q does not have a corresponding transition under frame $\mathbf{1}$.

In contrast to unicast actions, the support of the subjects of broadcast actions is always included in the support of the process generating the action. This result is used in the proof of the scope extension case of Theorem 11 to show that a scope extension does not enable any additional broadcast communication.

Lemma 14 *If $\Psi \triangleright P \xrightarrow{\overline{1K}(\nu \bar{a})N} P'$ or $\Psi \triangleright P \xrightarrow{?KN} P'$ then $n(K) \subseteq n(P)$.*

Proof By induction on the derivation, using Definition 6 at the base cases. □

4 Modelling network topology changes

When modelling wireless protocols, one important concern is dealing with connectivity changes. We here give general descriptions of methods of modelling different connectivity configurations using assertions.

The main idea is to allow for different generations of assertions by tagging assertions with a time. Only the most recent generation is used; a generation is made obsolete by composition with an assertion from a later generation. We here consider broadcast connectivity, but this technique can also be used in other scenarios where there is a need to retract assertions. In the following we assume a set of terms $\mathbf{B} \subseteq \mathbf{T}$ used as broadcast channels and in prefixes; we let B, B' range over elements of \mathbf{B} .

4.1 Simple topology

Here assertions are finite sets of connectivity information ($M \dot{\prec} K$ resp. $K \dot{\succ} M$), labelled with a time, with the empty set at time 0 as the unit assertion. Assertion composition intuitively computes the union of all connectivity

information labelled with the most recent generation. The sets \mathbf{C} and \mathbf{A} are defined using constructors operating on terms. We define substitution on \mathbf{C} and \mathbf{A} homomorphically on their structure. For simplicity, we assume that no rewriting happens in broadcast output, i.e. $\dot{\prec}$ is the equality relation of \mathbf{B} .

Formally,

$$\begin{aligned} \mathbf{C} &\triangleq \{\perp\} \cup \{\text{currentGeneration}(g) : g \in \mathbb{N}\} \cup \\ &\quad \{K \dot{\succ} M : K, M \in \mathbf{T}\} \cup \{M \dot{\prec} K : K, M \in \mathbf{T}\} \\ \mathbf{A} &\triangleq \mathbb{N} \times \mathcal{P}_{\text{fin}}(\{K \dot{\succ} M : K, M \in \mathbf{T}\}) \\ \mathbf{1} &\triangleq \langle 0, \emptyset \rangle \end{aligned}$$

$$\langle g, S \rangle \otimes \langle g', T \rangle \triangleq \begin{cases} \langle g, S \rangle & \text{if } g > g' \\ \langle g', T \rangle & \text{if } g < g' \\ \langle g, S \cup T \rangle & \text{if } g = g' \end{cases}$$

$$\begin{aligned} \langle g, S \rangle \vdash \text{currentGeneration}(g') &\text{ iff } g = g' \\ \langle g, S \rangle \vdash B \dot{\prec} B' &\text{ if } B = B' \\ \langle g, S \rangle \vdash B \dot{\succ} B' &\text{ if } B \dot{\succ} B' \in S \text{ and } n(B) \subseteq n(B') \end{aligned}$$

Proposition 15 *Given \mathbf{T} with a substitution function satisfying the requirements of Sect. 2, the definitions of $\mathbf{C}, \mathbf{A}, \otimes, \mathbf{1}$, and \vdash as above and $(M \dot{\leftrightarrow} N) \triangleq \perp$ satisfy the requirements of a broadcast psi-calculus.*

The assertion $\langle g, \{B \dot{\succ} B'\} \rangle$ states that B' is in-connected to B in generation g if $n(B) \subseteq n(B')$. The condition $\text{currentGeneration}(g)$ is used to test whether g is the most recent generation. It is needed for assertion equivalence to be compositional: without this condition, we would have $\langle 0, \{M \dot{\succ} K\} \rangle \simeq \langle 1, \{M \dot{\succ} K\} \rangle$ and $\langle 0, \{M \dot{\succ} K\} \rangle \otimes \langle 1, \{K \dot{\succ} M\} \rangle \not\approx \langle 1, \{M \dot{\succ} K\} \rangle \otimes \langle 1, \{K \dot{\succ} M\} \rangle$, contradicting compositionality.

As an example, we can define a topology controller (assuming a suitable encoding of the τ prefix):

$$\begin{aligned} T = & \langle \langle 1, \emptyset \rangle \rangle \mid \tau . (\langle \langle 2, \{K \dot{\succ} M, K \dot{\succ} N\} \rangle \rangle \\ & \mid \tau . (\langle \langle 3, \{K \dot{\succ} M\} \rangle \rangle)) \end{aligned}$$

In $P \mid T$, the process P broadcasts on K while T manages the topology. Initially, $\mathcal{F}(T) = \langle 1, \emptyset \rangle$ and the broadcast is disconnected; after $T \xrightarrow{\tau} T'$ then $\mathcal{F}(T') = \langle 2, \{K \dot{\succ} M, K \dot{\succ} N\} \rangle$ and a broadcast on K can be received on both M and N , and after $T' \xrightarrow{\tau} T''$ then a broadcast can be received only on M , since $\mathcal{F}(T'') = \langle 3, \{K \dot{\succ} M\} \rangle$.

Such a connectivity controller can also implement standard mobility models [12] over a discretised finite space. More fine-grained mobility models can be implemented by associating a generation with each possible connection, together with a flag for whether the connection is possible or not. In such a model, assertion $\langle \langle 0, M \dot{\succ} K, \text{true} \rangle \rangle$ states that the link $M \dot{\succ} K$ is enabled in its generation 0.

