Simple representative instantiations for multicast protocols

Javier Esparza and Monika Maidl

LFCS, School of Informatics, University of Edinburgh

Abstract. We present a formal model for multicast network protocols working on arbitrary tree structures. We give sufficient conditions under which correctness of the protocol for all structures reduces to correctness for the structures with at most one layer of internal nodes. If additional conditions hold, we can reduce further to correctness for one single structure. All these results can be applied to (an abstract version of) the Pragmatic General Multicast protocol.

In the last years, much effort has been devoted to the verification of parameterised distributed systems, i.e., distributed systems designed to work correctly independently of the number of processes taking part in them. Classical examples of these systems are distributed algorithms for leader election, byzantine agreement, or distributed termination, and communication protocols, like cache coherence or network protocols. In this paper we study multicast network protocols working on tree structures. Actually, these protocols run on arbitrary networks, but use a distribution tree to broadcast messages. We assume that this tree has been already been established. In these protocols, data are exchanged between a sender (the root of the tree) and several receivers (the leaves) via network elements (the internal nodes). Messages flowing from sender to receivers can be multicasted, i.e., simultaneously sent to several successors. An examples ofsuch a multicast network protocol is PGM (Pragmatic General Multicast) [S+00].

Verifying a property ϕ of a parameterised system consists of checking that ϕ for all possible structures (in our case, for all possible trees). This may be difficult, and so a common approach is to first reduce the task to checking ϕ for a restricted class of structures (see for instance [EN95]). This is also the approach of this paper.

We first provide a general formal model for multicast networks. The only assumption is that messages can overtake other messages and can get lost, but cannot be duplicated. This is a reasonable assumption for protocols in which channels are just an abstraction for a routing mechanism that may send different messages—or different fragments of the same message—through various routes. We define a notion of simulation that preserves stuttering-invariant linear-time properties, i.e., if the simulating structure satisfies the property, then the simulated structure also satisfies it.

Equipped with this formal setting, we identify general sufficient conditions for a protocol P to be *collapsable*, meaning that a property holds for all instantiations $\mathcal{N}(P)$ of P if and only if it holds for the set of instantiations $\mathcal{N}^1(P)$

with at most one level of internal elements. We prove that for every instantiation T in $\mathcal{N}(P)$ there is an instantiation T' in $\mathcal{N}^1(P)$ that simulates T.[1] Then, all instantiations in $\mathcal{N}^1(P)$ satisfy ϕ if and only if all instantiations in $T \in \mathcal{N}(P)$ satisfy ϕ, because simulation preserves properties and $\mathcal{N}^1(P) \subseteq \mathcal{N}(P)$.

In particular, our conditions are satisfied by network elements that only perform 'forwarding', i.e., only forward messages down from the parent to all children, and up from some child to the parent. Hence, our result can be applied to telecommunication protocols that do not assume that router support can be used and that use a fixed distribution tree. We show that they are also satisfied by the PGM protocol, where network elements have a much richer functionality, which is used to make communication between the sender and receivers more reliable and efficient.

While the collapse of $\mathcal{N}(P)$ to $\mathcal{N}^1(P)$ removes the problem of dealing with different tree topologies, it still leaves us with an infinite number of possible instantiations. This cannot be avoided as long as more receivers can generate more behaviour. However, we prove that if the number of different messages that can circulate is finite and receivers and network elements can repeatedly send the same message upwards, then the verification task can be further reduced: Given a property ϕ, all instantiations of $\mathcal{N}^1(P)$ satisfy ϕ if and only if one single universal instantiation U, which depends on Φ. Again, we prove that if the number of messages in the PGM protocol is bounded, then the result can be applied, and the protocol has a universal instantiation.

The paper is structured as follows. In section 1 we introduce (a version of) the PGM protocol, in order to introduce network protocols and have a rich running example for our definitions and results. Section 2 contains our formal model. Our notions of property and simulation are given in Section 3. Section 4 presents the sufficient conditions for a network to be collapsable. This section is divided into two parts; the first part deals with the special case, in which network elements can only forward messages, and the second deals with the general case. The section also shows that the PGM protocol satisfies the conditions. Finally, section 5 presents the universal instance that can simulate any other instantiation of a collapsable protocol, assuming that the number of messages is finite. The paper contains two appendices. The first one compares our version of the PGM protocol with the more standard one. The second contains the proofs of the theorems.

1 The PGM protocol

The Pragmatic General Multicast (PGM) protocol is a reliable multicast protocol for the distribution of information from multiple senders to multiple receivers.

[1] Notice that we always speak of instantiations *of the same protocol*. Given a protocol P, one can always find another protocol P' such that every $T \in \mathcal{N}(P)$ is simulated by some $T' \in \mathcal{N}^1(P')$ by making the sender, receiver and internal processes more complicated.

It is designed in order to minimize loading of the network due to acknowledgment messages or retransmissions of lost packets, and has been presented to the Internet Engineering Task Force as an open reference specification. We consider the following abstract, untimed variant of the protocol for one sender. We have a tree of processes connected by bidirectional channels. Messages can get lost and can overtake each other (i.e., can be delivered in a different order than they are sent), but cannot be duplicated. The root of the tree is the sender, and the leaves the receivers. The other processes are internal *network elements*. The source multicasts a numbered sequence of data packets called odata(*nr*,*trl*) (for original data) within a transmit window; *nr* is the number of the package, and *trl* is the left-hand edge of the sender's window at the moment of sending it. Network elements forward these packets down the distribution tree. If a receiver detects that packet *nr* is missing from the sequence, it repeatedly sends a *primary negative acknowledgment* (pnak(*nr*)) to its parent, requesting a repair. Each network element that receives a pnak forwards it to its parent, and multicasts a nak-confirmation (ncf(*nr*)) to its children; it then keeps sending secondary nak (snak(*nr*)) to its parent (which are forwarded upwards, but do not generate confirmations) until it receives a *nak-confirmation* itself. When the source receives a pnak or snak it provides a repair (rdata(*nr*,*trl*)), which is multicasted downwards to the processes that requested them. There is a final feature called *nak-anticipation*: A receiver may receive a confirmation to a pnak sent by *another* receiver. Anticipating its own future need for a repair, it repeatedly sends snak's to its parent, until either the original data or the repair arrives. Notice that odata-, rdata- and ncf-messages travel *downwards* (from sender to receivers) while pnak- and snak-messages travel *upwards* (from receiver to sender).

Formally, the protocol is given by three agents describing the sender, the receivers, and the network elements, whose descriptions can be found in Table 1. Every agent has a set of *variables* and a set of (atomic) *transitions*. Transitions are guarded by either boolean expressions over the process variables or by the delivery of a message. In the initial state, all sets are empty, *odata*, *txw_trail* and *rxw_trail* are 0, and *WIN_SIZE* has some fixed value (window size).

Our version of the PGM protocol differs slightly from [S+00] in that we distinguish between primary and secondary nak messages. While our result also holds for the original version, our version allows for a generic proof, and it can simulate the original version except for a behaviour which is not desirable according to the specification: Our version is more economic in the number of nak messages sent than the original one. All these points are discussed in detail in the Appendix.

2 A formal model of network protocols

In this section we formalise the notions of tree networks, and of families of tree networks defined by a protocol.

Table 1. Agents of the PGM protocol

agent source
odata, *WIN_SIZE*, TXWTR: $\mathbb{N}$; *rec_nak*: set of $\mathbb{N}$
s1: in pnak(nr) $\vee$ in snak(nr) $\rightarrow$ out ncf(nr) downwards;
$txw\_trail < nr \rightarrow rec\_nak := \text{add}(rec\_nak, nr)$;
s2: is_in(nr, *rec_nak*) $\rightarrow$ $nr > txw\_trail \rightarrow$ out rdata(nr, *txw_trail*) downwards;
$rec\_nak := \text{remove}(rec\_nak, nr)$;
s3: length(*rec_nak*) $= 0 \rightarrow$ out odata(*odata*, *txw_trail*) downwards;
$odata := odata + 1$;
$odata + 1 > WIN\_SIZE + txw\_trail$
$\rightarrow txw\_trail := txw\_trail + WIN\_SIZE$;
endagent;

agent network_element
set_repair: set of $\mathbb{N}$; *set_interf*: set of ($\mathbb{N}$, channel_name)
e1: in pnak(nr) $\rightarrow$ $set\_repair := \text{add}(set\_repair, nr)$;
$set\_interf := \text{add}(set\_interf, (nr, c))$; [$c$ reception channel]
pnak $\notin set\_repair \rightarrow$ out pnak(nr) upwards;
out ncf(nr) downwards;
e2: in snak(nr) $\rightarrow$ $set\_interf := \text{add}(set\_interf, (nr, c))$; [$c$ reception channel]
out snak(nr) upwards;
e3: in rdata(nr, trl) $\rightarrow$ $set\_repair := \text{remove}(set\_repair, nr)$;
out rdata(nr,trl) to all channels c' s.t. $(nr, c') \in set\_interf$
$set\_interf := \text{remove}((nr, c'), set\_interf)$
e4: in ncf(nr) $\rightarrow$ $set\_repair := \text{remove}(set\_repair, nr)$;
e5: in odata(nr, trl) $\rightarrow$ out odata(nr, trl) downwards;
e6: is_in(nr, *set_repair*) $\rightarrow$ out snak(nr) upwards;
endagent;

agent receiver
rxw_trail, *set_nr*, *set_missing*: set of $\mathbb{N}$
r1: in odata(nr, trl) $\wedge$ $rxw\_trail < nr \rightarrow$ $rxw\_trail < trl \rightarrow rxw\_trail := trl$;
$set\_nr := \text{add}(set\_nr, nr)$;
for all $(rxw\_trail < i < nr \wedge i \notin set\_nr)$
$set\_missing := \text{add}(set\_missing, i)$;
$set\_missing := \text{remove}(set\_missing, nr)$;
$set\_smissing := \text{remove}(set\_smissing, nr)$
r2: in ncf(nr) $\wedge$ $rxw\_trail < nr \wedge nr \notin set\_nr \rightarrow$ $set\_smissing := \text{add}(set\_smissing, nr)$
for all $(rxw\_trail < i < nr \wedge i \notin set\_nr)$
$set\_smissing := \text{add}(set\_smissing, i)$;
r3: in rdata(nr, trl) $\wedge$ $rxw\_trail < nr \rightarrow$ $rxw\_trail < trl \rightarrow rxw\_trail := trl$;
$set\_nr := \text{add}(set\_nr, nr)$;
$set\_missing := \text{remove}(set\_missing, nr)$;
$set\_smissing := \text{remove}(set\_smissing, nr)$
r4: is_in(nr, *set_missing*) $\rightarrow$ $nr > rxw\_trail \rightarrow$ out pnak (nr) upwards;
$nr \leq rxw\_trail \rightarrow set\_missing := \text{remove}(set\_missing, nr)$;
r5: is_in(nr, *set_smissing*) $\rightarrow$ $nr > rxw\_trail \rightarrow$ out snak(nr) upwards;
$nr \leq rxw\_trail \rightarrow set\_smissing := \text{remove}(set\_smissing, nr)$;
endagent;

2.1 Messages, actions, events, and histories

Let M be a (possibly infinite) set of *messages*. We assume that M contains an 'empty' message, denoted by $\perp$. We assume that $M = M{\uparrow} \cup M{\downarrow} \cup \perp$ [2] where $M{\uparrow}$ and $M{\downarrow}$ are sets of *upward* and *downward* messages such that $M{\uparrow} \cap M{\downarrow} = \emptyset$. We model receiving or sending no message as receiving or sending the "empty" message $\perp$. An *abstract action*, or just an *action*, is a triple $(i, o_1, o_2) \neq (\perp, \perp, \perp)$ of messages such that $o_1 \in M{\uparrow} \cup \perp$ and $o_2 \in M{\downarrow} \cup \perp$. Intuitively, an action models receiving a message i, and sending a message o_1 upwards and a message o_2 downwards. A *trace* is a finite sequence of actions. We denote the set of all traces by TR.

Let Ch be a set of *downward channels*, and let $ch \notin Ch$ be an *upward channel*. A *concrete action* or *event* over ch, Ch is a fivetuple (i, c, o_1, o_2, C), where (i, o_1, o_2) is an abstract action, $c \in \{ch\} \cup Ch$, $C \subseteq Ch$, and moreover, either $c = ch$ and $i \in M{\downarrow} \cup \perp$, or $c \in Ch$ and $i \in M{\uparrow} \cup \perp$. (The intuition is that ch is the channel communicating with the parent, and Ch the channels communicating with the children.) An event corresponds to a process receiving message i through channel c, sending o_1 through the upward channel ch, and sending o_2 through a subset C of downward channels. We denote the set of all events by E. A *history* is a finite sequence of events. We denote the set of all histories by H.

Given an event $e = (i, c, o_1, o_2, C)$ over ch, Ch, we define the *action corresponding to* e as (i, o_1, o_2), and denote it by $a(e)$. Given a channel c', we define the c'*-action* corresponding to e, denoted by $c'(e)$, as the pair $(c'(i), c'(o_1, o_2))$ given by: (1) $c'(i) = i$ if $c' = c$ and $c'(i) = \perp$ otherwise, and (2) $c'(o_1, o_2) = o_1$ if $c' = ch$, $c'(o_1, o_2) = o_2$ if $c' \in C$, and $c'(o_1, o_2) = \perp$ otherwise.

Given a history $h = e_1 \dots e_n$, we define its *associated trace* as the sequence $tr(h) = a(e_1) \dots a(e_n)$. Given a channel c, at most one message is sent through c during an action, and hence the projection of history h onto channel c is a sequence $tr(h, c) = c(e_1) \dots c(e_n)$. We call such sequences *projected traces* and denote them by $TRproj$.

2.2 Agents and processes

In multicast protocols, like the PGM, agents are defined to work independently of the identity and number of their upward and downward channels, because the architecture of the network is not known *a priori*. Our notion of *agents* intends to be very general, while respecting this limitation.

An *agent* is a pair $A = (\rho, f)$, where $\rho\colon TR \times M \to 2^{M{\uparrow} \times M{\downarrow}}$ is the *input/output relation*, and $f\colon TRproj \times Act \times Bool \to Bool$ is the *filter*. Let us explain this definition. Intuitively, an agent A selects the events that can be executed as a function of the past history, and of the current input message i. After receiving i, the agent selects an event in two steps. First, it nondeterministically selects the messages o_1, o_2 to be sent upwards and downwards, respectively, as a function of the current trace tr of actions and the input message i. Formally, $(o_1, o_2) \in$

[2] Throughout the paper we identify $\perp$ and the set $\{\perp\}$. This should cause no confusion.

$\rho(tr, i)$. In the second step, the agent determines the subset of downward channels through which the message o_2 is sent. The agent examines each channel $c \in Ch$, and decides whether to send o_2 through it or not depending on the action $a = (i, o_1, o_2)$, the projection $tr(h, c)$ of h on channel c, and on whether the input i came via the channel c or not. Formally, o_2 is sent through c if $f(tr(h, c), a, b) = \mathit{true}$, where b is true iff i arrived through channel c.

