Formal Models of the OSPF Routing Protocol

Jack Drury
Data61, CSIRO, Sydney, Australia
Jack.Drury@data61.csiro.au

Peter Höfner
Research School of Computer Science
ANU, Canberra, Australia
Data61, CSIRO, Sydney, Australia
Computer Science and Engineering
UNSW, Sydney, Australia
Peter.Hoefner@anu.edu.au

Weiyou Wang
Data61, CSIRO, Sydney, Australia
Weiyou.Wang@data61.csiro.au

We present three formal models of the OSPF routing protocol. The first two are formalised in the timed process algebra T-AWN, which is not only tailored to routing protocols, but also specifies protocols in pseudo-code that is easily readable. The difference between the two models lies in the level of detail (level of abstraction). From the more abstract model we then generate the third model. It is based on networks of timed automata and can be executed in the model checker Uppaal.

1 Introduction

Finding routes between nodes within networks is one of the most common tasks in networking. To solve this problem, researchers and engineers have designed dozens of different routing protocols, which specify how routers communicate with each other and distribute information that enables them to select routes between any two nodes. Even though most of these protocols are based on 'simple' techniques such as Dijkstra's shortest path algorithm [11] or the Bellman-Ford algorithm[1], it seems incredibly hard to ensure that the protocols are functionally correct. For example, routing protocols regularly establish non-optimal routes [31], the protocol AODV can yield routing loops [19] or the Border Gateway Protocol (BGP) can exhibit persistent route oscillations [41].

These examples show that careful protocol analysis is essential. However, a formal analysis needs to be based on an unambiguous model, which often does not exist as "despite the maturity of formal description languages and formal methods for analyzing them, the description of real protocols is still overwhelmingly informal" [42].

We present formal and unambiguous models for the Open Shortest Path First (OSPF) routing protocol [33, 9], a widely used interior gateway protocol. While we are not the first to create formal models for OSPF (see Section 7), it is our belief that we present the first formal model that covers not only the core functionality of OSPF, but also most other details defined in the protocol standard (Section 4). As usual, that standard [33, 9] is written in English prose and contains ambiguities, which we had to resolve.

Our detailed model is written in the process algebra T-AWN [6, 7], which we briefly describe in Section 3. It is a variant of standard process algebras, such as CCS [30], CSP [23] or ACP [4]. As for any process algebra, T-AWN's semantics is completely formal and hence our model is absolutely precise (no contradictions, no under-/over-specifications) and free of ambiguities. Moreover, T-AWN is particularly tailored for routing protocols such as OSPF and defines the protocol in pseudo-code that is easily readable by any network or software researcher/engineer.

Modelling many aspects and details of a routing protocol is important, but for an (initial) formal analysis a more abstract model can be extremely useful. For example, it allows an analysis using model

[1] First proposed by Shimbel [39]; later Bellman [3], Ford [16] and Moore [34] published the same algorithm.

A. Fehnker and H. Garavel (Eds.):
Models for Formal Analysis of Real Systems (MARS 2020)
EPTCS 316, 2020, pp. 72–120, doi:10.4204/EPTCS.316.4

© J. Drury, P. Höfner & W. Wang
This work is licensed under the
Creative Commons Attribution License.

checking techniques, which is infeasible with too many details present, due to state space problems. To this end we present, in Section 5, a second model, also written in T-AWN, that abstracts from aspects such as the retransmission of lost messages, the periodic refreshing of link information and the internal hierarchy of subnets. Last but not least, in Section 6, we translate that model into a network of timed automata, which can be executed by the model checker Uppaal [25, 2]. All three models are presented in full in the appendices.

2 The Open Shortest Path First (OSPF) Routing Protocol

The Open Shortest Path First (OSPF) protocol [33] is a widely used proactive, link-state routing protocol used to distribute routing information throughout a single autonomous system. Being a link-state protocol means that the routers (nodes in the network) exchange topological information about one-hop links with their neighbours, i.e. the nodes within transmission range. That information is flooded through the network so that (eventually) every router has a complete picture of available links in the network. This image is used to calculate the shortest/best routes between any two nodes, usually using a variant of Dijkstra's algorithm [11].

OSPF routers use HELLO *messages* to discover neighbouring nodes. HELLO messages are broadcast at regular intervals by every node and contain, next to the sender's identification, a list of all those nodes that the sender has received HELLO messages from. These messages are not flooded through the network, but are sent only a single hop from their origin.

A node will distribute information about the state of its connections by sending *Link State Advertisements* (LSAs).

Data received by a router that is used for determining the network topology is stored in the *Link-State Database (LSDB)*. This database represents the router's current view of the network topology; it contains the most recently received LSAs from each unique originator. Next to the LSDB, every node maintains a list of discovered neighbours, with whom routing information may be exchanged.

Network nodes also use HELLO messages (or lack thereof) to determine whether neighbours have become inactive or lost connectivity. When a node receives the first HELLO message from a neighbour, it learns of that neighbour's existence. If the node finds its own IP address listed within that HELLO message—meaning that the neighbour is aware of the node's existence—it checks whether it needs to form an adjacency with that neighbour.