4.2 Scoped topology

As a variation of the example above, we define a model where every name d corresponds to a broadcast channel with dynamic topology. The use of a name in the broadcast channel allows us to restrict its scope.

$$\begin{aligned}
\mathbf{B} &\triangleq \{\text{Bs}(d) : d \in \mathcal{N}\} \cup \{\text{Br}(M, d) : M \in \mathbf{T}, d \in \mathcal{N}\} \cup \mathcal{N} \\
\mathbf{C} &\triangleq \{\perp\} \cup \{\text{currentGeneration}(g, K) : g \in \mathbb{N}, K \in \mathbf{T}\} \cup \\
&\quad \{\text{Bs}(M) \dot{<} K) : M, K \in \mathbf{T}\} \cup \{K \dot{>} \text{Br}(M, N) : M, N, K \in \mathbf{T}\} \\
\mathbf{A} &\triangleq \mathbf{T} \rightarrow_{\text{fin}} \mathbb{N} \times \mathcal{P}_{\text{fin}}(\{\text{Conn}(M, N) : M, N \in \mathbf{T}\}) \\
\mathbf{1} &\triangleq \emptyset \\
(\Psi \otimes \Psi')(M) &\triangleq \\
&\begin{cases} \langle g, S \rangle & \text{if } \Psi(M) = \langle g, S \rangle \wedge (M \notin \text{dom}(\Psi') \vee \\ & \quad (\Psi'(M) = \langle j, T \rangle \wedge g > g')) \\ \langle g', T \rangle & \text{if } \Psi'(M) = \langle g', T \rangle \wedge (M \notin \text{dom}(\Psi) \vee \\ & \quad (\Psi(M) = \langle g, S \rangle \wedge g < g')) \\ \langle g, S \cup T \rangle & \text{if } \Psi(M) = \langle g, S \rangle \wedge \Psi'(M) = \langle g, T \rangle \end{cases} \\
\Psi &\vdash \text{currentGeneration}(g, d) \text{ if } \Psi(d) = \langle g, S \rangle \\
\Psi &\vdash \text{Bs}(c) \dot{<} d \text{ if } c = d \\
\Psi &\vdash c \dot{>} \text{Br}(N, d) \text{ if } c = d \text{ and } \Psi(c) = \langle g, S \rangle \text{ with } \text{Conn}(N, d) \in S
\end{aligned}$$

Proposition 16 *Given \mathbf{T} with a substitution function satisfying the requirements of Sect. 2, the definitions of \mathbf{C} , \mathbf{A} , \otimes , $\mathbf{1}$, and \vdash as above and $(M \dot{\leftrightarrow} N) \triangleq \perp$ satisfy the requirements of a broadcast psi-calculus.*

We can then define a topology controller which gradually changes the topology from fully disconnected to “ a listens on d and b listens on d ”:

$$\begin{aligned}
T &= (\langle d \mapsto \langle 1, \emptyset \rangle \rangle \mid \tau. (\langle d \mapsto \langle 2, \{\text{Conn}(a, d)\} \rangle \\
&\quad \mid \tau. (\langle d \mapsto \langle 3, \{\text{Conn}(a, d), \text{Conn}(b, d)\} \rangle)))
\end{aligned}$$

We now put a process P inside the scope of d in parallel with the topology controller as $(\nu d)(P \mid T)$. This ensures that P can communicate using broadcast on channel d while letting T , but not the environment, influence the topology. Moreover, no process in the environment can receive broadcasts from P (unless having previously received the bound name d). In this way, scoped topology enables hierarchical modelling of subsystems using wireless broadcast.

5 The LUNAR protocol in psi

In this section we present a model of the LUNAR routing protocol for mobile ad hoc networks [32,33]. LUNAR is intended for small wireless networks, ca 15 nodes, with a network diameter of 3 hops. It does not handle route reparation, caching, etc., and routes must be re-established every few seconds. It is reasonably simple in comparison with many other ad hoc routing protocols and allows us to focus on properties such as dynamic connectivity and broadcasting. It has previously been verified in [35,36] using SPIN and UPPAAL; our model is significantly more succinct and at an abstraction level closer to the specification.

The LUNAR protocol is at “layer 2.5”, between the link and network layers in the Internet protocol stack. Addressing is by pairs of MAC/Ethernet addresses and 64-bit selectors, similarly to the IP address and port number used in UDP/TCP. The selectors are used to find the appropriate packet handler through the Forwarding Information Base (FIB) table.

Below, we define a psi-calculus for modelling the LUNAR protocol. In an effort to keep our model simple, we abstract from details such as time-to-live (TTL) fields in messages, optional protocol fields, and globally unique host identifiers. These abstractions are similar to those made in [35,36]. We do not deal with time explicitly. In the SPIN verification, time is handled at an abstract level by using the Promela `timeout` predicate which is true when no other statement is executable, and checking that in this case, the protocol has succeeded in delivering a message (cf. Theorem 18).

5.1 The LUNAR broadcast psi-calculus

Channels are of two kinds: broadcast channels are terms node_i with (for simplicity) empty support, whose connectivity is given by the $\dot{>}$ and $\dot{<}$ predicates as defined in Sect. 4.1, and unicast channels that are pairs $\langle \text{sel}, \text{mac} \rangle$ where sel is a selector name and mac is a MAC address name. The sel part can also be a $\text{RouteOf}(\text{node}, \text{ip})$ construction, which looks up the route of an IP address ip in the routing table of the node node . Special channels $\langle \text{delivered}, \text{node}_i \rangle$ are used to signal delivery of a packet to the IP layer. Assertions are used to record requests originated at the local node with $\text{Redirected}(\text{node}, \text{sel})$, and with $\text{HaveRoute}(\text{node}, \text{destip}, \text{hops}, \text{sel})$ to specify found routes. The conditions contain predicates for testing whether a route has been found ($\text{HaveRoute}(\text{node}, \text{ip})$) and whether a selector has been used for a request originating at the local node ($\text{Redirected}(\text{node}, \text{sel})$), and to extract the forwarder of a route ($\langle \text{RouteOf}(\text{node}, \text{ip}), x \rangle \dot{\leftrightarrow} \langle \text{sel}, x \rangle$).

LUNAR protocol messages are of two types. The first is a route request message $\text{RREQ}(\text{selector}, \text{targetIP}, \text{replyTo})$, where the selector identifies the request, targetIP is the IP address the route should reach, and replyTo is the $\langle \text{sel}, \text{mac} \rangle$ channel the response should be sent to. The second is a route reply message, $\text{RREP}(\text{hops}, \text{fwdptr})$, where hops is the number of hops to the destination, and fwdptr is a forwarding pointer, i.e. a $\langle \text{sel}, \text{mac} \rangle$ channel where packets can be sent.