A *process* is a triple $P = (ch, Ch, A)$, where A is an agent, Ch is a set of channels, and ch is a channel that does not belong to Ch. The set of *transitions* of the process P is the subset of $H \times E \times H$ containing the triples $(h, e, h \cdot e)$ (also denoted by $h \xrightarrow{e} h \cdot e$), such that e is an event that can be selected by A when h is the past history of the process.

Example: We formalise the input/output relation ρ_n of the network element agent of the PGM protocol in our framework. (The filter can be formalised analogously.)

$$\begin{array}{ll}
\rho_n(tr, \mathrm{rdata}(nr, trl)) = (\bot, \mathrm{rdata}(nr, trl)) & \rho_n(tr, \mathrm{pnak}(nr)) \\
 & \quad = (\mathrm{pnak}(nr), \mathrm{ncf}(nr)) \\
\rho_n(tr, \mathrm{odata}(nr, trl)) = (\bot, \mathrm{odata}(nr, trl)) & \rho_n(tr, \mathrm{snak}(nr)) = (\mathrm{snak}(nr), \bot) \\
\rho_n(tr, \bot) = (\mathrm{snak}(nr), \bot) & \rho_n(tr, \mathrm{ncf}(nr)) = (\bot, \bot) \\
\quad \text{if } nr \in \mathit{set_repair}(tr) &
\end{array}$$

Notice that the value of *set_repair* is indeed fully determined by tr; $nr \in \mathit{set\_repair}(tr)$ iff tr contains the action $(\mathrm{pnak}(nr), \mathrm{pnak}(nr), \mathrm{ncf}(nr))$ (which adds nr to *set_repair*) and no later occurrence of the actions $(\mathrm{ncf}(nr), \bot, \mathrm{ncf}(nr))$ or $(\mathrm{rdata}(nr), \bot, \mathrm{rdata}(nr))$ (which remove nr). This is to say *set_repair* provides an abstract view on a trace sufficient to define $\rho_n(PGM)$.

2.3 Tree networks and protocols

Loosely speaking, a tree network is a network of processes with a tree topology; every process is connected to its parent and children by bidirectional channels.

Syntax A finite tree T is a set of nodes together with a partial order $\leq_T$ satisfying the usual tree condition: if $\mathbf{n}_1 \leq_T \mathbf{n}$ and $\mathbf{n}_2 \leq_T \mathbf{n}$, then $\mathbf{n}_1 \leq_T \mathbf{n}_2$ or $\mathbf{n}_2 \leq_T \mathbf{n}_1$. We denote the child relation by $\prec$ (i.e., $\mathbf{n} \prec \mathbf{n}'$ if $\mathbf{n} <_T \mathbf{n}'$ and there is no $\mathbf{n}''$ such that $\mathbf{n} <_T \mathbf{n}'' <_T \mathbf{n}'$). We write $p(n)$ for the *parent* of $\mathbf{n}$, i.e., for the unique $\mathbf{n}'$ such that $\mathbf{n}' \prec \mathbf{n}$ or for the symbol "$\_$" if there is not such $\mathbf{n}'$. If $\mathbf{n} \prec \mathbf{n}'$, then we call the pair $[\mathbf{n}, \mathbf{n}']$ a *channel*. We define the sets of downward channels of $\mathbf{n}$ in T as $Ch(\mathbf{n}, T) = \{[\mathbf{n}, \mathbf{n}'] \mid \mathbf{n} \prec \mathbf{n}'\}$, and let $ch(\mathbf{n}, T) = [p(\mathbf{n}), \mathbf{n}]$. If $\mathbf{n}$ is the root, then the channel $[\_, \mathbf{n}]$ is considered to connect to the environment. Throughout this paper, for simplicity we assume that the sender only uses its downward channels and receivers only use their upward channel. If no confusion is possible, we shorten $ch(\mathbf{n}, T)$ and $Ch(\mathbf{n}, T)$ to $ch(\mathbf{n})$ and $Ch(\mathbf{n})$, respectively.

A *tree network* is a pair (T, A), where T is a finite tree, and A is a mapping that associates to each node $\mathbf{n} \in N$ an agent $A(\mathbf{n})$. The mapping *Proc* associates to each node $\mathbf{n}$ the process $\mathit{Proc}(\mathbf{n}) = (ch(\mathbf{n}), Ch(\mathbf{n}), A(\mathbf{n}))$.

We call the root of the tree the *sender* of the network, and denote it by $\mathbf{s}$. Maximal nodes w.r.t. $\leq_T$ are called *receivers*. All other nodes are called *internal*.

A tree network (T, A) is *homogeneous* if $A(\mathbf{n}) = A(\mathbf{n}')$ for every two internal nodes $\mathbf{n}, \mathbf{n}'$, and $A(\mathbf{r}) = A(\mathbf{r}')$ for every two receivers $\mathbf{r}$ and $\mathbf{r}'$. A *protocol* P is a set of three agents A_s, A_n, A_r. A protocol P defines a family $\mathcal{N}(P)$ of homogeneous tree networks, namely the tree networks (T, A) satisfying $A(\mathbf{s}) = A_s$, $A(\mathbf{n}) = A_n$ for all internal elements $\mathbf{n}$, and $A(\mathbf{r}) = A_r$ for all receivers $\mathbf{r}$.

Semantics A *network event* of a tree network (T, A) is a pair $(\mathbf{n}, e)$, where $\mathbf{n}$ is a node of T and e is an event of $Proc(\mathbf{n})$. Intuitively, a network event models that the process $Proc(\mathbf{n})$ executes the event e. We denote the set of all network events by Nev. A *network history* is a finite sequence of network events. Given a network history $nh = (\mathbf{n}_1, e_1) \dots (\mathbf{n}_n, e_n)$ and a node $\mathbf{n}$, we define by $nh(\mathbf{n})$ the projection of nh onto the events executed by $\mathbf{n}$. Notice that $nh(\mathbf{n})$ is a history of the process $Proc(\mathbf{n})$.

Since messages can overtake other messages, a channel behaves like a multiset, which only retains the multiplicity of each message, but not their order. Loss of a message need not be modelled explicitely, because it can be simulated by never taking the message from the channel. The multiset of messages that are waiting for delivery in the channel c after the execution of nh is denoted by $M(nh, c)$. Formally, $M(nh, [\mathbf{n}, \mathbf{n}'])$ is the multiset of messages sent through channel c by $Proc(\mathbf{n})$ and $Proc(\mathbf{n}')$ during the network history nh, minus the multiset of messages received through c by the same processes, also during nh

A triple $(nh, (\mathbf{n}, e), nh')$, where nh, nh' are network histories, and $(\mathbf{n}, e)$ is a network event, is a transition of (T, A) if there is a transition $h \xrightarrow{e} h'$ of $Proc(\mathbf{n})$ satisfying the following conditions:

(1) $nh(\mathbf{n}) = h$, $nh'(\mathbf{n}) = h'$, and $nh(\mathbf{n}') = nh(\mathbf{n})$ for every $\mathbf{n}' \neq \mathbf{n}$.
(2) Let i be the message received by $Proc(\mathbf{n})$ in e, and let c be the channel through which i arrived. Then, either $i = \perp$, or $i \in M(nh, c)$. I.e., the message received by $Proc(\mathbf{n})$ in the event e was either empty or it was waiting for delivery in the channel c.

If $(nh, (\mathbf{n}, e), nh')$ is a transition of (T, A), then we write $nh \xrightarrow{e} nh'$. We write $nh \xrightarrow[\mathbf{n}_1 \dots \mathbf{n}_n]{e_1 \dots e_n} nh'$ if there are transitions $nh \xrightarrow[\mathbf{n}_1]{e_1} nh_1 \dots nh_{n-1} \xrightarrow[\mathbf{n}_n]{e_n} nh'$. W.l.o.g, we assume that every network history has at least one successor (if not, just add self-looping transitions with some special label).

Fair executions In our framework, a state of a tree network is given by a network history. An *execution* of T is an infinite sequence $\pi = nh_0 nh_1 nh_2 \dots$ of network histories (i.e., of states) such that for every $i \geq 0$ there is a network event $(e_i, \mathbf{n}_i)$ satisfying $nh_i \xrightarrow[\mathbf{n}_i]{e_i} nh_{i+1}$. An execution is *fair* with respect to a subset $\tilde{T} \subseteq T$ if for every node $\mathbf{n} \in \tilde{T}$ in the network, $\mathbf{n}_i = \mathbf{n}$ for infinitely many $i \geq 0$, i.e., if every node in $\tilde{T}$ executes infinitely many events.

3 Fair stuttering simulations for tree networks

Fix a protocol P. We are interested in properties that concern only the behaviours of the sender and the receivers, since these are the 'visible' elements of the protocol. So we consider properties Φ of the form $\Phi = \forall\mu\forall\sigma\phi$, where μ is a tuple of *message variables*, σ is a tuple of *receiver variables*, and ϕ is a stuttering invariant LTL temporal logic formula [Lam83]. The atomic propositions of ϕ can be indexed by the variables of μ, σ.

Example. Informally, the main property that the PGM protocol should satisfy is "for every receiver $\mathbf{r}$ and for every message m sent by the sender, eventually one of the following two holds: $\mathbf{r}$ receives m, or $\mathbf{r}$ knows that m is lost, and that the sender is not going to resend it in the future". The second possibility means that $\mathbf{r}$ finds out that the lower end of the sender's retransmission window (given by *txw_trail*) is larger than the number of m.

Given a tree network $T \in \mathcal{N}(P)$, we interpret atomic propositions over sets of network histories of T.[3] We say that T satisfies Φ if for all valuations *val* of the variables μ, σ, and for all executions π of T that are fair with respect to $\mathbf{s}$ and to all receivers in the image of *val*, $\pi \models \phi[\mu := val(\mu),\ \sigma := val(\sigma)]$.

So we ignore executions in which the specified receivers or the sender are not scheduled infinitely often, or are 'cut off from the network' after a certain time point, i.e. do no longer receive messages due, say, to a connection breakdown. But we do not require that all processes in the network execute infinitely often, so that the property still has to hold if some processes not mentioned in the property are cut off the network. We say that P satisfies Φ if T satisfies Φ for all $T \in \mathcal{N}(P)$, i.e., P satisfies a property if all the homogeneous tree networks of $\mathcal{N}(P)$ satisfy it.

A simulation of T by T' preserves a property Φ if $T' \models \Phi$ implies $T \models \Phi$. Our goal is to prove that all the tree networks of $\mathcal{N}(P)$ can be simulated by those in a subset $\mathcal{N}'(P) \subseteq \mathcal{N}(P)$, according to a notion of simulation that preserves properties. Once this is achieved, we can prove that P satisfies Φ by showing that the networks of $\mathcal{N}'(P)$ satisfy Φ.

We now define a stuttering version of simulation, similar to stuttering bisimulation [BCG88]. Since in the simulation we consider all actions of the simulated network T will have a corresponding sequence of actions in the simulating network T', we do not have to consider stuttering in T, which simplifies the definition. Since we only require fair paths to satisfy a property, we adapt the definition accordingly, like in fair simulation as introduced in [GL94], the coarsest simulation that preserves fair-ACTL*.

Given two networks T and T', let Im be a mapping that assigns to each receiver $\mathbf{r}$ of T a receiver $Im(\mathbf{r})$ of T', called the *image* of $\mathbf{r}$. Let $Match(nh, nh')$ be the relation between histories of T and T' given by $tr(nh(\mathbf{s})) = tr(nh'(\mathbf{s}')))$, where $\mathbf{s}$ and $\mathbf{s}'$ are the senders of T and T', respectively, and $tr(nh(\mathbf{r})) = tr(nh'(Im(\mathbf{r})))$ for all receivers $\mathbf{r}$.

[3] Once P is fixed, all networks (T, A) of $\mathcal{N}(P)$ share the same agent function A. So we shorten (T, A) to T.

Definition 1
A *fair stuttering simulation of T by T' with respect to Im* is a relation $R \subseteq NH \times NH'$ such that $R(nh, nh')$ implies:

- $Match(nh, nh')$.
- For every subset $\tilde{T}$ of T consisting of the sender and receivers, and for every execution π of T starting at nh which is fair with respect to $\tilde{T}$ there is an execution π' of T' starting at nh', fair with respect to $Im[\tilde{T}]$, and an increasing mapping $\sigma\colon \mathbb{N} \longrightarrow \mathbb{N}$ with $\sigma(0) = 0$ such that (a) for all $n \geq 0$, $R(\pi(n), \pi'(\sigma(n))$, and (b) for all $\sigma(n) < j \leq \sigma(n+1)$, $Match(\pi(n+1), \pi'(j))$ and (c) the values $(\sigma(i))_{i \in \mathbb{N}}$ increase over all bounds.

We say that T is simulated by T' (with respect to Im) if there is a mapping Im and a fair stuttering simulation R of T by T' with respect to Im such that $R(nh_0, nh'_0)$, where nh_0 and nh'_0 are the empty network histories of T and T'.

The following theorem describes properties that are preserved by fair stuttering simulation. It follows easily from the fact that for every execution π in T that is fair with respect to $\tilde{T}$, the violating valuation for receiver variables in Φ, there is an execution π' in T', which is fair with respect to $Im[\tilde{T}]$, such that the states of π and π' pointwise satisfy $Match$, respective to the fact that π' can contain stuttering steps (insertion of states s such that s and its predecessor satisfy $Match$). This implies that π' satisfies the same stuttering-invariant LTL properties as π.

Theorem 1 *Let T and T' be tree networks such that T is simulated by T' with respect to Im. Let $\Phi = \forall\ \mu\forall\ \sigma\phi$ be a formula such that ϕ is stuttering-invariant, and for all atomic propositions p, and for all network histories nh and nh', $Match(nh, nh')$ implies $nh \models p \Longleftrightarrow nh' \models p$. Then, $T' \models \Phi$ implies $T \models \Phi$.*

Finally, we have the result we were looking for:

Theorem 2 *Let $\mathcal{N}' \subseteq \mathcal{N}(P)$ such that each $T \in \mathcal{N}(P)$ is simulated by some $T' \in \mathcal{N}'$, and let Φ as is Theorem 1. Then $\mathcal{N}' \models \Phi$ implies $\mathcal{N}(P) \models \Phi$.*

4 Collapsable tree networks

In this section, we explore conditions on protocols that allow to flatten the tree hierarchy, i.e. that any tree network can be simulated by one with at most one layer of internal elements. So the subset of tree networks with only one (without a) level of internal element can be used, in the sense of Theorem 2 for properties that speak about the sender and receivers only.