The term *adjacency* describes the relationship between two neighbouring nodes that must exchange all of their topological information—the content of their LSDBs. Eventually, they will have an identical understanding of the network. Not all one-hop neighbours will form an adjacency. This is because network traffic can be greatly reduced by forming as few adjacencies as possible whilst still ensuring that all single-hop neighbours will end up with synchronised LSDBs (through a 'chaining' of adjacencies). When two nodes recognise that they need to become adjacent one of them will be nominated as *master*, the other as *slave*. It is the master that initiates the exchange of data by sending the necessary information using *Database Description (*DBD*) messages*. The slave sends DBD messages in response to DBD messages from the master.

Once a node has received a full description of the other node's LSDB, it compares that description with its own LSDB. To resolve inconsistencies between the two LSDBs, such as missing entries, it sends *Link State Request* (LSR) messages, asking for LSAs containing the newest information available. LSR messages are answered by *Link State Update* (LSU) messages; they contain the requested LSAs. Each receipt of an LSU message is acknowledged by a *Link State Acknowledgement* (LSACK) message.

Designated routers are used in combination with adjacency to further reduce network traffic. As they are not crucial in the understanding of the main functionality of OSPF we omit a detailed description.

Our model follows closely the OSPF specification as described in the RFC version 2 [33]; the amendments described in [9] are taken into account as well. However, many of these changes concern the introduction of IPv6 support—the original specification only allowed IPv4—; as we model IP addresses as arbitrary unique identifiers the amendments had little effect on our model.

3 The Specification Language T-AWN

One of the standard tools to describe interactions, communications and synchronisations between a collection of agents, processes or network nodes are process algebras. They provide algebraic laws that allow formal reasoning. We chose to model OSPF using T-AWN [6, 7], a *timed* process algebra specifically tailored for wireless networks in general and routing protocols in particular.

T-AWN 'extends' the (untimed) language AWN (Algebra of Wireless Networks) [14], by timing constructs. In fact, the syntax of both algebras is (nearly) identical. Both AWN and T-AWN provide the right level of abstraction to model key protocol features, while abstracting from implementation-related details. As its semantics is completely unambiguous, specifying a protocol in such a framework enforces total precision and the removal of any ambiguity. (T-)AWN is tailored for modelling and verifying routing and communication protocols and therefore offers primitives such as **unicast** and **multicast/groupcast**;it defines the protocol in a pseudo-code that is easily readable—the language itself is implementation independent. Currently the tool support is limited; that is why we present an executable model in Section 6.

(T-)AWN is a variant of standard process algebras (e.g. [23, 29, 4]) extended with a local broadcast mechanism and a novel *conditional unicast* operator—allowing error handling in response to failed communications while abstracting from link layer implementations of the communication handling—and incorporating data structures with assignments; its operational semantics is defined in [14].

We use an underlying data structure (described later) with several types, variables ranging over these types, operators and predicates. First order predicate logic yields terms (or *data expressions*) and formulas to denote data values and statements about them. The data structure has to contain the types DATA, MSG, IP and $\mathscr{P}(\texttt{IP})$ of *application layer data*, *messages*, *identifiers* and *sets of identifiers*. The messages comprise *data packets*, containing application layer data, and *control messages*. A network is modelled as a parallel composition of network routers. Several processes may run in parallel on a single router.

An entire network is modelled as an encapsulated parallel composition of network nodes.

Nodes can only communicate with their direct neighbours, i.e. with nodes that are currently within transmission range. There are three different ways for internode communication: broadcast, unicast, or an iterative unicast/multicast (called *groupcast* in (T-)AWN).

The *process expressions* are given in Table 1. They should be understandable without further explanation; we add a short description in Appendix A.

When designing or formalising a protocol in (T-)AWN, an engineer should not be bothered with timing aspects; except for functions and procedures that schedule tasks depending on the current time. Because of this, the only difference between the syntax of AWN and the one of T-AWN is that the later is equipped with a local timer now.

T-AWN assumes a discrete model of time, where each sequential process maintains the local variable now holding its local clock value—an integer. Only one clock for each sequential process is employed. All (sequential) processes in a network synchronise in taking time steps, and at each time step all local clocks are incremented by one time unit. For the rest, the variable now behaves as any other variable

Table 1: process expressions [20]

$X(exp_1,\ldots,exp_n)$	process name with arguments
$P+Q$	choice between process P and Q
$[\varphi]P$	conditional process (if-statement)
$[\![\mathtt{var} := exp]\!]P$	assignment followed by process P
$\textbf{broadcast}(ms).P$	broadcast *ms* followed by P
$\textbf{groupcast}(dests,ms).P$	iterative unicast or multicast to all destinations *dests*
$\textbf{unicast}(dest,ms).P \blacktriangleright Q$	unicast *ms* to *dest*; if successful proceed with P; otherwise with Q
$\textbf{send}(ms).P$	synchronously transmit *ms* to parallel process on same node
$\textbf{deliver}(data).P$	deliver data to application layer
$\textbf{receive}(\mathrm{msg}).P$	receive a message
ξ,P	process with valuation
$V \langle\!\langle W$	parallel procs. on the same node
$id:V:R$	node *id* running V with range R
$N \Vert M$	parallel composition of nodes
$[N]$	encapsulation

maintained by a process: its value can be read when evaluating guards, thereby making progress time-dependent, and any value can be assigned to it. T-AWN does not model clock drift nor clock skew implicitly: however, as `now` is a standard variable that can be changed, these concepts could be integrated in a model. As neither of the two concepts is of practical relevance for `OSPF` we ignore them.