The parameters of the LUNAR broadcast psi-calculus extend the simple topology calculus in Sect. 4.1. We define substitution in the standard way, as the syntactic replacement of names by terms. The sets \mathbf{T} , \mathbf{C} , and \mathbf{A} are defined recursively using constructors operating on terms in order to be closed under substitution.

$$\begin{aligned}
 \mathbf{T} &\triangleq \mathcal{N} \cup \{\text{node}_i : i \in \mathbb{N}\} \cup \{\text{delivered}\} \cup \\
 &\quad \{\text{RREQ}(Ser, \text{TargetIp}, Rep) : Ser, \text{TargetIp}, \\
 &\quad Rep \in \mathbf{T}\} \cup \\
 &\quad \{\text{RREP}(i, Fwd) : i, Fwd \in \mathbf{T}\} \cup \\
 &\quad \{\text{RouteOf}(Node, Ip) : Node, Ip \in \mathbf{T}\} \cup \\
 &\quad \{\langle Sel, N \rangle : Sel, N \in \mathbf{T}\} \cup \{N + 1 : N \in \mathbf{T}\} \cup \{0\} \\
 \mathbf{C} &\triangleq \{M = N, M \leftrightarrow N, \text{HaveRoute}(M, N), \\
 &\quad \text{Redirected}(M, N) : M, N \in \mathbf{T}\} \cup \\
 &\quad \{K \succ M : K, M \in \mathbf{T}\} \cup \{M \prec K : K, M \in \mathbf{T}\} \cup \\
 &\quad \{\text{currentGeneration}(g) : g \in \mathbb{N}\} \cup \{\neg\phi : \phi \in \mathbf{C}\} \\
 \mathbf{A} &\triangleq \mathbb{N} \times \mathcal{P}_{\text{fin}}(\{\langle K \succ M \rangle : K, M \in \mathbf{T}\}) \times \\
 &\quad \mathcal{P}_{\text{fin}}(\{\text{HaveRoute}(M, N_1, i, N_2) : i, M, N_1, \\
 &\quad N_2 \in \mathbf{T}\} \cup \{\text{Redirected}(M, N) : M, N \in \mathbf{T}\}) \\
 \mathbf{1} &\triangleq \langle 0, \emptyset, \emptyset \rangle
 \end{aligned}$$

$$\langle g, S, A \rangle \otimes \langle g', T, B \rangle \triangleq \begin{cases} \langle g, S, A \cup B \rangle & \text{if } g > g' \\ \langle g', T, A \cup B \rangle & \text{if } g < g' \\ \langle g, S \cup T, A \cup B \rangle & \text{if } g = g' \end{cases}$$

Given $\Psi = \langle g, S, A \rangle$, we let \mathcal{R}_Ψ be the symmetric and transitive closure of the relation

$$\begin{aligned}
 &\{(\langle a, b \rangle, \langle a, b \rangle) : a, b \in \mathcal{N}\} \\
 &\cup \{(\langle \text{delivered}, \text{node}_i \rangle, \langle \text{delivered}, \text{node}_i \rangle) : i \in \mathbb{N}\} \cup \\
 &\{(\langle \text{RouteOf}(\text{node}_i, a), x \rangle, \langle b, x \rangle) \\
 &\quad : i \in \mathbb{N}, j \in \mathbf{T}, \text{HaveRoute}(\text{node}_i, a, j, b) \in A\}
 \end{aligned}$$

Entailment is then defined as follows.

$$\begin{aligned}
 \Psi &\vdash a = a, a \in \mathcal{N} \\
 \Psi &\vdash M \leftrightarrow N \text{ iff } (M, N) \in \mathcal{R}_\Psi \\
 \langle g, S, A \rangle &\vdash \text{currentGeneration}(g) \\
 \Psi &\vdash M \prec N \text{ iff } M = N \\
 \langle g, S, A \rangle &\vdash M \succ N \text{ iff } M \succ N \in S \\
 &\quad \text{and } n(M) \subseteq n(N) \\
 \langle g, S, A \cup \{\text{HaveRoute}(\text{node}_i, a, j, b)\} \rangle & \\
 &\quad \vdash \text{HaveRoute}(\text{node}_i, a) \\
 \langle g, S, A \cup \{\text{Redirected}(\text{node}_i, s)\} \rangle & \\
 &\quad \vdash \text{Redirected}(\text{node}_i, s) \\
 \Psi &\vdash \neg\phi \text{ if not } \Psi \vdash \phi
 \end{aligned}$$

Theorem 17 *The LUNAR psi-calculus defined above satisfies all the requisites of a broadcast psi-calculus.*

This theorem has been formally proved in Isabelle/Nominal [25]. A sketch outlining the main ideas of the proof follows:

Proof (Sketch) The requisites on the support of the broadcast channels are immediate from the definition. It is straightforward to show the Abelian monoid laws for \otimes , $\mathbf{1}$. Transitivity and symmetry of channel equivalence holds by definition. The only non-trivial property is compositionality: we establish that $\Psi \otimes \Psi_1 \vdash \phi$ and $\Psi_1 \simeq \Psi_2$ implies $\Psi \otimes \Psi_2 \vdash \phi$ by induction on the structure of the condition ϕ . The only inductive step is for negation, and this follows by symmetry of \simeq . If

ϕ is a broadcast connectivity condition or $\text{currentGeneration}(g)$, the proof is by case distinction on the relative generations of Ψ_1, Ψ_2 , and Ψ . If ϕ is a channel equivalence, an inner induction on the length of the chain of the involved HaveRoute elements in $\Psi \otimes \Psi_1$ is necessary. Each such element is either in Ψ and therefore also in $\Psi \otimes \Psi_2$, or in Ψ_1 . In the latter case, Ψ_1 entails a channel equivalence from this element alone and therefore Ψ_2 entails the same. Thus, Ψ_2 must contain a suitable sequence of HaveRoute elements to derive this channel equivalence; this sequence is then in $\Psi \otimes \Psi_2$. \square

5.2 Representing process identifiers

We use process identifiers to improve the readability of the LUNAR protocol model. However, an astute reader will note that broadcast psi-calculi do not feature process identifiers—rather, replication is used as the mechanism for expressing infinite behaviour. In many other process calculi, process identifiers and recursion can be encoded in a standard fashion using replication, see e.g. [30]. Unfortunately, there is currently no proof that the same encodability results apply to broadcast psi-calculi.