Throughout this section we fix a protocol P consisting of agents (ρ_s, f_s), (ρ_n, f_n) and (ρ_r, f_r). Let $\mathcal{N}^0(P)$ be the set of tree networks in $\mathcal{N}(P)$ without internal elements, and let $\mathcal{N}^1(P)$ be the set of tree networks in $\mathcal{N}(P)$ with only one layer of internal elements. For simplicity, we only consider sender agents that always sends downward messages through all their downward channels.

We are interested in protocols like the PGM, where the sender exchanges messages with the receivers, and the primary functionality of internal elements is to *forward* these messages. Transition $\mathbf{e_5}$ of the PGM provides an example: An odata-message coming from above is simply passed on to all successors. However, in order to deal with lost messages and to improve the efficiency, internal elements can also perform other tasks. First, they can *filter* downward messages: Instead of sending a message to all its successors, they select a subset of them as recipients. An example is transition $\mathbf{e_3}$: rdata(nr,trl) is forwarded to the channels through which pnak(nr) or snak(nr) has been received but rdata(nr,trl) has not been sent afterwards. Moreover, besides forwarding with or without filtering, internal elements can also *generate* messages. Transition $\mathbf{e_6}$ is an example of 'spontaneous' generation, while in transition $\mathbf{e_1}$ reception of pnak(nr) triggers the generation of ncf(nr). The generation of messages can in general depend on the internal state of the component, i.e., its history. For example, generation of snak(nr) depends on whether nr is contained in the set *set_repair*.

We present our main result in two steps. First, in section 4.1 we consider internal agents that have only forwarding transitions (but possibly with filtering). We then consider protocols in which network elements can also generate messages. It is easy to see that for general protocols of this form there is no $n \geq 0$ such that $\mathcal{N}^n(P)$ simulates $\mathcal{N}(P)$. We identify a class of protocols P for which $\mathcal{N}(P)$ can be simulated by $\mathcal{N}^1(P)$, and show that the PGM belongs to it.

4.1 Forwarding agents

Intuitively, a forwarding agent is an agent that forwards incoming upward messages through its upward channel, and multicasts incoming downward messages through some of its downward channels. We also allow the agent to 'swallow' messages.

Definition 2
Formally, an agent (ρ, f) is a *forwarding agent* if for every trace tr and every message $i \in M$:

- If $i \in M\uparrow$, then $\rho(tr, i) \subseteq \{(i, \bot), (\bot, \bot)\}$;
- if $i \in M\downarrow$, then $\rho(tr, i) \subseteq \{(\bot, i), (\bot, \bot)\}$; and
- if $i = \bot$, then $\rho(tr, \bot) = \{(\bot, \bot)\}$.

A *forwarding protocol* is a protocol whose network element agent (but not necessarily its sender or receiver agents) is forwarding.

Let T be a tree network. We define the tree $\underline{T}$ as follows: $\underline{T}$ contains a sender $\underline{\mathbf{s}}$ and a receiver $\underline{\mathbf{r}}$ for every receiver $\mathbf{r}$ of T, but no internal elements; $\leq_{\underline{T}}$ is the projection of $\leq_T$ onto $\underline{T}$. So in $\underline{T}$, every receiver is a child of the sender.

In the Appendix, we define a relation $nh \lhd \underline{nh}$ between network histories nh and $\underline{nh}$, and show that $\lhd$ is a fair stuttering simulation of T by $\underline{T}$. We get as corollary:

Theorem 3 *If P is a forwarding protocol and Φ is like in Theorem 1, then $\mathcal{N}^0(P) \models \Phi$ if and only if $\mathcal{N}(P) \models \Phi$.*

4.2 Forwarding and generating agents

In this section, we consider internal agents that can not only forward but also generate messages. Clearly, we can no longer expect $\mathcal{N}(P)$ to be simulated by $\mathcal{N}^0(P)$, unless the messages generated by the internal elements have no effects whatsoever. We give conditions under which $\mathcal{N}(P)$ can be simulated by $\mathcal{N}^1(P)$. We then show that these conditions are satisfied by the PGM protocol.

Let $M_{gen} \subseteq M$ be the set of messages that can be generated by internal network elements. (Notice that the sender and the receivers can also generate messages, but these do not have to be in M_{gen}.)

Definition 3
A *forwarding and generating agent* (*f&g agent* for short) is an agent (ρ, f) such that for every trace $tr \in TR$ and for every message $i \in M$, the following conditions hold:

- If $i \in M\uparrow$ and $(o_1, o_2) \in \rho(tr, i)$, then $o_1 \in \{i, \perp\}$ and $o_2 \in M_{gen} \cup \{\perp\}$;
- if $i \in M\downarrow$ and $(o_1, o_2) \in \rho(tr, i)$, then $o_1 \in M_{gen} \cup \{\perp\}$ and $o_2 \in \{i, \perp\}$; and
- if $i = \perp$ and $(o_1, o_2) \in \rho(tr, \perp)$, then $o_1, o_2 \in M_{gen} \cup \{\perp\}$.

A *f&g protocol* is a protocol with a $f\&g$ internal element agent.

These definition allows messages to be generated as 'side-effects' of forwarding other messages, or spontaneously, i.e. without receiving input. In the PGM protocol, ncf-messages are generated as side-effects, while snak-messages are generated spontaneously.

Conditions on the protocol We define the class of simple protocols, for which we prove that $\mathcal{N}^1(P)$ simulates $\mathcal{N}(P)$. This requires some preliminaries.

Receiving a message has two effects. The first, immediate effect is that some messages are sent. The second effect is that the internal state of the process (given in our model by the history) changes. This change may enable the process to send messages, but they may be sent at a later stage. The change of state may also disable the emission of messages. An example of an enabling effect is given by transitions **e1** and **e6** of the PGM: Transition **e1** adds *nr* to *set_repair*, which enables the process to send snak through transition **e6**.

We need a definition implying the following intuitive idea: Receiving a message i in a subset M' may have arbitrary disabling effects, but it can only enable upward forwarding of i, or generation of messages in another subset M''.

Definition 4
Let (ρ, f) be an agent, tr be a trace of it, $M' \subseteq M$, and $M'' \subseteq M\uparrow$. Let $Rem(tr, M', M'')$ be the set of all traces resulting from tr by removing arbitrarily many actions of the form (i, o_1, o_2) such that either $i \in M'$ or $i = o_2 = \perp$ and $o_1 \in M''$. We say that M' *only enables* M'' in the agent (ρ, f) if for every trace tr and every $tr' \in Rem(tr, M', M'')$ the following conditions hold:

(1) If $i \in M' \cap M{\uparrow}$ and $(o_1, o_2) \in \rho(tr, i)$, then $o_1 \in M'' \cup \{i, \bot\}$ and $o_2 = \bot$;
if $i \in M' \cap M{\downarrow}$ and $(o_1, o_2) \in \rho(tr, i)$, then $o_1 \in M'' \cup \bot$ and $o_2 = \bot$;
(2) if $i \in M \setminus \bot$, then $\rho(tr, i) \subseteq \rho(tr', i)$; and
$\rho(tr, \bot) \setminus \rho(tr', \bot) \subseteq \{(o, \bot) \mid o \in M''\}$.

Let us see why this definition captures the intuition above. Condition (1) expresses that the *immediate* effect of receiving $i \in M'$ can only be forwarding i up, or the generation of an upward message $o_1 \in M''$. Condition (2) expresses that the *long-term* effect of receiving i can only the be generation of upward messages in M'', let us see why. Suppose first that $M'' = \emptyset$. Then (2) implies that receiving messages in M' can only have disabling effects: Whatever we can do after receiving the messages (given by $\rho(tr, i)$), we can also do without receiving them (given by $\rho(tr', i)$). Now consider the general case. Since $\rho(tr, \bot) \setminus \rho(tr', \bot) \subseteq \{(o, \bot) \mid o \in M''\}$, receiving messages in i may now enable actions $(\bot, o, \bot)$ for $o \in M''$. And since in $Rem(tr, M, M'')$ we now allow to remove actions $(\bot, o, \bot)$ for $o \in M''$, we make sure that such actions themselves only enable actions of the same kind.

We are almost ready to present the definition of simple protocols. Let M_{fil} be the subset of $M{\downarrow}$ containing the messages that can be filtered when being forward downwards. More precisely, for all actions a in which the downwards output message belongs to $M{\downarrow} \setminus M_{fil}$, there is no filtering, i.e., $f(tr, a, b) = true$ for all traces tr and all boolean values b. We now define:

Definition 5
Let $P = (A_s, A_n, A_r)$ be a f&g protocol. P is *simple* if:

(a) In A_n (the network element agent), M_{gen} only enables $\emptyset$, and
(b) in A_r (the receiver agent), $M_{fil} \cup M_{gen}$ only enables $M_{gen} \cap M{\uparrow}$.

If condition (a) is dropped, the following scenario becomes possible: A network element $\mathbf{n}$ spontaneously generates a message o_1 and sends it upwards to its parent. The parent forwards it up and, in the next step, generates itself another copy of o_1; all predecessors of $\mathbf{n}$ behave in the same way. This produces a cascade of upward messages, and the number of o_1's received by the sender depends on the position of $\mathbf{n}$ in the tree structure. It is easy to see that this makes it impossible to simulate arbitrary structures by structures with at most one layer of internal elements.

If condition (b) is dropped, the following scenario becomes possible: A downward message $m \in M_{fil}$ is filtered out by a network element $\mathbf{n}$, i.e., $\mathbf{n}$ does not forward m down to any of its successors. The receivers that get m react by sending upwards a message $o_1 \notin M_{gen}$. Network elements that receive o_1 forward it up and, in the next step, generate a message $o_1' \in M_{gen}$ and send it upwards. Again, the number of o_1' messages received by the sender depends on the position of $\mathbf{n}$ in the tree structure.

The simulation We show that $\mathcal{N}^1(P)$ simulates $\mathcal{N}(P)$ if P is simple. Let T be a tree network. We define the tree $\underline{T}$ as follows: $\underline{T}$ contains a sender $\underline{\mathbf{s}}$. As network

elements can generate messages, the simulating tree network now has to contain a network element $\underline{\mathbf{n}}$ for every network element $\mathbf{n}$ in T and $\underline{\mathbf{n}}$ is a child of $\underline{\mathbf{s}}$. Moreover, a network element $\underline{\mathbf{n}}$ has to have the same receivers below itself as $\mathbf{n}$ does, because the actions of all these receivers can affect $\mathbf{n}$. So for every $\mathbf{r} > \mathbf{n}$, $\underline{\mathbf{n}}$ needs to have a child $\underline{\mathbf{r}}(\mathbf{n})$ that acts like $\mathbf{r}$. Thereby, we assume for simplicity that all receivers have an internal element as parent.

Figure 1 displays an example tree T and its flattening $\underline{T}$. In the Appendix, we define a simulation relation $nh \lhd \underline{nh}$ between network histories nh of T and $\underline{nh}$ of $\underline{T}$. The receiver $\underline{\mathbf{r}}(p(\mathbf{r}))$ exactly simulates $\underline{\mathbf{r}}$, and so we abbreviate $\underline{\mathbf{r}}(p(\mathbf{r}))$ by $\underline{\mathbf{r}}$. The other copies $\underline{\mathbf{r}}(\mathbf{n})$ are guaranteed to be able to execute the same actions as $\mathbf{r}$ except actions (i, o_1, o_2) such that $i \in M_{fil} \cup M_{gen}$ or actions $(\bot, o, \bot)$ such that $o \in M_{gen}$.

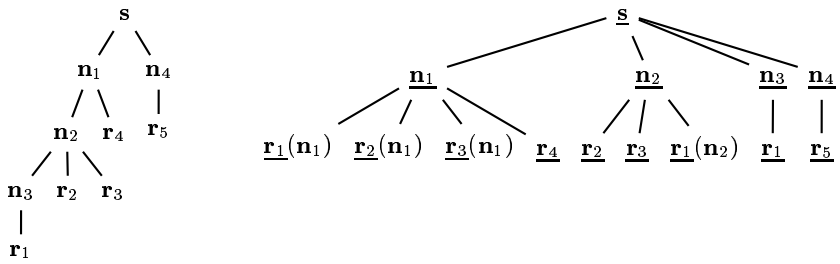

Fig. 1. Example of a tree T and its flattening $\underline{T}$

We obtain that for simple f&g protocols, it suffices to consider trees with only one level of internal elements:

Theorem 4 *If P is a simple f&g protocol and Φ is as in Theorem 1 then $\mathcal{N}^1(P) \models \Phi$ holds if and only if $\mathcal{N}(P) \models \Phi$.*

The PGM protocol We sketch why the PGM example as presented in Section 1 is simple. In the PGM protocol, M_{gen} consists of the messages of form snak(nr) and ncf(nr), and M_{fil} consists of the messages of form rdata(nr, trl). Let us show that the protocol satisfies the conditions of Definition 5.
First, the network element agent (ρ_n, f_n) of page 6 is f&g by definition. For (a): When forwarding snak(nr), only a message in M_{gen} (namely snak(nr)) is sent upwards, and nothing downwards, and on reception of ncf(nr) nothing is sent. This shows that condition (1) of Definition 4 holds. In order to see that messages in M_{gen} do not have long-term enabling effects (condition (2)), first note that ρ_n can be defined in terms of the variable *set_repair*, as explained on 6.[4] Reception of snak(nr) does not change *set_repair*, and on reception of ncf(nr), nr is removed from *set_repair*: So if $tr' \in Rem(tr, M_{gen}, \emptyset)$, i.e., if tr' is obtained by removing actions with input snak(nr) or ncf(nr) then in tr' the same actions as in tr (and

[4] The variable *set_interf* is irrelevant for ρ_n, it only affects the filter function f_n.

possibly more) are enabled. For (b): On reception of rdata(nr,trl) or ncf(nr), nothing is sent (condition (1)). In order to see that messages in M_{gen} do not have long-term enabling effects, notice first that ρ_n can be defined in terms of the variables of the agent. Since reception of rdata(nr,trl) adds nr to *set_nr* and removes it from *set_missing* and *set_smissing*, it only disables actions. Reception of ncf(nr) only enables actions of the form $(\perp, \text{snak}(nr'), \perp)$. Finally, sending of $(\perp, \text{snak}(nr'), \perp)$ does not change the variables of the receiver, and so it does not enable or disable actions. So the only actions that can be enabled in tr but not in $tr' \in Rem(tr, M_{fil} \cup M_{gen}, M_{gen} \cap M\uparrow)$ are of the form $(\perp, \text{snak}(nr'), \perp)$.

5 Reduction of one-layer tree networks

Given a simple protocol P and a property Φ, we have shown that checking Φ reduces to proving that $\mathcal{N}^1(P) \models \Phi$. We now introduce the class of *iteration protocols*, and show that for them $\mathcal{N}^1(P) \models \Phi$ reduces to proving $U_k \models \Phi$, where U_k is a particular instantiation of $\mathcal{N}^1(P)$ that depends on the number k of receiver variables used in Φ. Combining the two results we have that a simple iteration protocol P satisfies Φ if and only if $U_k \models \Phi$.