4 A Detailed Model of OSPF

Our first model, written in the specification language T-AWN (Section 3), is a detailed rendering of OSPF as described in the *Request For Comments* (RFC) [33, 9], the de-facto standard; only a few abstractions are made. The model is about 165 lines long, split over 9 processes, and uses around 30 functions. Compared to the 244 pages of the RFC, this is a significant reduction in size, while being more precise. We do not advocate to abandon a specification written in English prose as it often describes the intention and the intuition of protocol designs. However, we believe that many problems of protocol development and specification could be avoided if formal protocol descriptions, such as the one presented here, would accompany the textual specification.

In this section, we describe the abstractions and deviations from the RFC taken, present the overall structure of our model and discuss in detail the T-AWN process `HELLO`, which models all the actions after the receipt of a HELLO message. The full model is given in Appendix C.

4.1 Abstractions & Deviations

In OSPF only certain pairs of nodes will form adjacencies—not to be confused with node connectivity—and share topological information. Although the procedure to find these pairs varies with the capabilities of the network in question, such determination always relies on information provided by a network administrator. As we cannot model the intention of a network administrator, we assume the existence of a predicate $\mathtt{adj}(\mathrm{ip},\mathrm{ip}')$, which evaluates to true if the nodes ip and ip$'$ are supposed to form an adjacency.

To enhance scalability, OSPF allows networks to be split into multiple *areas*. Some of the nodes that are connected to at least two areas will summarise and share intra-area topological information between the adjoining areas. Usually, the introduction of areas greatly reduces the number of messages sent. For the moment we assume a single area; we plan to add this feature in the future.

On top of these abstractions we also make the following assumptions.

(a) We assume that every message has a payload of arbitrary size. In reality, as the content of LSDBs can be very large, a node splits the information into smaller fragments before sending and the receivers reassemble these fragments. However, this feature is independent of the core functionality of OSPF.
(b) Messages are complete and correct and do not become corrupted (we do allow packet loss, though).
(c) Our model focuses on the sharing of LSAs and synchronisation of LSDBs. Since LSAs are only shared via bidirectional connections, we assume all connections to be bidirectional.
(d) We abstract from the so-called fight-back mechanism (Sect. 13.4 of [33]), which handles the receipt of self-originated LSAs.

4.2 Overall Structure

The detailed model consists of the 9 processes `OSPF`, `HELLO`, `DBD`, `SNMIS`, `REQ`, `UPD`, `ACK`, `QMSG` and `QSND`.

- The main process `OSPF` reads a single message from the input queue. Depending on the type of the message it calls other processes, such as `HELLO`. The process also sends HELLO messages at periodic intervals and maintains the node's LSDB by removing dead neighbours.
- The process `HELLO` describes all the actions performed when a HELLO message is received. This includes updating the relevant inactivity timer and, if the HELLO message is from a previously unknown neighbour, updating the node's neighbour list.
- The process `DBD` handles incoming DBD messages. Sending requests for those LSAs from DBD message that do not have a matching partner in the node's LSDB is among its actions.
- The process `SNMIS` describes the actions required when a sequence number mismatch occurs during the adjacency-establishment procedure.
- The process `REQ` describes the actions following the receipt of an LSR message, such as finding the requested LSAs in the node's local data, and sending them back to the sender of the message.
- The process `UPD` manages incoming LSU messages, including the installation of up-to-date LSAs in the node's LSDB, and broadcasting the updated information.
- The process `ACK` describes the actions taken to handle the receipt of an LSA.
- The process `QMSG` models a message queue. As before we have two instances, a queue for incoming messages and one for outgoing processes.
- The process `QSND` is responsible for sending outgoing messages that were generated by the other processes described above.

4.3 Handling HELLO Messages

The process `OSPF` handles the receipt of a message using the expression **receive**(msg). In case the received message is a HELLO message, which is tested by the command **[**$msg = \texttt{hello}(\text{ips},\text{sip})$**]**, the process `HELLO` is called. Here, sip is the unique identifier of the sender (usually its IP address) and ips is the actual content of the HELLO message, listing all one-hop neighbours of sip.

Next to the content of the received HELLO message (ips and sip), `HELLO` maintains the following data: the node's own identifier ip, a set nbrs of *neighbour structures*, the Link-State Database lsdb (see Section 2) and a time value hello_t that indicates the time at which ip should send the next HELLO message—remember that HELLO messages are sent periodically.

Each neighbour structure stored in nbrs contains, next to the identifier nip of the neighbour, additional information such as an inactivity timer. The inactivity timer indicates when ip should consider a node nip inactive, and hence when the entry concerning nip should be removed from nbrs. There are seven other fields that we only discuss in the appendix.