To introduce process identifiers on a more sound theoretical foundation, we combine broadcast psi-calculi with higher-order psi-calculi [23], an orthogonal extension of psi-calculi which allows terms to act as handles to invoke the behaviour of processes. In this setting, process identifiers are simply terms.

Briefly, higher-order psi-calculi introduce the notion of a *clause* $M \Leftarrow P$, meaning that the term M is a handle for invoking P . We extend the entailment relation \vdash so that assertions can entail clauses in addition to conditions. Agents are extended with *invocations* $\text{run } M$, and a single new rule is added to the semantics:

$$\text{INVOCATION} \frac{\Psi \vdash M \Leftarrow P \Psi \triangleright P \xrightarrow{\alpha} P'}{\Psi \triangleright \text{run } M \xrightarrow{\alpha} P'}$$

The calculi that result from adding the above-mentioned extensions to broadcast psi-calculi will be referred to as *higher-order broadcast psi-calculi*. We use Isabelle/Nominal to formally prove that all the meta-theoretical results presented in Sect. 3.1 apply not only to broadcast psi-calculi, but also to higher-order broadcast psi-calculi—hence, we feel justified in claiming that broadcast and higher-order are orthogonal extensions. The proof scripts are available online [25].

Further, higher-order psi-calculi feature a lifting technique whereby an arbitrary first-order psi-calculus can be lifted to a corresponding canonical higher-order psi-calculus, extending it with parametrised clauses. In a canonical higher-order psi-calculus, sets of parametrised clauses on the form $M(N) \Leftarrow P$ are added to the assertions, such that $\{M(N) \Leftarrow P\} \vdash M(N[\tilde{x} := \tilde{T}]) \Leftarrow P[\tilde{x} := \tilde{T}]$.

Fig. 1 Part 0: the initialisation step at the node that wishes to discover a route

$$\text{LunARP}(mynode, mymac, destip) \Leftarrow (\nu rchosen, schosen) \left(\begin{array}{l} !\langle rchosen, mymac \rangle(x) . \text{SRrepHandler}(mynode, mymac, destip, x) \\ | \langle \text{Redirected}(mynode, schosen) \rangle \\ | mynode \langle \text{RREQ}(schosen, destip, \langle rchosen, mymac \rangle) \rangle . \mathbf{0} \end{array} \right)$$

Fig. 2 Part 1: RREQ packet handler, and Part 2: Target found branch

$$\text{ReqHandler}(mynode, mymac, myip, \text{RREQ}(schosen, destip, repchn)) \Leftarrow \begin{array}{l} \text{if } \langle \text{Redirected}(mynode, schosen) \rangle \text{ then } \mathbf{0} \\ \text{else } \tau . \left(\langle \text{Redirected}(mynode, schosen) \rangle \mid \right. \\ \quad \text{if } destip = myip \text{ then} \quad \quad \quad /* \text{ Part 2: Target found */} \\ \quad \quad (\nu rchosen) \\ \quad \quad \left(\begin{array}{l} !\langle rchosen, mymac \rangle(x) . \text{IPdeliver}(x, mynode) \\ | repchn \langle \text{RREP}(0, \langle rchosen, mymac \rangle) \rangle . \mathbf{0} \end{array} \right) \\ \quad \text{else} \\ \quad \quad (\nu rchosen) \\ \quad \quad \left(\begin{array}{l} !\langle rchosen, mymac \rangle(x) . \text{IRrepHandler}(mymac, repchn, x) \\ | mynode \langle \text{RREQ}(schosen, destip, \langle rchosen, mymac \rangle) \rangle . \mathbf{0} \end{array} \right) \end{array}$$

$$\text{IRrepHandler}(mymac, repchn, \text{RREP}(hops, fwdptr)) \Leftarrow (\nu rchosen) \left(\begin{array}{l} !\langle rchosen, mymac \rangle(x) . \overline{fwdptr} x . \mathbf{0} \\ | repchn \langle \text{RREP}(hops + 1, \langle rchosen, mymac \rangle) \rangle . \mathbf{0} \end{array} \right)$$

Fig. 3 Part 3: Intermediate RREP packet handler

$$\text{SRrepHandler}(mynode, mymac, destip, \text{RREP}(hops, fwdptr)) \Leftarrow (\nu rchosen) \left(\begin{array}{l} !\langle rchosen, mymac \rangle(x) . \overline{fwdptr} x . \mathbf{0} \\ | \langle \text{HaveRoute}(mynode, destip, hops, rchosen) \rangle \end{array} \right)$$

Fig. 4 Part 4: Source RREP packet handler

In the following, we will implicitly be representing clauses using this feature of the canonical higher-order calculus corresponding to the LUNAR broadcast psi-calculus of Sect. 5.1.

5.3 The psi-calculus model of the LUNAR protocol

Figures 1, 2, 3, 4, 5, 6, and 7 describe our psi-calculus model of the LUNAR protocol. Process declarations are of the form $M(\tilde{N}) \Leftarrow P$, where M is a process identifier (and also a term, implicitly included in \mathbf{T}), \tilde{N} a list of terms where occurrences of names are binding, and P is a process s.t. $\mathfrak{n}(P) \subseteq \mathfrak{n}(\tilde{N})$. In a process, we write $M(\tilde{N})$ for invoking a process declaration $M(\tilde{K}) \Leftarrow P$ such that $\tilde{N} = \tilde{K}[\tilde{x} := \tilde{L}]$ with $\tilde{x} = \mathfrak{n}(\tilde{K})$, resulting in the process $P[\tilde{x} := \tilde{L}]$. For our purposes, lists can be adequately represented using the pairing construct included in the term language. We write **if** φ **then** P **else** Q for **case** $\varphi : P \mid \neg\varphi : Q$ and assume a suitable encoding of the τ prefix.