A trace tr is *reachable* by an agent (ρ, f) if it is empty or if $tr = tr' \cdot (i, o_1, o_2)$ where tr' is reachable and $(o_1, o_2) \in \rho(tr', i)$. An agent *can resend upward messages* if for every reachable trace $tr \cdot (i, o_1, o_2)$ and for every $n \geq 1$, the trace $tr \cdot (i, o_1, o_2) \cdot (\perp, o_1, \perp)^n$ is also reachable.

Definition 6
A protocol is an *iteration protocol* if $M\uparrow$ is a finite set, and both the receiver and the network element agents can resend upward messages.

Definition 7
Let P be an iteration protocol P, and let u and g be the sizes of $M\uparrow$ and $M\uparrow \cap M_{gen}$, respectively. For each $n \geq 1$, we define the *universal instance* U_k of $\mathcal{N}^1(P)$ as follows: The sender has $k + g + u$ children, all of them network elements; each of the first $k + u$ network elements has $u + 1$ children; and each of the other g network elements has u children. For $1 \leq i \leq k$, we denote the first child of the i-th network element by $\underline{\mathbf{r}}_i$. We also denote the tuple $(\underline{\mathbf{r}}_1, \ldots, \underline{\mathbf{r}}_k)$ by $\underline{\mathbf{R}}$.

In order to prove that U_k can simulate the behaviour of instances T in $\mathcal{N}^1(P)$, we show that fair executions of T are *stuttering-included* in the fair executions of U_k in the following sense. Let Im be a function mapping any k-sized subset $\tilde{T}$ of receivers of T to receivers of U_k. Let $Match(nh, nh')$ hold if $tr(nh)(\mathbf{r}) = trl(nh'(Im(\mathbf{r})))$ for all $\mathbf{r} \in \tilde{T}$. The requirement is that for any execution π which is fair with respect to $\tilde{T}$, there is an execution π' of U_k, fair with respect to $Im[\tilde{T}]$, and an increasing mapping $\sigma : \mathbb{N} \to \mathbb{N}$ with $\sigma(0) = 0$, such that (a) $Match(\pi(i), \pi'(\sigma(i)))$ holds for all i, and (b) for all $\sigma(i) < j \leq \sigma(i+1)$, $Match(\pi(i+1), \pi'(j))$ and (c) the values $(\sigma(i))_{i \in \mathbb{N}}$ increase over all bounds.

Intuitively, u-many receivers can mimic the behaviour of any set R of receivers as follows: For every $m \in M\uparrow$ there is a receiver $\mathbf{r}(m) \in R$ that first outputs m. We simulate all transitions of $\mathbf{r}(m)$ until the first output of m, and switch to iterating $(\bot, m, \bot)$ afterwards as required, and so can simulate all actions of receivers in R. By using this observation, we obtain (proof in the Appendix):

Theorem 5 *Let P be a simple f&g iteration protocol and $T \in \mathcal{N}^1(P)$. The fair executions of T are stuttering-included in the fair executions of U_k with respect to the mapping Im given by $Im(\mathbf{r}_i) = \underline{\mathbf{r}}_i$ for every receiver $\mathbf{r}_i$ of $\mathbf{R}$.*

Analogously to Theorem 1, stuttering-inclusion of fair executions implies that any formula Φ as in Theorem 1 with k receiver variables that holds for U_k also holds for T. (Note that a set of valuations of receiver variables contains at most k receivers.) As U_k is in $\mathcal{N}^1(P)$, we obtain:

Theorem 6 *Let P be a simple f&g iteration protocol, and let Φ be a property $\Phi = \forall\mu\forall\sigma\phi$ as in Theorem 1 such that σ is a tuple of n receiver variables. Then, $\mathcal{N}^1(P) \models \Phi$ if and only if $U_n \models \Phi[\sigma := \underline{\mathbf{R}}]$ and so, by Theorem 4, P satisfies Φ if and only if $U_n \models \Phi[\sigma := \underline{\mathbf{R}}]$.*

If in the PGM we bound the number of possible message numbers, i.e., instead of $nr \in \mathbb{N}$ we say $nr \in [1..n]$ for some number n, then we obtain a simple iteration protocol: The conditions of Definition 6 hold because of transitions **e6**, **r4** and **r5**.

6 Conclusions and related work

We have provided a general formal model of multicast network protocols. and have proved a general theorem showing that for a simple class of protocols, the verification problem reduces to the analysis of instantiations with at most one layer of internal elements between sender and receivers. For a smaller class we have also proved that the verification reduces to the analysis of one single instantiation. Protocols whose internal elements just forward messages, like the fit easily in our class. In fact, we have shown that the PGM protocol, whose internal elements exhibit a far more complicated behaviour, also fit in it.

As future work, we would like to explore whether our results can also be used for protocols that use local error recovery, which is a possible extension of the PGM protocol, and whether our approach can be extended to the analysis of timed protocols.

Related work. Some work on regular model-checking has addressed the problem of automatically verifying systems with a parameterised tree structure [BT02]. However, these techniques still seem far from being able to attack systems of the complexity of the PGM. There are also some papers on the analysis of the PGM protocol. However, so far they have concentrated on analysing the behaviour of a fixed instance, and so this work has a different nature to the work carried out

here. In [BBP02], the timed behaviour of a small instance of a simplified model is analysed. In the ADVANCE project, the untimed behaviour of a system that can simulate the universal instance U_1 has been studied. By our results this system can simulate *any* instance with respect to the main property of the protocol, since this property only involves one receiver variable. Unfortunately, at the time of writing this paper this instance is still out of the reach of the automatic tools.

Acknowledgements

This work has been supported by the FET Project ADVANCE, contract No IST-1999-29082.

References

[BBP02] Bérard, B., Bouyer, P. and Petit, A. *Analysing the PGM protocol with UPPAAL.* In: *2nd Workshop on Real-Time Tools.* Dep. Information Technology, Uppsala Univ., 2002, Tech. Report 2002-025.

[BCG88] Browne, M. C., Clarke, E. and Grumberg, O. *Characterizing finite Kripke structures in propositional temporal logic.* Theoretical Computer Science, 59: 115–131, 1988.

[BT02] Bouajjani, A. and Touili, T. *Extrapolating tree transformations.* In: *Proc. 14th Intl. Conf. on Comp. Aided Verif.* 2002, LNCS 2404.

[EN95] Emerson, E. A. and Namjoshi, K. S. *Reasoning about rings.* In: *Proc. 22th ACM Conf. on Principles of Programming Languages.* 1995.

[GL94] Grumberg, O. and Long, D. E. *Model checking and modular verification.* TOPLAS, 16(3): 843–871, 1994.

[Lam83] Lamport, L. *What good is temporal logic?* In: *Proc. IFIP 9th World Computer Congress.* 1983.

[Mai02] Maidl, M. *Simple representative instantiations for the PGM protocol*, 2002. Available at http://www.dcs.ed.ac.uk/monika.

[S+00] Speakman, T. et al. *PGM reliable transport protocol specification*, 2000. RFC 3208 (experimental) of the IETF, available at: http://www.ietf.org/rfc.html.

Appendix

Discussion of our PGM model

Our version of the PGM differs from that of [S+00] in the treatment of negative acknowledments. We claim that our version gets rid of some undesirable behaviours, admits a simpler collapsability proof, and satisfies the main property of the protocol if and only if the version of [S+00] does. We now substantiate these claims.

The original PGM specification [S+00], shown in Table 2, exhibits only one kind of negative acknowledgement: nak.. Our specification contains two kinds: pnaks and snaks. A process **n** generates a snak(nr) message when it (or its parent process if **n** is a receiver) has received the message pnak(nr) before (or a

message pnak(*nr'*) for some $nr' \geq nr$ if $\mathbf{n}$ is a receiver). The difference between pnaks and snaks is that forwarding of snaks does not add *nr* to *set_repair*, and does not generate any confirmation downwards.

The only behaviour of the specification of [S$^+$00] which is observable by senders or receivers and cannot be produced in our specification is the following: In both [S$^+$00] and our specification, the following behaviour is possible: (1) A network element $\mathbf{n}$ sends a nak(nr) to its parent $p(\mathbf{n})$; (2) $p(\mathbf{n})$ forwards a rdata(nr, trl) message down to $\mathbf{n}$; (3) $\mathbf{n}$ receives the nak(nr) message sent in (1). In [S$^+$00], $p(n)$ can now (4) generate a nak(nr) message, even though a repair rdata(nr, trl) has already been forwarded by $p(\mathbf{n})$ to $\mathbf{n}$. In our specification, no such message can be generated. Since the aim of the PGM protocol is to send as few nak-messages as possible and (4) is redundant, we conclude that the (1)-(2)-(3)-(4) behaviour of the [S$^+$00] specification is undesirable. All other behaviours of the original specification, as far as senders and receivers are concerned, can also be simulated in our version (this can be formulated as a simulation invariant). Using this invariant, it is easy to prove that our version satisfies the main property of the protocol if and only if the version of [S$^+$00] does.

In [Mai02], we have proved that the specification of [S$^+$00] is also collapsable, i.e., that $\mathcal{N}^1(P)$ simulates $\mathcal{N}(P)$. The reader is invited to compare this proof with the one given in this paper. Since the specification of [S$^+$00] is not a simple protocol, the proof is much more ad-hoc, and provides little insight about why collapsability is possible.

Proof of Theorem 3

We define a relation $\lhd$ between the network histories of T and $\underline{T}$, and then prove that it is a fair stuttering bisimulation with respect to the mapping Im given by $Im(\mathbf{r}) = \underline{r}$ for every receiver $\mathbf{r}$. Theorem 3 is an immediate corollary of this result and Theorem 2.

Let us first describe the proof technique we use to prove that a relation R is a fair stuttering simulation of T by T' with respect to Im. We proceed in several steps:

- For each network event $(\mathbf{n}, e)$ of T, we define a (possibly empty) sequence $\xi = (\mathbf{n}'_1, e'_1) \ldots (\mathbf{n}'_l, e'_l)$ of network events of $\underline{T}$, such that if $\mathbf{n}$ is the sender or a receiver, then $\mathbf{n}'_1 = \underline{\mathbf{n}}$ (and hence ξ is not empty).
- We prove that if $R(nh, \underline{nh})$ and $nh \xrightarrow[\mathbf{n}]{e} nh'$, then
  (1) $\underline{nh} \xrightarrow[\mathbf{n}'_1 \ldots \mathbf{n}'_l]{e'_1 \ldots e'_l} \underline{nh'}$ is defined and $R(nh', nh'_l)$.
  (2) For every $1 \leq k \leq l$, let $\underline{nh'_k}$ be the trace such that $nh' \xrightarrow[\mathbf{n}'_1 \ldots \mathbf{n}'_k]{e'_1 \ldots e'_k} nh'_k$, then $Match(nh', \underline{nh'_k})$ holds.

R is a fair stuttering simulation because (1) and (2) allow to recursively construct an execution π' of T' for every execution π of T such that the requirements on π' of Definition 1 are satisfied. Given π and assuming that $\sigma(0), \cdots, \sigma(n)$ and $\pi'(0) \cdots \pi'(\sigma(n))$ have been defined, let ξ be sequence assigned to the network

Table 2. Original PGM specification

agent source
odata, *WIN_SIZE*, *txw_trail*: $\mathbb{N}$; *rec_nak*: set of $\mathbb{N}$
s1: in nak(nr) $\rightarrow$ out ncf(nr) downwards;
txw_trail $< nr \rightarrow$ *rec_nak* := add(*rec_nak*, nr);
s2: is_in(nr, *rec_nak*) $nr >$ *txw_trail* $\rightarrow$ out rdata(nr, *txw_trail*) downwards;
rec_nak := remove(*rec_nak*, nr);
s3: length(*rec_nak*) $= 0 \rightarrow$ out odata(*odata*, *txw_trail*) downwards;
odata := *odata* + 1;
odata + 1 $>$ *WIN_SIZE* + *txw_trail* $\rightarrow$
txw_trail := *txw_trail* + *WIN_SIZE*;
endagent;
agent network_element
set_repair, *set_interf*: set of ($\mathbb{N}$, channel_names)
e1: in nak(nr) $\rightarrow$ *set_repair* := *add*(*set_repair*, nr);
set_interf := add(*set_interf*, (nr, c)); [c reception channel];
out nak(nr) upwards;
out ncf(nr) downwards;
e2: in rdata(nr, *trl*) $\rightarrow$ *set_repair* := remove(*set_repair*, nr);
out rdata(nr) to all channels c' s.t. $(nr, c') \in$ *set_interf*
set_interf := remove((nr, c'), *set_interf*)
e3: in ncf(nr) $\rightarrow$ *set_repair* := remove(*set_repair*, nr);
e4: in odata(nr, *trl*) $\rightarrow$ out odata(nr, *trl*) downwards;
e5: is_in(nr, *set_repair*) $\rightarrow$ out nak(nr) upwards;
endagent;

agent receiver
rxw_trail: $\mathbb{N}$; *set_nr*, *set_missing*: set of $\mathbb{N}$
r1: in odata(nr, *trl*) $\wedge$ *rxw_trail* $< nr \rightarrow$ *rxw_trail* $<$ *trl* $\rightarrow$ *rxw_trail* := *trl*;
set_nr := add(*set_nr*, nr);
for all (*rxw_trail* $< i < nr \wedge i \notin$ *set_nr*)
set_missing := add(*set_missing*, i);
set_missing := remove(*set_missing*, nr);
r2: in ncf(nr) $\wedge$*rxw_trail* $< nr \wedge nr \notin$ *set_nr* *set_missing* := add(*set_missing*, nr)
for all (*rxw_trail* $<$ i $< nr \wedge$ i $\notin$ *set_nr*)
set_missing := add(*set_missing*, i);
r3: in rdata(nr, *trl*) $\wedge$ *rxw_trail* $< nr \rightarrow$ *rxw_trail* $<$ *trl* $\rightarrow$ *rxw_trail* := *trl*;
set_nr := add(*set_nr*, nr);
set_missing := remove(*set_missing*, nr);
r4: is_in(nr, *set_missing*) $\rightarrow$ $nr >$ *rxw_trail* $\rightarrow$ out nak(nr) upwards;
$nr \leq$ *rxw_trail* $\rightarrow$ *set_missing* :=
remove(*set_missing*, nr);
endagent;

event $\pi(n)$; we define $\sigma(n+1) \stackrel{\text{def}}{=} \sigma(n) + |\xi|$, and define $\pi'(\sigma(n)+i)$ to be the result of executing $(\mathbf{n}'_1, e'_1) \cdots (\mathbf{n}'_i, e'_i)$. The required properties (a)–(c) hold by (1)and (2) and because if π is fair with respect to some subset $\tilde{T}$ consisting of senders and receivers, then obviously π' is fair with respect to $Im[\tilde{T}]$, and $(\sigma_i)_{i\in\mathbb{N}}$ exceeds all bounds.