HELLO first checks whether the sender sip of the HELLO message is already known,i.e. whether the node has received a HELLO message sent by sip before. This check is performed using the predicate nbrExist, which takes the set nbrs of neighbour structures and the sender identifier sip as parameters.

If sip is unknown (Line 1) the sender is added to the list of known neighbours nbrs, using the function newNBR (Line 2). Afterwards the process determines what to do, by calling the process HELLO—as the predicate nbrExist(nbrs, sip) now evaluates to true, the actions between Lines 5 and 21 are performed.

Process 9 Handling HELLO messages

```
HELLO(ips,sip,ip,nbrs,lsdb,hello_t) ≝
 1. [ ¬nbrExist(nbrs,sip) ]        /* the sender sip is unknown */
 2.   [[nbrs := newNBR(nbrs,sip)]]
 3.   HELLO(ips,sip,ip,nbrs,lsdb,hello_t)
 4. + [ nbrExist(nbrs,sip) ]       /* sip is a known neighbour */
 5.   [[nbrs := setINACTT(nbrs,sip,now+rtdead_intvl)]]
 6.   [[ns := getNS(nbrs,sip)]]
 7.   (
 8.     [ ip ∈ ips ∧ adj(ip,sip) ∧ ns = “Init” ]        /* 2-WayReceived, start adjacency-forming */
 9.       [[nbrs := setNS(nbrs,sip,“ExStart”)]]
10.       [[nbrs := incDDSQN(nbrs,sip)]]
11.       [[nbrs := setDDT(nbrs,sip,now+rxmt_intvl)]]
12.       send(sndmsg(genDBD(nbrs,lsdb,sip,ip),{sip})) . OSPF(ip,nbrs,lsdb,hello_t)
13.     + [ ip ∈ ips ∧ adj(ip,sip) ∧ ns ≥ “ExStart” ]     /* Adjacency-forming already underway */
14.       OSPF(ip,nbrs,lsdb,hello_t)
15.     + [ ip ∈ ips ∧ ¬adj(ip,sip) ]        /* 2-WayReceived, adjacency-forming not needed */
16.       [[nbrs := setNS(nbrs,sip,“2-Way”)]]
17.       OSPF(ip,nbrs,lsdb,hello_t)
18.     + [ ip ∉ ips ]        /* 1-WayReceived */
19.       [[nbrs := initNBR(nbrs,sip,“Init”)]]
20.       OSPF(ip,nbrs,lsdb,hello_t)
21.   )
```

If the neighbour exists the process first resets the inactivity timer for sip in Line 5: the value of now + rtdead_intvl identifies the latest time when the ip should have received another HELLO message from sip. In Line 6, the *neighbour state* for node sip is distilled from nbrs and stored in the local variable ns. The neighbour state is one of the additional pieces of information in the neighbour structure. It is a string representing the adjacency status of ip and sip; its value is one of the following.

Init: ip has received a HELLO message from sip, but a bidirectional communication is not yet verified;

2-Way: communication has been identified as bidirectional, but adj does not allow adjacency.

The remaining values indicate different stages during the creation of an adjacency between sip and ip.

ExStart: the first step of the adjacency-establishment procedure, which decides which router is the master, and which is the slave;

Exchange: the two routers ip and sip exchange descriptions of their entire LSDBs, using DBD messages;

Loading: nodes compare the received LSDB description with their own and send out LSR *messages* to find missing information, collect newer information or to resolve conflicting information;

Full: adjacency is established between ip and sip; the nodes now have identical LSDBs.

Depending on the current status between ip and sip, the process HELLO performs different actions.

Let us first consider the case where the receiver's ip is listed in the incoming HELLO message (ip $\in$ ips). Lines 9 to 12 are executed when the two routers want to form an adjacency, but the process has not yet begun (the status of ns is "Init"). In this case, the process updates the value of ns to "ExStart" (Line 9), increments *the Database Description Sequence Number* (DDSQN), which is used to ensure that both nodes correctly communicate a description of their entire LSDB (Line 10), resets the timer for the retransmission of DBD messages (Line 11), generates and sends a DBD message (Line 12) and then returns to the main process OSPF. For the sake of readability we do not detail how DBD messages are created. However, it is worth mentioning that all messages generated by processes such as HELLO are not sent directly to nodes in transmission range, but they are first passed to another process running on the same node, using the primitive **send**. It is that process that **broadcast**s, **groupcast**s and **unicast**s the message, respectively. The reason for that design is two-fold: First, it reflects reality as on every router there exists an output queue, which collects and sends messages, using a lower-level protocol such as CSMA [24]. Second, sending messages takes time (in reality and in T-AWN); if a message would be sent from the process HELLO it could potentially block other incoming message (and their effects).

If the nodes sip and ip have already started the adjacency-establishment procedure—the neighbour state is either "ExStart", "Exchange", "Loading" or "Full"—then no further action is required and the process returns to the main process OSPF (Line 14). If ip is listed in ips but the two routers do not want to form an adjacency (Line 15), the neighbour state ns is set to "2-Way" (Line 16). No other action is performed and the protocol returns to the process OSPF.