Our model of the protocol closely follows the informal protocol description in [32, Section 4]. Each figure in

$$\text{IPdeliver}(x, node) \Leftarrow \overline{\langle \text{delivered}, node \rangle} x . \mathbf{0}$$

Fig. 5 Part 5: IP delivery

our model corresponds to one or more of part 0–5 of the protocol description. To allocate a selector, we simply bind a name; to associate (or bind) a selector with a packet handler, we use a replicated process that receives on the unicast channel described by the pair of the selector and our MAC address. An example of this can be seen in the LunARP process declaration in Fig. 1. The description in [32, Section 4, step 0.a] says “Allocate an unused ‘receiver chosen’ selector S and binds it to a transient ‘source RREP packet handler’”, which in our process declaration corresponds to the binding of $rchosen$ and the subprocess $!\langle rchosen, mymac \rangle(x) . \text{SRrepHandler}(mynode, mymac, destip, x)$.

In the informal protocol description [32], the FIB is “abused” (in steps 0.b and 1.b) by installing a null packet handler for the selector created when sending a route request. This FIB entry is only used to detect and avoid circular forwarding of route requests. We model this by an explicit assertion and a matching condition. An example can be seen in the subprocess $\langle \text{Redirected}(mynode, schosen) \rangle$ of the LunARP process declaration, and the test on the first line of the ReqHandler process declaration (Fig. 2) using the $\langle \text{Redirected}(mynode, schosen) \rangle$ condition.

The routing table is modelled using assertions, which illustrates how these can be used as a global data structure. Additions to the routing table are done in the SRrepHandler process definition (Fig. 4), which adds $\langle \text{HaveRoute}(mynode, destip, hops, rchosen) \rangle$ to the environment. Such assertions together form the routing table, which is tested in the IPtransmit pro-

Fig. 6 Broadcast handler

$$\text{BrdHandler}(mynode, mac, ip) \Leftarrow \underline{mynode}(\lambda s, t, r) \text{RREQ}(s, t, r) \cdot \left(\begin{array}{l} \text{RreqHandler}(mynode, mac, ip, \text{RREQ}(s, t, r)) \\ | \text{BrdHandler}(mynode, mac, ip) \end{array} \right)$$

Fig. 7 IP transmission: if have route, send it to local forwarder, else ask for route

$$\text{IPtransmit}(mynode, mymac, destip, pkt) \Leftarrow \begin{array}{l} \text{if HaveRoute}(mynode, destip) \text{ then } \overline{\text{RouteOf}(mynode, destip), mymac} \text{ pkt} \cdot \mathbf{0} \\ \text{else LunARP}(mynode, mymac, destip) \end{array}$$

cess definition (Fig. 7) using the $\text{HaveRoute}(mynode, destip)$ condition.

For simplicity, we do not model route timeouts and the deletion of routes, but this could be done using the mechanism in Sect. 4.

The LUNAR procedure for route discovery starts when a node wants to send a message to a node it does not already have a route to (Fig. 7, **else** branch). It then (Fig. 1) associates a fresh selector with a response packet handler and broadcasts a Route Request (RREQ) message to its neighbours. A node that receives a RREQ message (Fig. 2) for its own IP address sets up a packet handler to deliver IP packets and includes the corresponding selector in a response Route Reply (RREP) message to the reply channel found in the RREQ message. If the RREQ message was not for its own IP address, the message is re-broadcast after replacing the reply channel with a freshly allocated reply selector and its own MAC address. When such an intermediary node receives a RREP message (Fig. 3), it increments the hop counter and forwards the RREP message to the source of the original RREQ message. When the originator of a RREQ message eventually receives the matching RREP (Fig. 4), it installs a route and informs the IP layer about it. The message can then be resent (Fig. 7, **then** branch) and delivered (Fig. 5) by unicast messages through the chain of intermediary forwarding nodes.

We show the basic correctness of the model by the following theorem, which in essence corresponds to the correct operation of an ad hoc routing protocol [36, Definition 1]: if there is a path between two nodes, the protocol finds it, and it is possible to send packets along the path to the destination node.

The system to analyse consists of n nodes with their respective broadcast handler; node 0 attempts to transmit a packet to the IP address of node n .

$$\text{Spec}_n(pkt, ip_0, \dots, ip_n) \Leftarrow (vmac_0, \dots, mac_n) \left(\begin{array}{l} \prod_{0 \leq i \leq n} \text{BrdHandler}(node_i, mac_i, ip_i) \\ | \text{IPtransmit}(node_0, mac_0, ip_n, pkt) \end{array} \right)$$

Theorem 18 *If Ψ connects $node_0$ and $node_n$ via a node $node_i$ (i.e. $\Psi \vdash node_0 \dot{\succ} node_i$ and $\Psi \vdash node_i \dot{\succ} node_n$), then*

$$\Psi \mid (vip_0, \dots, ip_n) \text{Spec}_n(pkt, ip_0, \dots, ip_n)$$

$$\Longrightarrow \overline{\text{delivered}, node_n} \text{pkt} \rightarrow \Psi \mid (vip_0, \dots, ip_n) S$$

and $\mathcal{F}(S) \vdash \text{HaveRoute}(node_0, ip_n)$, where \Longrightarrow stands for an interleaving of τ and broadcast output transitions.

Proof By following transitions. \square

The SPIN verification performed in [36] checks the same reachability property, for up to five nodes. Our analysis is valid for any n , but is limited to a configuration where the sender (node 0) and the receiver (node n) are only separated by a single node. This limitation is due to the labour of manually following transitions in a non-trivial specification. We are currently working on remedies for this: firstly by extending our symbolic semantics for pi-calculi [15], secondly by implementing the symbolic semantics in our tool for automatic verification [13], and thirdly and orthogonally, by implementing the LUNAR model in Isabelle/Nominal. These remedies are still work in progress. In the Isabelle approach, we hope to prove the following conjecture.