Given a network history nh of some network and a set of channels C, let $M\uparrow(nh, C)$ ($M\downarrow(nh, C)$) denote the sum over $c \in C$ of the multisets of upward (downward) messages contained in c after the network history nh. We define the relation $\lhd$ between the network histories of T and $\underline{T}$. $nh \sim \underline{nh}$ holds if the following conditions are satisfied:

(i) If $\mathbf{n}$ is the sender or a receiver, then $tr(nh(\mathbf{n})) = tr(\underline{nh}(\underline{\mathbf{n}}))$.
Informally: if $\mathbf{n}$ is a sender or a receiver, then the agents at the nodes $\mathbf{n}$ and $\underline{\mathbf{n}}$ have executed exactly the same sequence of actions. This guarantees that $nh \lhd nh'$ implies $Match(nh, nh')$.

(ii) For every receiver r, $M\downarrow(nh, C_{\mathbf{r}}) \subseteq M\downarrow(\underline{nh}, [\underline{s}, \underline{\mathbf{r}}])$, where $C_{\mathbf{r}}$ is the set of channels that appear in the path of T that connects $\mathbf{s}$ to $\mathbf{r}$.
Informally: any downward message in any of the channels connecting $\mathbf{s}$ to r can also be found in the channel $[\underline{s}, \underline{\mathbf{r}}]$.

(iii) $M\uparrow(nh, C) \subseteq M\uparrow(\underline{nh}, \underline{C})$, where C, $\underline{C}$ are the sets of channels of T and $\underline{T}$, respectively.
Informally: any upward message in some channel of T can also be found in some channel of $\underline{T}$.

Theorem 7 *The relation $\lhd$ is a fair stuttering simulation relation of T by $\underline{T}$ with respect to the mapping Im given by $Im(\mathbf{r}) = \underline{r}$ for every receiver $\mathbf{r}$.*

Proof: Assume that $nh \lhd \underline{nh}$. For every network event $(\mathbf{n}, e)$ of T such that $nh \xrightarrow[\mathbf{n}]{e} nh'$ is defined, we define a network event $(\underline{\mathbf{n}}, \underline{e})$ of $\underline{T}$ such that $\underline{nh} \xrightarrow[\underline{\mathbf{n}}]{\underline{e}} \underline{nh}'$ is defined and$nh' \lhd \underline{nh}'$.

To show that $\underline{nh} \xrightarrow[\underline{\mathbf{n}}]{\underline{e}} \underline{nh}'$ is defined, we must show that $\underline{nh}(\underline{\mathbf{n}}) \xrightarrow{\underline{e}} \underline{nh}'(\underline{\mathbf{n}})$ is defined and that $i \in M(\underline{nh}, c)$. We distinguish three different cases:

1. Receiver:
The network event is $(\mathbf{r}, e)$ for some receiver $\mathbf{r}$. Let $e = (i, c, o_1, o_2, C)$. Since receivers do not use downward channels, we have $c = [p(\mathbf{r}), \mathbf{r}]$ and $C_2 = \emptyset$ and $o_2 = \perp$.
Define $\underline{e} = (i, [\underline{\mathbf{s}}, \underline{\mathbf{r}}]), o_1, o_2, \emptyset)$. We show that $\underline{nh}(\underline{\mathbf{r}}) \xrightarrow{\underline{e}} \underline{nh}'(\underline{\mathbf{r}})$ is defined. For this we have to show:
   - $i \in M\downarrow(\underline{nh}, [\underline{s}, \underline{\mathbf{r}}])$.
     Since $nh \xrightarrow[\mathbf{r}]{e} nh'$ is defined, i is in the channel $[p(\mathbf{r}), \mathbf{r}]$ before the execution of e, and so in particular $i \in M\downarrow(nh, C_{\mathbf{r}})$. The result now follows from property (ii).
   - $(o_1, o_2) \in \rho_r(tr(\underline{nh}(\mathbf{r})), i)$.
     Since $nh \xrightarrow[\mathbf{r}]{e} nh'$ is defined, we have $(o_1, o_2) \in \rho_r(tr(nh(\mathbf{r})), i)$. The result now follows from property (i).

It remains to show that $nh' \lhd \underline{nh'}$.

(i) $tr(nh'(\mathbf{n})) = tr(\underline{nh'}(\underline{\mathbf{n}}))$ if $\mathbf{n}$ is the sender or a receiver.
$tr(nh(\mathbf{n})) = tr(\underline{nh}(\underline{\mathbf{n}}))$ holds by hypothesis. Since the action corresponding to the two events e and $\underline{e}$ is (i, o_1, o_2), we have

$$tr(nh'(\mathbf{n})) = tr(nh(\mathbf{n})) \cdot (i, o_1, o_2) = tr(\underline{nh}(\underline{\mathbf{n}})) \cdot (i, o_1, o_2) = tr(\underline{nh'}(\underline{\mathbf{n}}))$$

(ii) $M\!\downarrow\!(nh', C_{\mathbf{r}}) \subseteq M\!\downarrow\!(\underline{nh'}, [\underline{s}, \underline{\mathbf{r}}])$.
$M\!\downarrow\!(nh, C_{\mathbf{r}}) \subseteq M\!\downarrow\!(\underline{nh}, [\underline{s}, \underline{\mathbf{r}}])$ holds by hypothesis. For the receivers that do not execute the event e, we have

$$M\!\downarrow\!(nh', C_{\mathbf{r}}) = M\!\downarrow\!(nh, C_{\mathbf{r}}) \subseteq M\!\downarrow\!(\underline{nh}, [\underline{s}, \underline{\mathbf{r}}]) = M\!\downarrow\!(\underline{nh'}, [\underline{s}, \underline{\mathbf{r}}])$$

because e does not add or remove any message to the channels of $C_{\mathbf{r}}$, and $\underline{e}$ does not add or remove any message to the channel $[\underline{s}, \underline{\mathbf{r}}]$. For the receiver r that executes the event e, we have: $M\!\downarrow\!(nh', C_{\mathbf{r}}) = M\!\downarrow\!(nh, C_{\mathbf{r}}) \ominus \{i\} \subseteq M\!\downarrow\!(\underline{nh}, [\underline{s}, \underline{\mathbf{r}}]) \ominus \{i\} = M\!\downarrow\!(\underline{nh'}, [\underline{s}, \underline{\mathbf{r}}])$.

(iv) $M\!\uparrow\!(nh', C) \subseteq M\!\uparrow\!(\underline{nh'}, \underline{C})$.
$M\!\uparrow\!(nh, C) \subseteq M\!\uparrow\!(\underline{nh}, \underline{C})$ holds by hypothesis. We have: $M\!\uparrow\!(nh', C) = M\!\uparrow\!(nh, C) \oplus \{o_1\} \subseteq M\!\uparrow\!(\underline{nh}, \underline{C}) \oplus \{o_1\} = M\!\uparrow\!(\underline{nh'}, \underline{C})$. (Note that $i \in M\!\downarrow$.)

2. Sender:
The network event is $(\mathbf{s}, e)$ for the sender $\mathbf{s}$. Let $e = (i, c, o_1, o_2, C)$. Since the sender has no upward channels, i must be an upward message and $o_1 = \bot$. Moreover, since we assume that the sender only broadcasts messages, we have $C = Ch(\mathbf{s})$.
Since $nh \xrightarrow{e} nh'$ is defined, we have $i \in M\!\uparrow\!(nh, [\mathbf{s}, \mathbf{n}])$ for some child $\mathbf{n}$ of $\mathbf{s}$. Since $nh \lhd \underline{nh}$, by property (iii) we have $i \in M(\underline{nh}, [\mathbf{s}, \mathbf{r}])$ for some channel $[\mathbf{s}, \mathbf{r}]$ of $\underline{T}$. (Notice that all channels of $\underline{T}$ are of the form $[\mathbf{s}, \mathbf{r}]$.). We define the network event $(\underline{\mathbf{n}}, \underline{e})$ as follows: $\underline{\mathbf{n}} = \mathbf{s}$, and $\underline{e} = (i, [\mathbf{s}, \mathbf{r}], o_1, o_2, Ch(\mathbf{s}))$. By (i), $(o_1, o_2) \in \rho_s(tr(\underline{nh}(\mathbf{r})), i)$, and hence $\underline{nh}(\underline{\mathbf{s}}) \xrightarrow{\underline{e}} \underline{nh'}(\underline{\mathbf{s}})$ is defined.
We have to show that $nh' \lhd \underline{nh'}$.

(i) $tr(nh'(\mathbf{s})) = tr(\underline{nh'}(\underline{\mathbf{s}}))$ because the events e and $\underline{e}$ have the same action.

(ii) $M\!\downarrow\!(nh', C_{\mathbf{r}}) \subseteq M\!\downarrow\!(\underline{nh'}, [\underline{s}, \underline{\mathbf{r}}])$: For all $\mathbf{r}$, $M\!\downarrow\!(nh', C_{\mathbf{r}}) = M\!\downarrow\!(nh, C_{\mathbf{r}}) \oplus o_2 \subseteq M\!\downarrow\!(\underline{nh}, [\underline{s}, \underline{\mathbf{r}}]) \oplus o_2 = M\!\downarrow\!(\underline{nh'}, [\underline{s}, \underline{\mathbf{r}}])$.

(iii) $M\!\uparrow\!(nh', C) \subseteq M\!\uparrow\!(\underline{nh'}, \underline{C})$: $M\!\uparrow\!(nh', C) = M\!\uparrow\!(nh, C) \ominus i \subseteq M\!\uparrow\!(\underline{nh}, \underline{C}) \ominus i = M\!\uparrow\!(\underline{nh'}, \underline{C})$.

3. Network element:
The network event is $(\mathbf{n}, e)$ for some network element $\mathbf{n}$. Let $e = (i, c, o_1, C_1, o_2, C_2)$. Define ξ as the empty sequence. As hence $\underline{nh'} = \underline{nh}$, it remains to show that $nh' \lhd \underline{nh}$. Properties (i) and (ii) hold because they hold for nh and $\underline{nh}$ by hypothesis, and because the event e is neither executed by a sender or by a receiver. Let us now consider the other two properties. There are three possible cases.

- $i = \bot$. Then, since ρ_n is forwarding, e is the empty event, and so $nh' = nh \lhd \underline{nh} = \underline{nh'}$.

- $i \in M\uparrow$, i.e., i is an upward message. Then $e = (i, c, i', \bot, \emptyset)$, where $c \in Ch(\mathbf{n})$, and i' can be i or $\bot$.
  For property (ii), we have to show $M\downarrow(nh', C_{\mathbf{r}}) \subseteq M\downarrow(\underline{nh'}, [\underline{\mathbf{s}}, \underline{\mathbf{r}}])$ for every receiver $\mathbf{r}$. This follows immediately from $M\downarrow(nh, C_{\mathbf{r}}) \subseteq M\downarrow(\underline{nh}, [\underline{\mathbf{s}}, \underline{\mathbf{r}}])$, because e does not remove or add upward messages.
  For property (iii), we have to show $M\uparrow(nh', C) \subseteq M\uparrow(\underline{nh'}, \underline{C})$. Since e removes i from one channel and adds it to another one or adds nothing, we have $M\uparrow(nh', C) = M\uparrow(nh, C) = M\uparrow(\underline{nh'}, \underline{C}) = M\uparrow(\underline{nh}, \underline{C})$.
- $i \in M\downarrow$, i.e., i is a downward message. Then $e = (i, ch(\mathbf{n}), \bot, i', C)$, where $C \subseteq Ch(\mathbf{n})$, and i' can be i or $\bot$. (In the latter case, we can assume $C = \emptyset$.)
  For property (ii), we have to show $M\downarrow(nh', C_{\mathbf{r}}) \subseteq M\downarrow(\underline{nh'}, [\underline{\mathbf{s}}, \underline{\mathbf{r}}])$ for every receiver $\mathbf{r}$. For every receiver r, we have $M\downarrow(nh', C_{\mathbf{r}}) \subseteq M\downarrow(nh, C_{\mathbf{r}})$, because e removes the message i from $ch(\mathbf{n})$, and adds it to the channels C. So $M\downarrow(nh', C_{\mathbf{r}}) \subseteq M\downarrow(nh, C_{\mathbf{r}}) \subseteq M\downarrow(\underline{nh}, [\underline{\mathbf{s}}, \underline{\mathbf{r}}]) = M\downarrow(\underline{nh'}, [\underline{\mathbf{s}}, \underline{\mathbf{r}}])$.
  For property (iii), we have to show $M\uparrow(nh', C) \subseteq M\uparrow(\underline{nh'}, \underline{C})$. This follows immediately from $M\uparrow(nh, C) \subseteq M\uparrow(\underline{nh}, \underline{C})$, because e does not remove or add downward messages.

□

Proof of Theorem 4

We define a relation $\lhd$ between the network histories of T and $\underline{T}$, and then prove that it is a fair stuttering bisimulation with respect to the mapping Im given by $Im(\mathbf{r}) = \underline{r}$ for every receiver $\mathbf{r}$. Theorem 3 is an immediate corollary of this result and Theorem 2.

We say $nh \lhd \underline{nh}$ if:

(i) $tr(nh(\mathbf{s})) = tr(\underline{nh}(\underline{\mathbf{s}})$ and for all $\mathbf{r}$, $tr(nh(\mathbf{r})) = tr(\underline{nh}(\underline{\mathbf{r}}))$.
For $\mathbf{n} \neq p(\mathbf{r})$ and $\mathbf{n} \leq \mathbf{r}$, $tr(\underline{nh}(\underline{\mathbf{r}}(\mathbf{n})) = Rem(tr(nh(\mathbf{r})), M_{fil} \cup M_{gen})$.
For all n, $tr(\underline{nh}(\underline{\mathbf{n}})) = Rem(tr(nh(\mathbf{n})), M_{gen})$.

(ii) For all $\mathbf{n}$, $M\downarrow(nh, C^{\mathbf{s}}_{\mathbf{n}}) \setminus M_{gen} \subseteq M\downarrow(\underline{nh}, [\underline{\mathbf{s}}, \underline{\mathbf{n}}]) \setminus M_{gen}$ and for all $\mathbf{r} < \mathbf{n}$, $M\downarrow(nh, C^{\mathbf{n}}_{\mathbf{r}}) \setminus (M_{fil} \cup M_{gen}) \subseteq M\downarrow(\underline{nh}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}(\mathbf{n})]) \setminus (M_{fil} \cup M_{gen})$, where for $\mathbf{n} < \mathbf{n}'$, $C^{\mathbf{n}}_{\mathbf{n}'}$ denotes the set of channels between $\mathbf{n}$ and $\mathbf{n}'$. Moreover, for all $\mathbf{r}$ and $\mathbf{n} = p(\mathbf{r})$, $M\downarrow(nh, [\mathbf{n}, \mathbf{r}]) \subseteq M\downarrow(\underline{nh}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}])$.