The last possible sequence of actions (Lines 19 and 20) is executed if ip is not listed in ips and therefore the sender of the HELLO message is not aware of the receiver's existence. In this case the process sets ns to "Init" for the neighbour structure of sip and returns to OSPF. If it was the case that prior to the receipt of this HELLO message the formula ns $>$ "2-Way" evaluates to true, it means that the two nodes had, at the very least, attempted to form an adjacency. Hence Line 19 also wipes any information related to previous adjacency-forming attempts.

5 An Abstract Model of OSPF

The model described above is too complex for automated analyses such as model checking. For example, a network consisting of only two nodes needs to exchange 18 messages before reaching a steady state. One of the main reasons for this complexity is that there is a lot of redundancy built into OSPF to protect against message loss. Another contributing factor is that fresh LSAs are created and distributed regularly, even when no topological changes have occurred. If we assume that all sent messages are received[2] we can simplify the model without changing its core functionality. For example, we no longer need retransmission lists and timers, and we only need to generate LSAs in case of a topological change.

The simplified model comes in at roughly 70 lines, splits over 7 processes and uses only 12 functions. As before, we only present the handling of HELLO messages; the full model is given in Appendix B.

5.1 Abstractions & Deviations

As this model is an abstraction of the previous model, we make the same assumptions as in Section 4.

Additionally, we now assume that all sent messages are received. This means that we no longer need to maintain retransmission lists, as their sole purpose is to hold messages that are yet to be acknowledged.

[2] This is one of the fundamental features of (T-)AWN.

We are also able to remove acknowledgement messages for the same reason. We assume that all LSR messages will be handled and responded to correctly. Along with that, we further assume that an LSR message can contain multiple LSA headers. By doing so, a router can request all LSAs needed by sending a single request. According to the RFC all routers should refresh their LSAs at regular intervals by sending out LSU messages, even if nothing has changed. We assume that this is no longer necessary and remove this functionality. Finally, we assume that each node will form adjacencies with all its one-hop neighbours. In terms of exchanging topological information, this is equivalent to assuming that all connections are point-to-point.

5.2 Handling HELLO Messages

The overall structure is the same as for the detailed model (Section 4.2), with the exception that the processes SNMIS and ACK are not needed any longer. The simplified process HELLO is shown in Process 2. As we no longer have to check whether nodes form adjacencies the process simplifies drastically.

Process 2 Handling HELLO messages (simplified)

$\texttt{HELLO}(\mathrm{ips},\mathrm{sip},\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t}) \stackrel{def}{=}$

1. $[\ \neg\texttt{nbrExist}(\mathrm{nbrs},\mathrm{sip})\ ]$ /* the sender sip is unknown */
2. $[\![\mathrm{nbrs} := \texttt{newNBR}(\mathrm{nbrs},\mathrm{sip})]\!]$
3. $[\![\mathrm{nbrs} := \texttt{setINACTT}(\mathrm{nbrs},\mathrm{sip},\texttt{now}+\texttt{rtdead_intvl})]\!]$
4. $[\![\mathrm{lsa} := \texttt{newLSA}(\mathrm{ip},\texttt{now},\mathrm{nbrs})]\!]$ /* generate new LSA */
5. $[\![\mathrm{lsdb} := \texttt{install}(\mathrm{lsdb},\{\mathrm{lsa}\})]\!]$ /* install LSA */
6. $\textbf{send}(\texttt{sndmsg}(\texttt{upd}(\{\mathrm{lsa}\},\mathrm{ip}),\{nip \mid (nip,*)\in \mathrm{nbrs}\}))\ .$ /* send LSA */
7. $\textbf{send}(\texttt{sndmsg}(\texttt{dbd}(\{\texttt{hdr}(lsa) \mid lsa \in \mathrm{lsdb}\},\mathrm{ip}),\{\mathrm{sip}\}))\ .$ /* send DBD back to sip */
8. $\texttt{OSPF}(\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$
9. $+\ [\ \texttt{nbrExist}(\mathrm{nbrs},\mathrm{sip})\ ]$ /* sip is a known neighbour */
10. $[\![\mathrm{nbrs} := \texttt{setINACTT}(\mathrm{nbrs},\mathrm{sip},\texttt{now}+\texttt{rtdead_intvl})]\!]$
11. $\texttt{OSPF}(\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$

The only case distinction that remains is to check whether router ip has seen messages from the sender sip of the HELLO message (Lines 1 and 9). In the case that ip has previously received a message from the sender ($\texttt{nbrExist}(\mathrm{nbrs},\mathrm{sip})$), the inactivity timer is reset (Line 10) and the process returns to the OSPF process (Line 11), ready to receive another message. These two lines are the only actions that remain of Lines 5–21 in the detailed Process 9. Otherwise, if the message received is the first HELLO message from sip ($\neg\texttt{nbrExist}(\mathrm{nbrs},\mathrm{sip})$), the router ip creates a new neighbour structure (Line 2) and initialises its inactivity timer (Line 3). Taking into account that in the detailed model the process HELLO is called in Line 3, these two lines correspond to Line 2 and Line 5 of Process 9, respectively. Afterwards, ip generates a new LSA (Line 4), adds it to its own database lsdb (Line 5) and sends it out to all of the neighbours it is aware of, which are determined by the formula $\{\mathrm{nip} \mid (\mathrm{nip},*)\in \mathrm{nbrs}\}$; here, all one-hop neighbours will form an adjacency and there is no lengthy establishment procedure, therefore the nodes are considered adjacent upon discovery.