Conjecture 19 If Ψ connects $node_0$ and $node_n$ via k proxy nodes pn_1, \dots, pn_k , where $\{pn_1, \dots, pn_k\} \subseteq \{node_1, \dots, node_{n-1}\}$ (i.e. $\Psi \vdash node_0 \dot{\succ} pn_1, pn_1 \dot{\succ} pn_2, \dots, pn_{k-1} \dot{\succ} pn_k, pn_k \dot{\succ} node_n$), then

$$\Psi \mid (vip_0, \dots, ip_n) \text{Spec}_n(pkt, ip_0, \dots, ip_n)$$

$$\Longrightarrow \overline{\text{delivered}, node_n} \text{pkt} \rightarrow \Psi \mid (vip_0, \dots, ip_n) S$$

and $\mathcal{F}(S) \vdash \text{HaveRoute}(node_0, ip_n)$, where \Longrightarrow stands for an interleaving of τ and broadcast output transitions.

The definition of BrdHandler illustrates a peculiarity of broadcast semantics: a reader well versed in pi-calculus specifications with replication and recursion may consider a more concise variant of the definition using replication instead of recursion, e.g.

$$\text{BrdHandler}'(mynode, mac, ip) \Leftarrow$$

$$\! \underline{mynode}(\lambda s, t, r) \text{RREQ}(s, t, r) \cdot$$

$$\text{RreqHandler}(mynode, mac, ip, \text{RREQ}(s, t, r))$$

However, when the input prefix is over a broadcast channel, as is the case here, the two are not equivalent since a single communication with $\text{BrdHandler}'$ may result in arbitrarily many RreqHandler processes, while BrdHandler only results in one.

6 Related work

Process calculi with broadcast communication go back to the early 1980s. Milner developed SCCS [20] as a generalisation of CCS [19] to include multiway communication, of which broadcast can be seen as a special case. At the same time, Austria and Boudol presented MEIJE [2] as a semantic basis for high-level hardware definition languages.

The first process calculus to seriously consider broadcast with an asynchronous parallel composition was CBS [26, 27]. Its development is recorded in a series of papers, examining it from many perspectives. The main focus is on employing broadcast as a high-level programming paradigm. CBS was later extended to the pi-calculus in the $b\pi$ formalism [7]. Here, the broadcast communication channels are names that can be scoped and transmitted between agents. The main point of this work is to establish a separation result in expressiveness: in the pi-calculus, broadcast cannot be uniformly encoded by unicast.

Recent advances in wireless networks have created a renewed interest in the broadcast paradigm. The first process calculus with this in mind was probably CBS[‡] [22]. This is a development of CBS to include varying interconnection topologies. Input and output are performed on a universal ether, and transitions are indexed with topologies that are sets of connectivity graphs; the connectivity graph matters for the input rule (reception is possible from any connected location). Main applications are on cryptography and routing protocols in mobile ad hoc wireless networks. CBS[‡] has been followed by several similar calculi. In CWS [16, 18], the focus is on modelling low-level interference. Communication actions have distinct beginnings and endings, and two actions may interfere if one begins before another has ended. The main result is an operational correspondence between a labelled semantics and a reduction semantics. CMAN [10] is a high-level formalism extended with data types, just as the applied pi-calculus extends the original pi-calculus. Data can contain constructors and destructors. There are results on properties of weak bisimulation and an analysis of a cryptographic routing protocol. In the ω -calculus [31], emphasis is on expressing connectivity using sets of group names. An extension also includes separate unicast channels, making this formalism the first to accommodate both multicast and unicast in wireless networks. There are results about strong bisimulation and a verification of a mobile ad hoc network leader election protocol through weak bisimulation.

RBPT [9] is similar and uses an alternative technique to represent topology changes, leading to smaller state spaces, and is also different in that it can accommodate an asymmetric neighbour relation (to model the fact that A can send to B but not the other way).

$bA\pi$ [11] is an extension of the applied pi-calculus [1] with broadcast, where connectivity information appears explicitly in the process terms and can change non-deterministically during execution. The claimed result of the paper is proving that a weak labelled bisimulation, for which connectivity is irrelevant, coincides with barbed equivalence. However, for the same reasons as in the applied pi-calculus (cf. [4]), labelled bisimilarity is not compositional in $bA\pi$, so the correspondence does not hold. A suggested fix is to remove communication of unicast channels from the calculus. We would finally mention CMN [17]. The claimed result is to compare two different kinds of semantics for a broadcast operation, but it is in error. The labelled transition semantics contains no rule for merging two inputs as in our BRMERGE. As a consequence, parallel composition fails to be associative. Consider the situation where P does an output and Q and R both do inputs. A broadcast communication involving all three agents can be derived from $(P|Q) | R$ but not from $P | (Q|R)$, since in the latter agent the component $Q|R$ cannot make an input involving both Q and R .

It is interesting to compare these formalisms and our broadcast psi from a few important perspectives. Firstly, the broadcast channels are explicitly represented in ω , $b\pi$, CWS, and CMN; they are mobile (in the sense that they can be transmitted) only in $b\pi$. In ω , only unicast channels can be communicated. In broadcast psi, channels are represented as arbitrary mobile data terms that may contain any number of names. Secondly, the data transmitted in CMAN and $bA\pi$ are akin to the applied pi-calculus where data are drawn from an inductively defined set and contain names that may be scoped. In ω and $b\pi$ data are single names which may be scoped; in the other calculi data cannot contain scoped names. In broadcast psi data are arbitrary terms, drawn from a nominal set, and may include higher-order objects as well as bound names. Finally, node mobility is represented explicitly as particular semantic rules in CMAN, CMN, $bA\pi$, and ω , and implicitly in the requirements of bisimulation in CBS[‡] and RBPT. In this respect, broadcast psi-calculi are similar to the latter: connectivity is determined by the assertions in the environment, and in a bisimulation, these may change after each transition.