(iii) For all $\mathbf{n}$, $M\uparrow(nh, C^{\mathbf{n}}) \setminus M_{gen} \subseteq M\uparrow(\underline{nh}, C^{\underline{\mathbf{n}}}) \setminus M_{gen}$, where $C^{\mathbf{n}}$ is the set of channels below $\mathbf{n}$, and $M\uparrow(nh, Cint) \subseteq M\uparrow(\underline{nh}, \underline{Cint})$, where $Cint$ is the set of channels leading to internal elements. Moreover, for all $\mathbf{r}$ and $\mathbf{n} = p(\mathbf{r})$, $M\uparrow(nh, [\mathbf{n}, \mathbf{r}]) \subseteq M\uparrow(\underline{nh}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}])$.

(iv) For all $\mathbf{r}$, $tr(\underline{nh}(p(\underline{\mathbf{r}})), [p(\underline{\mathbf{r}}), \underline{\mathbf{r}}]) = tr(nh(p(\mathbf{r})), [p(\mathbf{r}), \mathbf{r}])$. I.e. the channels $[p(\underline{\mathbf{r}}), \underline{\mathbf{r}}]$ and $[p(\mathbf{r}), \mathbf{r}]$ satisfy the same filter conditions.

Using the same proof technique as in the proof of Theorem 7, we show that $\lhd$ is a fair stuttering simulation.

Theorem 8 *If P is a simple f&g protocol, the relation $\lhd$ is a fair stuttering simulation relation of T by $\underline{T}$ with respect to the mapping Im given by $Im(\mathbf{r}) = \underline{r}$ for every receiver $\mathbf{r}$.*

Proof:

1. Sender:
   The network event is $(\mathbf{s}, e)$, where $e = (i, c, \perp, o, Ch(\mathbf{s}))$ for some $c \in Ch(\mathbf{s})$. As $c \in \underline{Cint}$, by (iii) it follows that $i \in M\uparrow(nh, Cint)$, and as all channels leading to internal elements are children of $\underline{\mathbf{s}}$, there is $\underline{\mathbf{n}}$ such that $i \in M(\underline{nh}, [\underline{\mathbf{s}}, \underline{\mathbf{n}}])$. Let $\underline{e} = (i, [\underline{\mathbf{s}}, \underline{\mathbf{n}}], \perp, o, Ch(\mathbf{s}))$. By (i), $(\perp, o) \in \rho_s(tr(\underline{nh}(\mathbf{s})), i)$, and hence $\underline{nh} \xrightarrow[\underline{\mathbf{s}}]{\underline{e}} \underline{nh'}$ is defined.
   We have to show $nh' \lhd \underline{nh'}$.
   (i) $tr(nh'(\mathbf{s})) = tr(\underline{nh'}(\underline{\mathbf{s}})$ holds as e and $\underline{e}$ execute the same action $(i, \perp, o)$.
   (ii) For all $\mathbf{n}$, $M\downarrow(nh', C^{\mathbf{s}}_{\mathbf{n}}) \setminus M_{gen} \subseteq M\downarrow(\underline{nh'}, [\underline{\mathbf{s}}, \underline{\mathbf{n}}]) \setminus M_{gen}$:
   For $o \in M_{gen}$, the claim is clear. Otherwise: $M\downarrow(nh', C^{\mathbf{s}}_{\mathbf{n}}) \setminus M_{gen} = M\downarrow(nh, C^{\mathbf{s}}_{\mathbf{n}}) \setminus M_{gen} \oplus o \subseteq M\downarrow(\underline{nh}, [\underline{\mathbf{s}}, \underline{\mathbf{n}}]) \setminus M_{gen} \oplus o = M\downarrow(\underline{nh'}, [\underline{\mathbf{s}}, \underline{\mathbf{n}}]) \setminus M_{gen}$.
   (iii) $M\uparrow(nh', Cint) \subseteq M\uparrow(\underline{nh'}, \underline{Cint})$: $M\uparrow(nh', Cint) = M\uparrow(nh, Cint) \ominus i \subseteq M\uparrow(\underline{nh}, \underline{Cint}) \ominus i = M\uparrow(\underline{nh'}, \underline{Cint})$.
   (iv) Trivial as no internal element executes a step.
2. Receiver:
   The network event is $(\mathbf{r}, e)$, where $e = (i, ch(\mathbf{r}), o, \perp, \emptyset)$.
   For some $\mathbf{n} < \mathbf{r}$, we define an event $e_{\mathbf{n}}$ which is executed by $\underline{\mathbf{r}}(\mathbf{n})$; the sequence ξ consists of $(\underline{\mathbf{r}}, e_{p(\mathbf{r})})$ followed by all $(\underline{\mathbf{r}}(\mathbf{n}), e_{\mathbf{n}})$ for all $\mathbf{n} \neq p(\mathbf{r})$ for which the event $e_{\mathbf{n}}$ is defined.
   By (i), $tr(nh(\mathbf{r})) = tr(\underline{nh}(\underline{\mathbf{r}}))$ holds, and therefore $\rho_r(tr(nh(\mathbf{r})), i) = \rho_r(tr(\underline{nh}(\underline{\mathbf{r}})), i)$; by (ii), $i \in M(\underline{nh}, [p(\underline{\mathbf{r}}), \underline{\mathbf{r}}])$. So we can define $e_{p(\mathbf{r})} = (i, ch(\underline{\mathbf{r}}), o, \perp, \emptyset)$.
   For $\mathbf{n} \neq p(\mathbf{r})$, we have to consider different cases:
   Case $i \notin M_{gen} \cup M_{fil}, i \neq \perp$: Then by (ii), $i \in M\downarrow(nh, C^{\mathbf{n}}_{\mathbf{r}})$ for all $\mathbf{n} < \mathbf{r}$, and we can define $e_{\mathbf{n}} = (i, ch(\underline{\mathbf{r}}(\mathbf{n})), o, \perp, \emptyset)$ for all $\mathbf{n} \leq \mathbf{r}$, because as $M_{gen} \cup M_{fil}$ only enables messages in $M_{gen} \cap M\uparrow$, and $i \neq \perp$, and $\rho(tr(\underline{nh}(\underline{\mathbf{r}}(\mathbf{n})) \in Rem(tr(nh(\mathbf{r})), M_{fil} \cup M_{gen}, M_{gen})$, it follows taht $(o, \perp) \in \rho(tr(\underline{nh}(\underline{\mathbf{r}}(\mathbf{n}))), i)$.
   Case $i \in M_{gen} \cup M_{fil}$: In this case, let $e_{\mathbf{n}}$ be undefined for all $\mathbf{n} \neq p(\mathbf{r})$.
   Case $i = \perp$: If $o \notin M_{gen}$, then because $M_{gen} \cup M_{fil}$ only enables messages in $M_{gen} \cap M\uparrow$, $(o, \perp) \in \rho_n(tr(\underline{nh}(\underline{\mathbf{r}}(\mathbf{n}))), \perp)$ must hold. We define $e_{\mathbf{n}} = (i, ch(\underline{\mathbf{r}}(\mathbf{n})), o, \perp, \emptyset)$ for all $\mathbf{n} \leq \mathbf{r}$. If $o \in M_{gen}$, then let $e_{\mathbf{n}}$ be undefined for all $\mathbf{n} \neq p(\mathbf{r})$.
   We have to show that $nh' \lhd \underline{nh'}$ holds.
   (i) $tr(nh'(\mathbf{r})) = tr(\underline{nh'}(\underline{\mathbf{r}}))$ holds as e and $e_{p(\mathbf{r})}$ have the same action, and for $\mathbf{n} \leq \mathbf{r}$ and $\mathbf{n} \neq p(\mathbf{r})$, $tr(\underline{nh'}(\underline{\mathbf{r}}(\mathbf{n}))) \in Rem(tr(nh'(\mathbf{r})), M_{gen} \cup M_{fil}, M_{gen})$ holds because if $i \notin M_{fil} \cup M_{gen}$ and the action is not $(\perp, o, \perp)$ for $o \in M_{gen}$, then the action of $e_{\mathbf{n}}$ equals the action of e for all $\mathbf{n} < \mathbf{r}$
   (ii) For all $\mathbf{n} < \mathbf{r}$, $M\downarrow(nh', C^{\mathbf{n}}_{\mathbf{r}}) \setminus (M_{fil} \cup M_{gen}) \subseteq M\downarrow(\underline{nh'}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}(\mathbf{n})]) \setminus (M_{fil} \cup M_{gen})$ and $M\downarrow(nh', [p(\mathbf{r}), \mathbf{r}]) \subseteq M\downarrow(\underline{nh'}, [p(\underline{\mathbf{r}}), \underline{\mathbf{r}}])$.
   For $i \in M_{fil} \cup M_{gen}$, the first claim holds trivially, and for $i \notin M_{fil} \cup M_{gen}$,

$M\downarrow(nh', C^{\mathbf{n}}_{\mathbf{r}}) \setminus (M_{fil} \cup M_{gen}) = M\downarrow(nh, C^{\mathbf{n}}_{\mathbf{r}}) \setminus (M_{fil} \cup M_{gen}) \ominus i \subseteq M\downarrow(\underline{nh}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}(\mathbf{n})]) \setminus (M_{fil} \cup M_{gen}) \ominus i = M\downarrow(\underline{nh'}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}(\mathbf{n})]) \setminus (M_{fil} \cup M_{gen})$.
For the second claim: $M\downarrow(nh', [p(\mathbf{r}), \mathbf{r}]) = M\downarrow(nh, [p(\mathbf{r}), \mathbf{r}]) \ominus i \subseteq M\downarrow(\underline{nh}, [p(\underline{\mathbf{r}}), \underline{\mathbf{r}}]) \ominus i = M\downarrow(\underline{nh'}, [p(\underline{\mathbf{r}}), \underline{\mathbf{r}}])$.

(iii) For all $\mathbf{n} < \mathbf{r}$, $M\uparrow(nh', C^{\mathbf{n}}) \setminus M_{gen} \subseteq M\uparrow(\underline{nh'}, C^{\underline{\mathbf{n}}}) \setminus M_{gen}$, and for $\mathbf{n} = p(\mathbf{r})$, $M\uparrow(nh', [\mathbf{n}, \mathbf{r}]) \subseteq M\uparrow(\underline{nh'}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}])$.
For $o \in M_{gen}$, the first claim is trivial, and if $o \notin M_{gen}$, then $e_{\mathbf{n}}$ is defined for all $\mathbf{n} < \mathbf{r}$, because if $i \in M_{fil} \cup M_{gen}$, then $o \in M_{gen}$ follows, as $M_{fil} \cup M_{gen}$ only enables $M_{gen} \cap M\uparrow$.
$M\uparrow(nh', C^{\mathbf{n}}) \setminus M_{gen} = M\uparrow(nh, C^{\mathbf{n}}) \setminus M_{gen} \oplus o \subseteq M\uparrow(\underline{nh}, C^{\underline{\mathbf{n}}}) \setminus M_{gen} \oplus o = M\uparrow(\underline{nh}, C^{\underline{\mathbf{n}}}) \setminus M_{gen}$.
Let $\mathbf{n} = p(\mathbf{r})$: $M\uparrow(nh', [\mathbf{n}, \mathbf{r}]) = M\uparrow(nh, [\mathbf{n}, \mathbf{r}]) \oplus o \subseteq M\uparrow(\underline{nh}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}]) \oplus o = M\uparrow(\underline{nh'}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}])$.

(iv) Trivial since no internal element executes a step.

3. Network element:
Let the network event be $(\mathbf{n}, e)$. We have to consider different cases:
Case $i \in M\uparrow \setminus M_{gen}$: In this case, $e = (i, c, i, m, C)$ for some $c \in Ch(\mathbf{n})$ and $C \subseteq Ch(\mathbf{n})$, and $m \in M_{gen} \cup \{\bot\}$. As $c \in C^{\mathbf{n}}$, by (iii) it follows that $i \in M\uparrow(\underline{nh}, \underline{C^{\mathbf{n}}})$, i.e. $i \in M(\underline{nh}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}(\mathbf{n})])$ for some $\mathbf{r} > \mathbf{n}$. Moreover, if $c = [\mathbf{n}, \mathbf{r}']$ for some $\mathbf{r}'$ such that $\mathbf{n} = p(\mathbf{r}')$, then by (iii), $i \in M\uparrow(\underline{nh}, [\underline{\mathbf{n}}, \underline{\mathbf{r}'}])$; so in that case the channel leading to the receiver $\underline{\mathbf{r}'}$ can be chosen.
Let $\underline{e} = (i, [\underline{\mathbf{n}}, \underline{\mathbf{r}}(\mathbf{n})], i, m, C')$, where C' is the subset of $Ch(\underline{\mathbf{n}})$ of channels c' such that $f_n(tr(\underline{nh}(\underline{\mathbf{n}}), c'), (i, o, o'), c = c') = true$, and let ξ consist of $(\underline{\mathbf{n}}, \underline{e})$. As M_{gen} only enables $\emptyset$, and $tr(\underline{nh}(\underline{\mathbf{n}})) \in Rem(tr(nh(\mathbf{n})), M_{gen}, \emptyset)$, $(i, m) \in \rho_n(tr(\underline{nh}(\underline{\mathbf{n}})), i)$, and hence $\underline{nh} \xrightarrow[\underline{\mathbf{n}}]{\underline{e}} \underline{nh'}$ is defined. By (iv) it follows that $[\mathbf{n}, \mathbf{r}] \in C$ iff $[\underline{\mathbf{n}}, \underline{\mathbf{r}}] \in C'$ for all $\mathbf{r} > \mathbf{n}$.
We have to show that $nh' \lhd \underline{nh'}$.

(i) $tr(\underline{nh'}(\underline{\mathbf{n}})) \in Rem(tr(nh'(\mathbf{n})), M_{gen}, \emptyset)$ holds as e and $\underline{e}$ have the same action.

(ii) We have to show that for all $\mathbf{n}'$, $M\downarrow(nh', C^{\mathbf{s}}_{\mathbf{n}'}) \setminus M_{gen} \subseteq M\downarrow(\underline{nh'}, [\underline{\mathbf{s}}, \underline{\mathbf{n}'}]) \setminus M_{gen}$ and $M\downarrow(nh', C^{\mathbf{n}'}_{\mathbf{r}}) \setminus (M_{fil} \cup M_{gen}) \subseteq M\downarrow(\underline{nh'}, [\underline{\mathbf{n}'}, \underline{\mathbf{r}}(\mathbf{n})]) \setminus (M_{fil} \cup M_{gen})$ and for $\mathbf{r}$ such that $\mathbf{n} = p(\mathbf{r})$, $M\downarrow(nh', [\mathbf{n}, \mathbf{r}]) \subseteq M\downarrow(\underline{nh'}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}])$.
The first two claims follows directly from the hypothesis as $m \in M_{gen}$.
If $[\mathbf{n}, \mathbf{r}] \notin C$, the last claim follows from the hypothesis; otherwise, $[\underline{\mathbf{n}}, \underline{\mathbf{r}}] \in C'$, and $M\downarrow(nh', [\mathbf{n}, \mathbf{r}]) = M\downarrow(nh, [\mathbf{n}, \mathbf{r}]) \oplus m \subseteq M\downarrow(\underline{nh}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}]) \oplus m = M\downarrow(\underline{nh'}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}])$.