Then, in Line 7, the receiving router ip sends a DBD message to the originator sip of the HELLO message so that they can compare their LSDBs. This is similar to Line 12 of Process 9. However, the data structure of the detailed model is more complex and therefore more assignments and different functions are needed. Finally the process returns to OSPF.

6 An Executable Model of OSPF

We translate the simplified T-AWN model into an executable model, which can be used by Uppaal [2, 25], which is an established model checker for *networks of timed automata*. Uppaal is well suited for an automated analysis of (T-)AWN specifications as it provides (a) two synchronisation mechanisms, which translate to uni- and broadcast communication; (b) common data structures, such as arrays and structs, and a C-like programming language to define updates on these data structures; (c) mechanisms for time.

6.1 Uppaal's Specification Language

Uppaal accepts networks of timed automata with guards and data structures as input. The state of the system is determined, in part, by the values of data variables that can be either shared between automata, or local. As for (T-)AWN, we assume a data structure with several types, variables ranging over these types, operators and predicates.

Each automaton is a graph, with locations, and edges between locations. Every edge has a guard—if not specified the guard is `true`—, optionally a synchronisation label, and an update. Synchronisation occurs via so-called channels. For each channel a there is one label $a!$ to denote the sender, and $a?$ to denote the receiver. Transitions without labels are internal; all others use one of two types of synchronisation.

In *binary handshake* synchronisation, one automaton having an edge with a label that has the suffix ! synchronises with a single other automaton that has an edge with the same label including a ?-suffix. These two transitions synchronise when both guards are true in the current state, and only then. When the transition is taken both locations change, and the updates will be applied to the state variables; first the updates on the !-edge, then the updates on the ?-edge. If there is more than one possible pair, then the transition is selected non-deterministically.

In *broadcast* synchronisation, one automaton with a !-labelled edge synchronises with a set of other automata. The initiating automaton can change its location, and apply its update, if the guard on its edge evaluates to true. It does not require a second synchronising automaton. Automata with a matching ?-labelled edge have to synchronise if their guard is currently true. They change their location and update the state. The automaton with the !-edge will update the state first, followed by the other automata.

6.2 T-AWN to Uppaal

The translation from (T-)AWN to Uppaal is more or less straightforward. A possible translation from AWN to Uppaal is sketched in [32].[3]

The main process `OSPF` (including the subprocesses it calls, such as `HELLO`) of every node in the network is modelled as a single automaton, each having its own data structures such as an LSDB and message buffers. A node's output queue is modelled as a separate automaton. As a slight optimisation we are able to avoid another automaton for the input queue. A network of n nodes is modelled by n copies of the two automata; each tagged with a unique identifier.

The implementation of the data structure defined in T-AWN is straightforward. An LSDB for example is an array of LSAs, one entry for every node in the network. An LSA is given by the data type

```
typedef struct
{ int [0,N-1] ip;            // identifier of originator
  int[0,age_bound+1] age;    // time when the LSA is originated
} LSA;
```

where `age` denotes the time when the LSA was generated.

[3] A similar translation is also mentioned in [13].

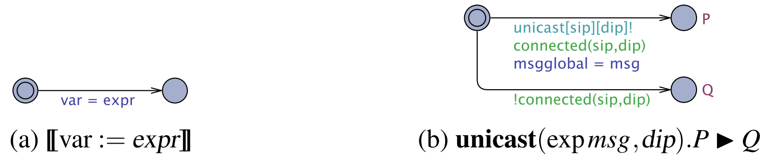

(a) $[\![\text{var} := \textit{expr}]\!]$ (b) **unicast**$(\exp msg, dip).P \blacktriangleright Q$

Figure 1: Translating T-AWN to Uppaal

An assignment $[\![\text{var} := \textit{expr}]\!]$ is translated into an edge with a single update, as depicted in Figure 1a. Similarly, the translation of a conditional $[\varphi]$ yields an edge with a single guard. Translating the sending primitives is more complicated as one has to take care of the intended destinations, the connectivity between nodes as well as the exchange of the actual messages. Figure 1b depicts the translation schema for **unicast**: the locations named `P` and `Q` stand for sub-automata resulting from the translation of the subprocess P and Q; `unicast[sip][dip]` represents a unique channel where data is sent from node `sip` to `dip`; the predicate `connected` characterises the topology—only when the two nodes are connected is the message transferred. When the transition is taken, the sender `sip` copies its message to a global variable `msgglobal`, and the receiver copies it subsequently to its local buffer. As we currently do not have software supporting this translation we perform the translation manually; this allows us to optimise the Uppaal model. For example, a sequence of assignments can be combined into a single edge.