All calculi presented here use a kind of labelled transition semantics (LTS). $b\pi$, $bA\pi$, CBS[‡], CWS, and ω use it in conjunction with a structural congruence (SC), and the rest (including broadcast psi) do not use a SC. In our experience, SC is efficient in that the definitions become more compact and easy to understand, but introduces severe difficulties

Table 3 Comparison of some process algebras for wireless broadcast

Calculus	Broadcast channels	Scoped data	Mobility	Semantics
$bA\pi$	–	Term	In semantics	LTS + SC and RS
CBS [#]	–	–	In bisimulation	LTS + SC
CWS	Constant	–	–	LTS + SC and RS
CMAN	–	Term	In semantics	LTS and RS
CMN	Name	–	In semantics	LTS and RS
ω	Groups	Name	In semantics	LTS + SC
RBPT	–	–	In bisimulation	LTS
Broadcast psi	Term	Term	In bisimulation	LTS

in making fully rigorous proofs. $bA\pi$, CWS, CMAN, and CMN additionally use a reduction semantics using structural congruence (RS) and prove its agreement with the labelled semantics. Table 3 summarises some of the distinguishing features of calculi for wireless networks.

Finally, broadcast psi is different from the other calculi for wireless broadcast in that there is no stratification of the syntax into processes and networks. There is just the one kind of agent, suitable for expressing both processes operating in nodes and behaviours of entire networks. In contrast, the other calculi has one set of constructs to express processes and another to express networks, sometimes leading to duplication of effort (e.g. there can be a parallel composition operator both at the process and at the network level). Our conclusion is that broadcast psi is conceptually simpler and more efficient for rigorous proofs, and yet more expressive.

7 Conclusion

We have extended the psi-calculi framework with broadcast communication and formally proved using Isabelle/Nominal that the standard congruence and structural properties of bisimilarity hold also after the addition. We have shown how node mobility and network topology changes can be modelled using assertions. Since bisimilarity is closed under all assertions, two bisimilar processes are equivalent in all initial topologies and for all node mobility patterns. We demonstrated expressive power by modelling the LUNAR protocol for route discovery in wireless ad hoc networks and verified a basic correctness property of the protocol.

The proofs of the meta-theoretical results in Sect. 3.1 [28] are formally verified in the interactive theorem prover Isabelle/Nominal. The full formalisation of broadcast psi-calculi amounts to ca 33,000 lines of Isabelle code, of which about 21,000 lines are reused from our earlier work [5].

The model of LUNAR is simplified for clarity and to make manual analysis more manageable. The simplifications are similar to those in the SPIN model by Wibling et al. [36], although we do not model timeouts. Their model [35] is ca 250 lines of SPIN code (excluding comments) while ours is approximately 30 lines. Our model could be improved at the cost of added complexity. For example, allowing broadcast channels to have non-empty support would let us hide broadcast actions, routing tables could be made local by including a scoped name per node, and route deletions could be modelled using generational mechanisms similar to Sect. 4.

We are currently working on extending the symbolic semantics for psi-calculi [15] with broadcast and implementing the semantics in our tool for automatic verification, the Psi-calculi Workbench [13]. We also plan to study weak bisimulation for the broadcast semantics. In order to model more aspects of wireless protocols, we would like to add general resource awareness (e.g. energy or time) to psi-calculi.

References


1. Abadi, M., Fournet, C.: Mobile values, new names, and secure communication. In: Proceedings of POPL '01, pp. 104–115. ACM (2001)
2. Austry, D., Boudol, G.: Algèbre de processus et synchronisation. Theor. Comput. Sci. **30**, 91–131 (1984)
3. Bengtson, J.: Formalising process calculi. PhD thesis, Uppsala University (2010)
4. Bengtson, J., Johansson, M., Parrow, J., Victor, B.: Psi-calculi: mobile processes, nominal data, and logic. In: Proceedings of LICS 2009, pp. 39–48. IEEE (2009)
5. Bengtson, J., Johansson, M., Parrow, J., Victor, B.: Psi-calculi: a framework for mobile processes with nominal data and logic. Logical Methods Comput. Sci. **7**(1) (2011)
6. Borgström, J., Huang, S., Johansson, M., Raabjerg, P., Victor, B., Åman Pohjola, J., Parrow, J.: Broadcast psi-calculi with an application to wireless protocols. In: Barthe, G., Pardo, A., Schneider, G. (eds.) Software Engineering and Formal Methods: SEFM 2011, vol. 7041 of Lec-