(iii) For all $\mathbf{n}$, $M\uparrow(nh', C^{\mathbf{n}}) \setminus M_{gen} \subseteq M\uparrow(\underline{nh'}, C^{\underline{\mathbf{n}}}) \setminus M_{gen}$, and $M\uparrow(nh', \mathit{Cint}) \subseteq M\uparrow(\underline{nh'}, \underline{\mathit{Cint}})$:
$M\uparrow(nh', C^{\mathbf{n}}) \setminus M_{gen} = M\uparrow(nh, C^{\mathbf{n}}) \setminus M_{gen} \ominus i \subseteq M\uparrow(\underline{nh}, C^{\underline{\mathbf{n}}}) \setminus M_{gen} \ominus i = M\uparrow(\underline{nh'}, C^{\underline{\mathbf{n}}}) \setminus M_{gen}$.
If $c \notin \mathit{Cint}$, then $M\uparrow(nh', \mathit{Cint}) = M\uparrow(nh, \mathit{Cint}) \oplus i \subseteq M\uparrow(\underline{nh}, \underline{\mathit{Cint}}) \oplus i = M\uparrow(\underline{nh'}, \underline{\mathit{Cint}})$.
Otherwise, the claim holds as $M\uparrow(nh', \mathit{Cint}) = M\uparrow(nh, \mathit{Cint})$.

(iv) We only have to consider the case that $\mathbf{r}$ is a child of $\mathbf{n}$. If $c = [\mathbf{n}, \mathbf{r}]$ then by definition the input channel of event $\underline{e}$ is $[\underline{\mathbf{n}}, \underline{\mathbf{r}}]$. This means that we get the following equations:
$tr(\underline{nh'}(\underline{\mathbf{n}}), [\underline{\mathbf{n}}, \underline{\mathbf{r}}]) = tr(\underline{nh}(\underline{\mathbf{n}}), [\underline{\mathbf{n}}, \underline{\mathbf{r}}])\cdot(i, m) = tr(nh(\mathbf{n}), [\mathbf{n}, \mathbf{r}])\cdot(i, m) = tr(nh'(\mathbf{n}), [\mathbf{n}, \mathbf{r}])$, where $m = \bot$ if $c \notin C$ and $c \notin C'$.
For $c \neq [\mathbf{n}, \mathbf{r}]$ we get the same equation, but with i set to $\bot$.

Case $i \in M_{gen} \cap M\uparrow$: in this case, $e = (i, c, i, \bot, \emptyset)$ for some $c \in Ch(\mathbf{n})$, as M_{gen} only enables $\emptyset$. If $c = [\mathbf{n}, \mathbf{r}']$ for some child $\mathbf{r}'$ of $\mathbf{n}$, then by (iii), $i \in M\uparrow(\underline{nh}, [\underline{\mathbf{n}}, \underline{\mathbf{r}'}])$; in this case, let $\underline{e} = (i, [\underline{\mathbf{n}}, \underline{\mathbf{r}'}], i, \bot, \emptyset)$ and let ξ consist of $(\underline{\mathbf{n}}, \underline{e})$. Otherwise, let ξ be empty.

We have to show that $nh' \lhd \underline{nh'}$.

(i) We have to show that $tr(\underline{nh'}(\underline{\mathbf{n}})) \in Rem(tr(nh'(\mathbf{n})), M_{gen}, \emptyset)$, which holds as the action of e is of form (i, o_1, o_2) for $i \in M_{gen}$.

(ii) Trivial as $i \notin M\downarrow$, and hence no message from $M\downarrow$ is removed from or added to some channel.

(iii) $M\uparrow(nh', C^{\mathbf{n}}) \setminus M_{gen} \subseteq M\uparrow(\underline{nh'}, C^{\underline{\mathbf{n}}}) \setminus M_{gen}$ follows directly from the hypothesis as $i \in M_{gen}$.
$M\uparrow(nh', Cint) \subseteq M\uparrow(\underline{nh'}, \underline{Cint})$:
If $c \neq [\mathbf{n}, \mathbf{r}']$ for some child $\mathbf{r}'$ of $\mathbf{n}$, $M\uparrow(nh', Cint) = M\uparrow(nh, Cint)$, as i is moved from one channel in $Cint$ to another channel in $Cint$. Otherwise, $M\uparrow(nh', Cint) = M\uparrow(nh, Cint) \oplus i \subseteq M\uparrow(\underline{nh}, \underline{Cint}) \oplus i = M\uparrow(\underline{nh'}, \underline{Cint})$.
If $c \neq [\mathbf{n}, \mathbf{r}']$ for some child $\mathbf{r}'$ of $\mathbf{n}$, then $M\uparrow(nh', [\mathbf{n}, \mathbf{r}]) \subseteq M\uparrow(\underline{nh'}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}])$ follows from the hypothesis.
Otherwise, if $c = [\mathbf{n}, \mathbf{r}']$, then $M\uparrow(nh', [\mathbf{n}, \mathbf{r}']) = M\uparrow(nh, [\mathbf{n}, \mathbf{r}']) \ominus i \subseteq M\uparrow(\underline{nh}, [\underline{\mathbf{n}}, \underline{\mathbf{r}'}]) \ominus i = M\uparrow(\underline{nh'}, [\underline{\mathbf{n}}, \underline{\mathbf{r}'}])$.

(iv) $tr(\underline{nh'}(p(\underline{\mathbf{r}})), [p(\underline{\mathbf{r}}), \underline{\mathbf{r}}]) = tr(nh'(p(\mathbf{r})), [p(\mathbf{r}), \mathbf{r}])$.
Only in case that $\mathbf{r}$ is a child of $\mathbf{n}$ has to be considered. If $c = [\mathbf{n}, \mathbf{r}]$ then by definition, the input channel of $\underline{e}$ is $[\underline{n}, \underline{\mathbf{r}}]$, and $tr(\underline{nh'}(\underline{\mathbf{n}}), [\underline{\mathbf{n}}, \underline{\mathbf{r}}]) = tr(\underline{nh}(\underline{\mathbf{n}}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}])\cdot(i, \bot) = tr(nh(\mathbf{n}), [\mathbf{n}, \mathbf{r}])\cdot(i, \bot) = tr(nh'(\mathbf{n}, [\mathbf{n}, \mathbf{r}])$.
If $c \neq [\mathbf{n}, \mathbf{r}]$, we get the same equations with i set to $\bot$.

Case $i \in M\downarrow \setminus M_{gen}$: In this case, $e = (i, ch(\mathbf{n}), o, i, C)$ for some $C \subseteq Ch(\mathbf{n})$ and $o \in M_{gen} \cup \bot$.

As $ch(\mathbf{n}) \in C^{\mathbf{s}}_{\mathbf{n}}$, by (ii) it follows that $i \in M\Downarrow(\underline{nh}, [\underline{\mathbf{s}}, \underline{\mathbf{n}}])$. So we can define $\underline{e} = (i, ch(\underline{\mathbf{n}}), o, i, \underline{C})$, where $\underline{C} \subseteq Ch(\underline{\mathbf{n}})$ is the set of channels c for which $f_n(tr(\underline{nh}(\underline{\mathbf{n}}), c), (i, o, i), false)$ returns true. By (i) and as M_{gen} only enables $\emptyset$, $(o, i) \in \rho_n(tr(\underline{nh}(\underline{\mathbf{n}})), i)$. So $\underline{nh} \xrightarrow{\underline{e}} \underline{nh'}$ is defined. Moreover, by (iv) it follows that $[\mathbf{n}, \mathbf{r}] \in C$ iff $[\underline{\mathbf{n}}, \underline{\mathbf{r}}] \in \underline{C}^{\underline{\mathbf{n}}}$.

We have to show that $nh' \lhd \underline{nh'}$.

(i) $tr(\underline{nh'}(\underline{\mathbf{n}})) \in Rem(tr(nh'(\mathbf{n})), M_{gen}, \emptyset)$:
Holds as e and $\underline{e}$ execute the same action.

(ii) We have to show that for all $\mathbf{n}'$, $M\Downarrow(nh', C^{\mathbf{s}}_{\mathbf{n}'}) \setminus M_{gen} \subseteq M\Downarrow(\underline{nh'}, [\underline{\mathbf{s}}, \underline{\mathbf{n}'}]) \setminus M_{gen}$:

For all $\mathbf{n}' \neq \mathbf{n}$, message i is at most removed from channels in $C^{\mathbf{s}}_{\mathbf{n}'}$, and the content of $[\underline{\mathbf{s}}, \underline{\mathbf{n}'}]$ is unchanged, and hence the claim holds by the hypothesis.
For $\mathbf{n}$: $M\downarrow\!(nh', C^{\mathbf{s}}_{\mathbf{n}}) \setminus M_{gen} = M\downarrow\!(nh, C^{\mathbf{s}}_{\mathbf{n}}) \setminus M_{gen} \ominus i \subseteq M\downarrow\!(\underline{nh}, [\underline{\mathbf{s}}, \underline{\mathbf{n}}]) \setminus M_{gen} \ominus i = M\downarrow\!(\underline{nh'}, [\underline{\mathbf{s}}, \underline{\mathbf{n}}]) \setminus M_{gen}$.
Next we have to show that for all $\mathbf{n}'$, $M\downarrow\!(nh', C^{\mathbf{n}'}_{\mathbf{r}}) \setminus (M_{fil} \cup M_{gen}) \subseteq M\downarrow\!(\underline{nh'}, [\underline{\mathbf{n}'}, \underline{\mathbf{r}}(\mathbf{n}')]) \setminus (M_{fil} \cup M_{gen})$.
For $\mathbf{n}' \neq \mathbf{n}$, i is only moved within channels of $C^{\mathbf{n}'}_{\mathbf{r}}$, or the content of all channels in $C^{\mathbf{n}'}_{\mathbf{r}}$ is unchanged, while the content of $[\underline{\mathbf{n}'}, \underline{\mathbf{r}}(\mathbf{n}')]$ does not change. For $\mathbf{n}$, we only have to consider the case $i \notin M_{fil} \cup M_{gen}$: $M\downarrow\!(nh', C^{\mathbf{n}}_{\mathbf{r}}) \setminus (M_{fil} \cup M_{gen}) = M\downarrow\!(nh, C^{\mathbf{n}}_{\mathbf{r}}) \setminus (M_{fil} \cup M_{gen}) \oplus i \subseteq M\downarrow\!(\underline{nh}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}(\mathbf{n})]) \setminus (M_{fil} \cup M_{gen}) \oplus i = M\downarrow\!(\underline{nh'}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}(\mathbf{n})]) \setminus (M_{fil} \cup M_{gen})$ for all $\mathbf{r} < \mathbf{n}$.
Last, we have to show that for all $\mathbf{r}$ such that $\mathbf{n} \prec \mathbf{r}$, $M\downarrow\!(nh', [\mathbf{n}, \mathbf{r}]) \subseteq M\downarrow\!(\underline{nh'}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}])$. if $[\mathbf{n}, \mathbf{r}] \in C$ then $[\underline{\mathbf{n}}, \underline{\mathbf{r}}] \in \underline{C}$.

(iii) We have to show that for all $\mathbf{n}'$, $M\uparrow\!(nh', C^{\mathbf{n}'}) \setminus M_{gen} \subseteq M\uparrow\!(\underline{nh}, C^{\underline{\mathbf{n}'}}) \setminus M_{gen}$. As $i \in M\downarrow$ and $o \in M_{gen}$, this follows from the induction hypothesis.
Next we have to show that $M\uparrow\!(nh', Cint) \subseteq M\uparrow\!(\underline{nh'}, \underline{Cint})$: $M\uparrow\!(nh', Cint) = M\uparrow\!(nh, Cint) \oplus o \subseteq M\uparrow\!(\underline{nh}, \underline{Cint}) \oplus o = M\uparrow\!(\underline{nh'}, \underline{Cint})$. Last, we have to show that for all $\mathbf{r}$ such that $\mathbf{n} \prec \mathbf{r}$, $M\uparrow\!(nh', [\mathbf{n}, \mathbf{r}]) \subseteq M\uparrow\!(\underline{nh'}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}])$. This follows from the fact that $i \notin M\uparrow$.

Case $i \in M_{gen} \cap M\downarrow$: By definition of simple protocol, $e = (i, ch(\mathbf{n}), \bot, \bot, \emptyset)$; note that since $i \in M\downarrow$, it cannot be sent upwards. Let ξ be empty. In order to show that $nh' \lhd \underline{nh'}$, we only have to consider (i), as no messages are added to or removed from channels:
$tr(\underline{nh'}(\underline{\mathbf{n}})) = tr(\underline{nh}(\underline{\mathbf{n}})) \in Remove(tr(nh(\mathbf{n})) \cdot (i, \bot, \bot), Mgen, \emptyset)$, as $i \in M_{gen}$.

Case $i = \bot$: In this case, $e = (\bot, ch(\mathbf{n}), o_1, o_2, C)$ for some $o_i \in M_{gen} \cup \bot$ and $C \subseteq Ch(\mathbf{n})$. By (i), $tr(\underline{nh}(\underline{\mathbf{n}})) \in Rem(tr(nh(\mathbf{n})), M_{gen}, \emptyset)$, and hence $(o_1, o_2) \in \rho_n(tr(\underline{nh}(\underline{\mathbf{n}})), \bot)$. So we can define $\underline{e} = (\bot, ch(\underline{\mathbf{n}}), o_1, o_2, C')$, where C' is the subset of $Ch(\underline{\mathbf{n}})$ of channels c' such that $f_n(tr(\underline{nh}(\underline{\mathbf{n}}), c'), (\bot, o_1, o_2), false) = true$; By (iv), $[\mathbf{n}, \mathbf{r}] \in C$ iff $[\underline{\mathbf{n}}, \underline{\mathbf{r}}] \in C'$. So $\underline{nh} \xrightarrow[\underline{\mathbf{n}}]{\underline{e}} \underline{nh'}$ is defined. We have to show that $nh' \lhd \underline{nh'}$.

(i) $tr(\underline{nh'}(\underline{\mathbf{n}})) \in Rem(tr(nh'(\mathbf{n})), M_{gen}, \emptyset)$ holds as e and $\underline{e}$ contain the same action.