A major difference between the (T-)AWN- and Uppaal models is that in Uppaal message queues can only be of finite length, which is in contrast to the T-AWN model. The local message buffer is modelled as a fixed-length array of messages. Uppaal will give a warning if during model checking an out-of-bounds error occurs, i.e. if the array was too small. Another difference is that (T-)AWN is based on discrete time, while Uppaal uses dense time. It is easy to see that (for our models) an embedding of discrete time into dense time does not cause any problems.

We translate the simpler model, presented in Section 5, into an executable Uppaal model. The resulting timed automaton modelling `OSPF`, including its subprocesses such as `HELLO`, is presented in Figure 2.

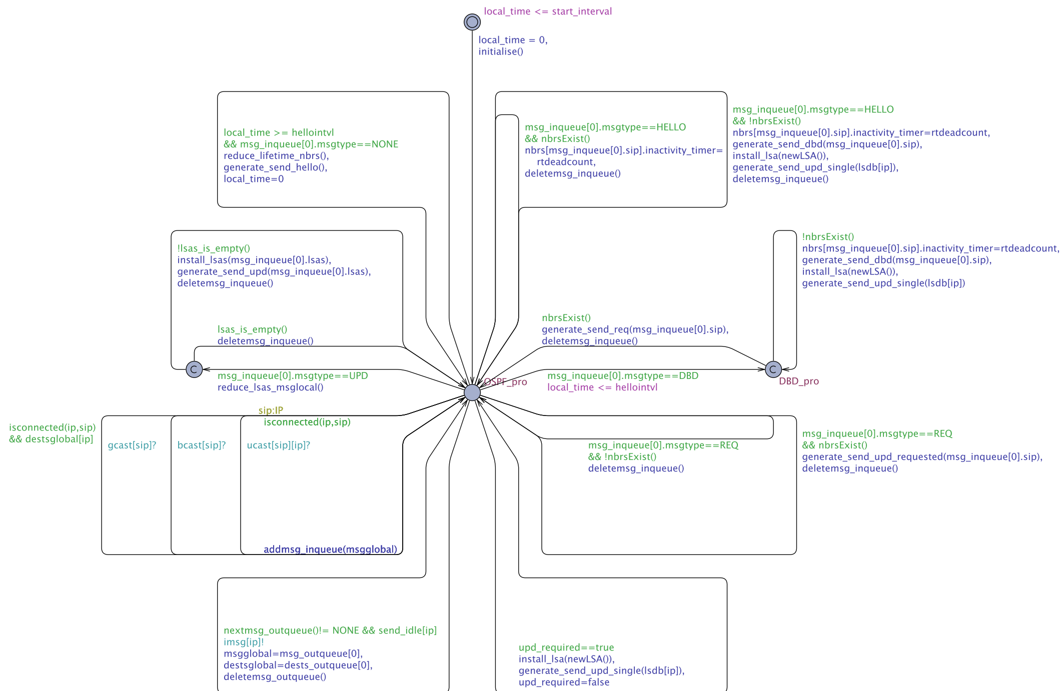

Figure 2: Executable Model of OSPF as timed automaton

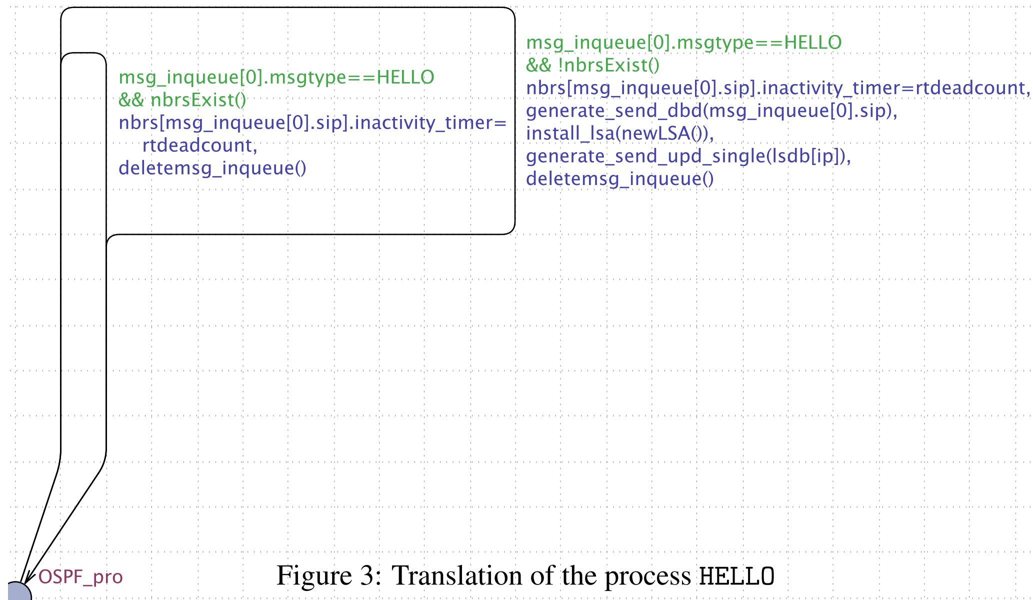

Figure 3: Translation of the process `HELLO`

We have assumed that each node will form adjacencies with all its one-hop neighbours; an assumption that significantly decreased the size of our model. For model checking, however, this assumption potentially increases the state space as creating and maintaining adjacencies requires a lot of message passing.