- ture Notes in Computer Science, pp. 74–89. Springer (2011)
7. Ene, C., Traian, M.: Expressiveness of point-to-point versus broadcast communications. In: Ciobanu, G., Paun, G. (eds.) Proceedings of FCT, vol. 1684 of Lecture Notes in Computer Science, pp. 258–268. Springer (1999)
 8. Gabbay, M., Pitts, A.: A new approach to abstract syntax with variable binding. *Formal Aspects Comput.* **13**, 341–363 (2001)
 9. Ghassemi, F., Fokink, W., Movaghar, A.: Restricted broadcast process theory. In: Cerone, A., Gruner, S. (eds.) Proceedings of SEFM 2008, pp. 345–354. IEEE Computer Society (2008)
 10. Godskesen, J.C.: A calculus for mobile ad hoc networks. In: Murphy, A.L., Vitek, J. (eds.) Proceedings of COORDINATION 2007, vol. 4467 of Lecture Notes in Computer Science, pp. 132–150. Springer (2007)
 11. Godskesen, J.C.: Observables for mobile and wireless broadcasting systems. In: Proceedings of COORDINATION 2010, vol. 6116 of Lecture Notes in Computer Science, pp. 1–15. Springer (2010)
 12. Godskesen, J.C., Nanz, S.: Mobility models and behavioural equivalence for wireless networks. In: Field, J., Vasconcelos, V.T. (eds.) Proceedings of Coordination Models and Languages, number 5521 in Lecture Notes in Computer Science, pp. 106–122. Springer (2009)
 13. Gutkovas, R.: Exercising psi-calculi: a psi-calculi workbench. M. Sc. thesis, Department of Information Technology, Uppsala University, June (2011)
 14. Johansson, M.: Psi-calculi: a framework for mobile process calculi. PhD thesis, Uppsala University, May (2010)
 15. Johansson, M., Victor, B., Parrow, J.: Computing strong and weak bisimulations for psi-calculi. *J. Logic Algeb. Program.* **81**(3), 162–180 (2012)
 16. Lanese, I., Sangiorgi, D.: An operational semantics for a calculus for wireless systems. *Theor. Comput. Sci.* **411**(19), 1928–1948 (2010)
 17. Merro, M.: An observational theory for mobile ad hoc networks (full version). *J. Inf. Comput.* **207**(2), 194–208 (2009)
 18. Mezzetti, N., Sangiorgi, D.: Towards a calculus for wireless systems. *Electron. Notes Theor. Comput. Sci.* **158**, 331–353 (2006)
 19. Milner, R.: A calculus of communicating systems, vol. 92 of Lecture Notes in Computer Science. Springer (1980)
 20. Milner, R.: Calculi for synchrony and asynchrony. *Theor. Comput. Sci.* **25**, 267–310 (1983)
 21. Milner, R., Parrow, J., Walker, D.: A calculus of mobile processes, part I/II. *J. Inf. Comput.* **100**, 1–77 (1992)
 22. Nanz, S., Hankin, C.: A framework for security analysis of mobile wireless networks. *Theor. Comput. Sci.* **367**(1–2), 203–227 (2006)
 23. Parrow, J., Borgström, J., Raabjerg, P., Åman Pohjola, J.: Higher-order psi-calculi. In: *Mathematical Structures in Computer Science*, FirstView, June (2013)
 24. Pitts, A.M.: Nominal logic, a first order theory of names and binding. *Inf. Comput.* **186**, 165–193 (2003)
 25. Åman Pohjola, J.: Higher-order broadcast psi-calculus formalisation. <http://www.it.uu.se/research/group/mobility/theorem/highorderbroadcast.tar.gz>, July (2013)
 26. Prasad, K.V.S.: A calculus of broadcasting systems. In: Abramsky, S., Maibaum, T.S.E. (eds.) TAPSOFT, vol. 1, volume 493 of Lecture Notes in Computer Science, pp. 338–358. Springer (1991)
 27. Prasad, K.V.S.: A calculus of broadcasting systems. *Sci. Comput. Program.* **25**(2–3), 285–327 (1995)
 28. Raabjerg, P., Åman Pohjola, J.: Broadcast psi-calculus formalisation. <http://www.it.uu.se/research/group/mobility/theorem/broadcastpsi>, July (2011)
 29. Sangiorgi, D.: On the bisimulation proof method. *Math. Struct. Comput. Sci.* **8**(5), 447–479 (1998)
 30. Sangiorgi, D., Walker, D.: *The π -Calculus: A Theory of Mobile Processes*. Cambridge University Press, Cambridge (2001)
 31. Singh, A., Ramakrishnan, C.R., Smolka, S.A.: A process calculus for mobile ad hoc networks. *Sci. Comput. Program.* **75**(6), 440–469 (2010)
 32. Tschudin, C.F.: Lightweight underlay network ad hoc routing (LUNAR) protocol. Internet Draft, Mobile Ad Hoc Networking Working Group, March (2004)
 33. Tschudin, C., Gold, R., Rensfelt, O., Wibling, O.: LUNAR: a lightweight underlay network ad-hoc routing protocol and implementation. In: Proceedings of NEW2AN'04, St. Petersburg, February (2004)
 34. Urban, C., Tasson, C.: Nominal techniques in Isabelle/HOL. In: Nieuwenhuis, R. (ed.) Proceedings of CADE 2005, vol. 3632 of Lecture Notes in Computer Science, pp. 38–53. Springer (2005)
 35. Wibling, O.: SPIN0 and UPPAAL ad hoc routing protocol models. http://www.it.uu.se/research/group/mobility/adhoc/gbt/other_examples (2004)
 36. Wibling, O., Parrow, J., Pears, A.: Automated verification of ad hoc routing protocols. In: Núñez, D.F.E.M. (eds.) *Formal Techniques for Networked and Distributed Systems (FORTE 2004)*, vol. 3235 of Lecture Notes in Computer Science, pp. 343–358. Springer (2004)

Author Biographies



Johannes Borgström is a research fellow in the mobility research group at the IT department, Uppsala University. He obtained a PhD on “Equivalences and calculi for formal verification of cryptographic protocols” from EPFL in 2008 and did postdoctoral research on programming languages and calculi at Microsoft Research in Cambridge and at the UPMARC Linnaeus center at Uppsala University. His research focuses on semantical issues, in particular in process calculi and probabilistic programming languages.



Shuqin Huang is an employee of China Construction Bank. She has been doing research in the mobility research group at the IT department, Uppsala University, in 2010. She obtained a Master Degree on Computer Science in 2011. Her research focused on process algebra, in particular in probabilistic languages.



Björn Victor is an associated professor in the mobility research group at the Department of IT, Uppsala University, where he is affiliated to, e.g. the UPMARC Linnaeus center and the Pro-Fun project. He was the principal author of the first automated tool for pi-calculus (the Mobility Workbench) and developed the Fusion calculus as part of his thesis work.



Magnus Johansson is a co-inventor of the psi-calculi framework and obtained his PhD in 2010 from Uppsala University. He is currently working at Fox Technologies, developing enterprise grade access management software.



Johannes Åman Pohjola is a PhD student in the mobility research group at the Department of IT, Uppsala University. He holds a Master's degree in computer science from Uppsala University. Research interests include concurrency theory and theorem proving.



Palle Raabjerg is a PhD student at the UPMARC Linnaeus center at Uppsala University and is affiliated with the mobility research group at the IT department. He has a Master's degree in computer science from Aalborg University. His research interests include type systems, process calculi, mechanised theorem proving, and tentatively human-computer interaction.



Joachim Parrow is professor in Computing Science at Uppsala University and has previously worked at the Royal Institute of Technology, Stockholm, the Swedish Institute of Computer Science, and the University of Edinburgh. He is a co-inventor of the pi-calculus and the Concurrency Workbench, and an ISI highly cited scientist.