(ii) We have to show that for all $\mathbf{n}'$, $M\downarrow\!(nh', C^{\mathbf{s}}_{\mathbf{n}'}) \setminus M_{gen} \subseteq M\downarrow\!(\underline{nh'}, [\underline{\mathbf{s}}, \underline{\mathbf{n}'}]) \setminus M_{gen}$ and $M\downarrow\!(nh', C^{\mathbf{n}'}_{\mathbf{r}}) \setminus (M_{fil} \cup M_{gen}) \subseteq M\downarrow\!(\underline{nh'}, [\underline{\mathbf{n}'}, \underline{\mathbf{r}}(\mathbf{n})]) \setminus (M_{fil} \cup M_{gen})$.
For all $\mathbf{r}$ such that $\mathbf{n} = p(\mathbf{r})$, we have to show that $M\downarrow\!(nh', [\mathbf{n}, \mathbf{r}]) \subseteq M\downarrow\!(\underline{nh'}, [\underline{\mathbf{n}}, \underline{\mathbf{r}}])$.
The first two claim hold as $o_2 \in M_{gen}$, and the last one holds as $[\mathbf{n}, \mathbf{r}] \in C$ implies $[\underline{\mathbf{n}}, \underline{\mathbf{r}}] \in C'$.

(iii) We have to show that for all $\mathbf{n}'$, $M\!\uparrow\!(nh, C^{\mathbf{n}'}) \setminus M_{gen} \subseteq M\!\uparrow\!(\underline{nh}, C^{\underline{\mathbf{n}'}}) \setminus M_{gen}$. As $o_1 \in M_{gen}$, this follows directly from the hypothesis. Next we have to show that $M\!\uparrow\!(nh', Cint) \subseteq M\!\uparrow\!(\underline{nh}', \underline{Cint})$: $M\!\uparrow\!(nh', Cint) = M\!\uparrow\!(nh, Cint) \oplus o_1 \subseteq M\!\uparrow\!(\underline{nh}, \underline{Cint}) \oplus o_1 = M\!\uparrow\!(\underline{nh}', \underline{Cint})$.
Last we have to show that for all children $\mathbf{r}$ of $\mathbf{n}$, $M\!\uparrow\!(nh', [\mathbf{n}, \mathbf{r}]) \subseteq M\!\uparrow\!(\underline{nh}', [\underline{\mathbf{n}}, \underline{\mathbf{r}}])$. This follows from the hypothesis as no upward message is added to $[\mathbf{n}, \mathbf{r}]$.

(iv) For all $\mathbf{r}$ such that $\mathbf{n} = p(\mathbf{r})$, $tr(\underline{nh}'(\underline{\mathbf{n}}), [\underline{\mathbf{n}}, \underline{\mathbf{r}}]) = tr(nh'(\mathbf{n}), [\mathbf{n}, \mathbf{r}])$: This follows from the fact that $[\mathbf{n}, \mathbf{r}] \in C$ implies $[\underline{\mathbf{n}}, \underline{\mathbf{r}}] \in C'$.

□

Proof of Theorem 5:

Let $M\!\uparrow = \{m_1, \ldots, m_u\}$ and $M_{gen} \cap M\!\uparrow = \{m_1, \ldots, m_g\}$. Let $T \in \mathcal{N}^1(P)$ and let $r_1, \ldots, r_k$ be receivers in T, and let π be an execution of T. We need to identify the following nodes: Remember that if P is simple, actions of ρ_n of form (i, o, o') such that $o \in M_{gen}$ satisfy $o' = \perp$. Actions of receivers trivially have this property.

- Let $\mathbf{n}(\mathbf{r}_i)$ be the parent of $\mathbf{r}_i$.
- For $m_j \in M\!\uparrow$, let $\mathbf{r}^i(m_j)$ be the receiver below $\mathbf{n}(\mathbf{r}_i)$ not equal to r_i that first performs an action with output $(m_j, \perp)$ within execution π.
- For every $m_h \in M_{gen} \cap M\!\uparrow$, let $\mathbf{m}(m_h)$ be the internal element not equal to some $\mathbf{n}(\mathbf{r}_i)$ that first outputs $(m_h, \perp)$.
- For every $m_j \in M\!\uparrow$, let $\mathbf{t}^h(m_j)$ be the child of $\mathbf{m}(m_h)$ that first outputs $(m_j, \perp)$.
- For every $m_j \in M\!\uparrow$, let $\mathbf{v}(m_j)$ be the first receiver not equal to some receiver below $\mathbf{n}(\mathbf{r}_i)$ for some i that outputs $(m_j, \perp)$, and let $\mathbf{o}_j$ be its parent.
- For every $m_j \in M\!\uparrow$, let $\mathbf{v}^i(m_j)$ be the first receiver below o_i not equal to $\mathbf{v}(m_i)$ that outputs $(m_j, \perp)$.

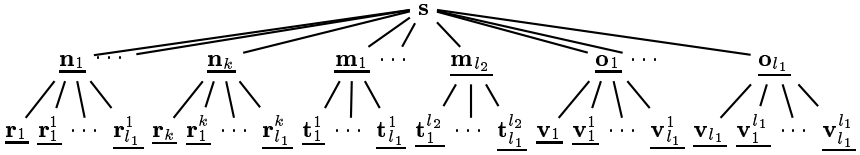

Fig. 2. U_k

Figure 6 gives names to the nodes in U_k. In the path π' of U_k corresponding to π,

- $\underline{\mathbf{r}}_i$ executes the actions of $\mathbf{r}_i$,

- $\underline{\mathbf{r}}_h^i$ executes the actions of $\mathbf{r}^i(m_h)$,
- $\underline{\mathbf{n}}_i$ executes the actions of $\mathbf{n}(r_i)$,
- $\underline{\mathbf{m}}_j$ executes the actions of $\mathbf{m}(m_j)$,
- $\underline{\mathbf{t}}_i^j$ executes the actions of $\mathbf{t}^h(m_i)$,
- $\underline{\mathbf{v}}_i$ executes the actions of $\mathbf{v}(m_i)$,
- $\underline{\mathbf{v}}_h^i$ executes the actions of $\mathbf{v}^i(m_h)$ and
- $\underline{\mathbf{o}}_i$ executes the actions of $\mathbf{o}_i$.

The intuition is that $\underline{\mathbf{r}}_1^i, \ldots, \underline{\mathbf{r}}_u^i$ mimic the behaviour of the neighbour receivers of $\mathbf{r}_i$ and $\underline{t}_1^j, \ldots, \underline{t}_u^j$ mimic the behaviour of the receivers below $\mathbf{m}(m_j)$, and $\underline{\mathbf{v}}_1^i, \ldots, \underline{\mathbf{v}}_u^i$ mimic the behaviour of the neighbour receivers of $v(m_i)$.

Note that there can be overlap: for example if $\mathbf{r}_i$ and $\mathbf{r}_h$ are neighbours, then $\mathbf{n}(\mathbf{r}_i) = \mathbf{n}(\mathbf{r}_h)$ and possibly $\mathbf{r}^i(m_j) = \mathbf{r}^h(m_j)$, and hence $\underline{n}_i$ and $\underline{\mathbf{n}}_h$ both have to execute the actions of $\mathbf{n}(\mathbf{r}_i)$, and $\underline{r}_j^i$ and $\underline{r}_j^h$ both have to execute the actions of $\mathbf{r}^i(m_j)$. Similarly, $\mathbf{t}^i(m_j) = \mathbf{v}_j$ is possible, which means that $m_i = o_j$.

This is not a problem as duplication of network elements with their receivers below is possible as the sender multicasts messages downwards to all network elements.

The path π' is defined by defining for every transition $\pi(h) \xrightarrow[\mathbf{n}]{e} \pi(h+1)$ a sequence of network events $\xi(h)$ in which all processes that have to execute the actions of $\mathbf{n}$ execute $a(e)$. π' is the execution corresponding to $\xi(0)\xi(1)\ldots$.

We define $\xi(h)$ in a way that the following holds:

(i) If $\pi(h) \xrightarrow[\mathbf{x}]{e} \pi(h+1)$ for some $\mathbf{x} = \mathbf{s}, \mathbf{n}(\mathbf{r}_i), \mathbf{r}_i$, then $\xi(h)$ contains the network event $(\mathbf{x}', e')$ for $\mathbf{x} = \underline{s}, \underline{n}_i$ or $\underline{r}_i$ respectively, where $a(e) = a(e')$. I.e. the corresponding nodes in U_k execute the same action.

(ii) If $\pi(h) \to \pi(h+1)$ is a transition in which a network element not equal to $\mathbf{n}(\mathbf{r}_1), \ldots, \mathbf{n}(\mathbf{r}_k)$ sends a message m_i upwards (i.e. a transition that can affect the sender), then $\xi(h)$ contains a transition of $\underline{m}_i$ sends the same message upwards.

More precisely, we define π' inductively as follows. Suppose $\xi(h')$ has been defined for all $h' < h$.

We have to distinguish between the following cases. Note that as explained above, the cases are not exclusive. $\xi(h)$ consists of all events specified in matching clauses.

- $\pi(h) \xrightarrow[\mathbf{x}]{e} \pi(h+1)$ for some $\mathbf{x} = \mathbf{s}, \mathbf{r}_i$. Let $\xi(h)$ consists of the network event $(\underline{r}_i, e')$ or $(\underline{s}, e')$ respectively, such that $a(e) = a(e')$.
- $\pi(h) \xrightarrow[\mathbf{n}(\mathbf{r}_i)]{e} \pi(h+1)$ for some i. Let $\xi(h)$ contains the network events $(\underline{n}_i, e')$, where $a(e) = a(e')$.
- If $\pi(h) \xrightarrow[\mathbf{r}^i(m_j)]{e} \pi(h+1)$ and h is smaller than or equal to the first number at which $\mathbf{r}^i(m_j)$ outputs $(m_j, \perp)$, then $\xi(h)$ contains $(\underline{\mathbf{r}}_j^i, e')$ such that $a(e) = a(e')$.

- Analogously, if $\pi(h) \xrightarrow[\mathbf{m}^i(m_h)]{e} \pi(h+1)$ and h is smaller than or equal to the first number at which $\mathbf{m}(m_h)$ outputs $(m_h, \bot)$, then $\xi(h)$ contains $(\underline{\mathbf{m}}_j, e')$, where $a(e') = a(e)$.
- To make transitions $\xi(h)$ of $\underline{\mathbf{m}_i}$ enabled, we need to simulate the receivers below $\mathbf{m}(m_i)$: If $\pi(h) \xrightarrow[\mathbf{t}^i(m_j)]{e} \pi(h+1)$ and h is smaller than or equal to the first number at which $\mathbf{m}(m_i)$ outputs $(m_i, \bot)$ and the number at which $\mathbf{t}^i(m_j)$ outputs $(m_j, \bot)$, then $\xi(h)$ contains $(\underline{\mathbf{t}}^i_j, e')$, where $a(e') = a(e)$.
- We need processes $\mathbf{v}_j$ in order to simulate receivers that are not children of some $\mathbf{n}(\mathbf{r}_i)$: If $\pi(h) \xrightarrow[\mathbf{v}(m_j)]{e} \pi(h+1)$ and h is smaller than the first number at which $\mathbf{v}(m_j)$ outputs $(m_j, \bot)$, then $\xi(h)$ contains $(\underline{\mathbf{v}_j}, e')$ such that $a(e) = a(e')$.
- The role of $\underline{\mathbf{o}_j}$ is to act exactly like $\mathbf{o}_j$, in particular to forward all messages that $\underline{\mathbf{v}_j}$ sends upwards: If $\pi(h) \xrightarrow[\mathbf{o}_j]{e} \pi(h+1)$, then $\xi(h)$ contains $(\underline{\mathbf{o}_j}, e')$ such that $a(e') = a(e)$.
- Finally, we need processes $\underline{\mathbf{v}}^i_j$ in order to simulate the receivers below $\mathbf{o}_i$ besides $\mathbf{v}(m_i)$: if $\pi(h) \xrightarrow[\mathbf{v}^i(m_j)]{e} \pi(h+1)$, and h is smaller than the first number at which $v(m_i)$ outputs $(m_i, \bot)$ and the number at which $\mathbf{v}^i(m_j)$ outputs $(m_j, \bot)$, then $\xi(h)$ consists of $(\underline{\mathbf{v}}^i_j, e')$ such that $a(e') = a(e)$.
- For all h such that $\pi(h) \xrightarrow[\mathbf{n}]{e} \pi(h+1)$ is such that a message is sent upwards, and $\xi(h)$ is not defined by the previous cases, we define $\xi(h)$ as follows:
  If $\mathbf{x}$ is a receiver below $\mathbf{n}(m_i)$ and the output is $(m_j, \bot)$, then let $\xi(h)$ consist of $(\underline{\mathbf{r}}^i_j, e')$, where $a(e') = (\bot, m_j, \bot)$. Such an event is executable by definition of iteration protocol, as $\underline{r}^i_j$ has been simulated in π' up to the first output of $(m_j, \bot)$.
  Analogously, if $\mathbf{x}$ is a network element not equal to some $\mathbf{n}(\mathbf{r}_i)$ and $(m_h, \bot)$ is sent for some $m_h \in M_{gen}$, then let $\xi(h)$ contain $(\underline{\mathbf{m}_h}, e')$, where $a(e') = (\bot, m_h, \bot)$.
  If $\mathbf{x}$ is a receiver below $\mathbf{m}_h$ and m_j is set upwards, then let $\xi(h)$ be $(\underline{t^h_j}, e')$ such that $a(e') = (\bot, m_j, \bot)$.
  If $\mathbf{x}$ is a receiver not below some $\mathbf{n}(r_i)$, and m_j is sent upwards, let $\xi(h)$ contain $(\underline{v}_j, e')$ such that $a(e') = (\bot, m_j, \bot)$.
  If $\mathbf{x}$ is a network element other than $\mathbf{n}(\mathbf{r}_i)$ and a message $m_j \notin M_{gen}$ is sent upwards, this means that $a(e)$ is a forwarding action. So some receiver not below $\mathbf{n}(\mathbf{r}_i)$ must have sent $(m_j, \bot)$, and by the previous clause this means that $\underline{v}_j$ has sent $(m_j, \bot)$. So we can define $\xi(h)$ to contain $(\underline{o}_j, e')$, where $a(e') = a(e)$.
  Finally, if $\mathbf{x}$ is a receiver below o_i other than v_i, and m_j is sent upwards, let $\xi(h)$ contain $(\underline{v^i_j}, e')$ such that $a(e') = (\bot, m_j, \bot)$.
- In all other cases, let $\xi(h)$ be empty.

It follows from the construction that all transitions in $\xi(h)$ are enabled.

The construction gives in an obvious way rise to an increasing $\sigma : \mathbb{N} \to \mathbb{N}$ such that $tr(\pi(l), \mathbf{s}) = tr(\pi'(\sigma(l)), \underline{\mathbf{s}})$ and $tr(\pi(l), \mathbf{r}_i) = tr(\pi'(\sigma(l)), \underline{\mathbf{r}_i})$ for all i, i.e. $Match(\pi(l), \pi'(\sigma(l))$, and for all $\sigma(l) < j \leq \sigma(l+1)$, $Match(\pi(l+1), \pi'(j))$.