6.3 Handling Hello Messages

To illustrate the relationship between T-AWN and Uppaal we present the translation of the `HELLO` process in Figure 3. That figure is an enlarged version of the right-upper corner of the automaton in Figure 2. For simplicity we skip the data structures and the definitions of functions occurring in the timed automaton; for the full description see Appendix D.

The two edges correspond with the two conditions in Lines 1 and 9 of the process `HELLO` (Process 2).

The outermost edge—the one that states `&& !nbrsExist()` in the second line—combines Lines 1–8 of Process 2 in a single edge. The guard first checks that the received message is indeed a HELLO message (`msg_inqueue[0].msgtype == HELLO`) and then validates whether it received a HELLO message from the sender using `&& !nbrsExist()`. The latter is a direct translation of Line 1. The remaining actions of Process 2 are translated into simple updates of the data structure. Additionally we have to delete the received HELLO message from the input queue. This is done by `deletemsg_inqueue`.

The innermost path corresponds directly to Lines 9–11. It describes the scenario where the router has seen a HELLO message from the sender before. As before, we need the additional action of deleting the HELLO message.

7 Related Work

For our modelling efforts we use the process algebra T-AWN, as well as networks of timed automata that form the input language of the model checker Uppaal.

Next to T-AWN, several process algebras modelling broadcast communication have been proposed, such as the Calculus of Broadcasting Systems (CBS) [37, 38], the $b\pi$-calculus [12], CBS# [36], the Cal-

culus of Wireless Systems (CWS) [28], the Calculus of Mobile Ad Hoc Networks (CMAN) [21], the Calculus for Mobile Ad Hoc Networks (CMN) [27], the ω-calculus [40], restricted broadcast process theory (RBPT) [17], $bA\pi$ [22] and the broadcast psi-calculi [5]. However, T-AWN is the only process algebra that provides all features needed to fully model routing protocols such as OSPF, namely data handling, (conditional) unicast and (local) broadcast. Moreover, all other process algebras lack the feature of guaranteed receipt of messages. Due to this, it is impossible to analyse properties such as route discovery in process algebras different to T-AWN: only in T-AWN a failure of route discovery can be interpreted as an imperfection in the protocol, rather than as a result of a chosen formalism not ensuring guaranteed receipt.

There are only a few other formal models for OSPF. Nakibly et al. [35] created a model of `OSPF` for the model checker CBMC [8][4]. Their model uses a fixed topology—similar to our model-checking models, but different to our T-AWN models—and abstracts from many crucial details such as HELLO messages and DBD messages. The online model—a link is given in [35]—analyses a tiny topology with three nodes. Another model, that can be fed into the Z3-Solver [10], is described in [26]. The authors claim that it is a detailed model covering concepts such as designated routers and claim that network topologies up to 30 nodes can be analysed. This would be very impressive, but the model is neither described in detail in [26] nor available online; hence a comparison to our model is impossible. It is our belief, that the formal models available are far too abstract to give guarantees about the correctness of OSPF. By providing more models we bridged this gap.

8 Conclusion and Future Work

We have presented three different models of the routing protocol OSPF. Two are written in the specification language T-AWN, a process algebra particularly tailored to routing protocols. The first model covers many details of OSPF and, to the best of our knowledge, is by far the most detailed model available. The second model is based on additional assumptions, making it significantly shorter and easier. From the simpler model we have derived a model that can be executed in the model checker Uppaal.
Future work is three-fold.

(a) We plan to add additional concepts to our detailed model. This includes protocol-specific details such as the fight-back mechanism (Sect. 13.4 of [33]) and a more dynamic definition of adjacency (`adj`), as well as more general concepts, such as *areas*, which are mentioned in Section 4 and *interfaces*. The latter may require a small adaptation of T-AWN.

(b) Moreover, we want to validate our models against real implementations. One of the main problems when creating formal specifications for, or implementations of routing protocols is that they always represent the personal view of the developer(s). This is due to the fact that standards are written in English prose. We plan to automatically compare our Uppaal model with off-the-shelf implementations of OSPF, such as the well-known Quagga implementation [33, 9], running on the network emulator CORE [1].

(c) Last, but certainly not least, we are going to analyse OSPF. This analysis will include classic properties such as the avoidance of routing loops, but also security aspects. For the latter we have to develop formal attack models.

Acknowledgments This work was conducted in partnership with the Defence Science & Technology Group and Data61, CSIRO, through the *Next Generation Technologies Fund*.

[4]CBMC is a bounded model checker accepting a simplified C-language as input.

References

[1] J. Ahrenholz, C. Danilov, T. Henderson & J. Kim (2008): *CORE: A Real-Time Network Emulator*. doi:10.1109/MILCOM.2008.4753614.

[2] G. Behrmann, A. David & K. G. Larsen (2004): *A Tutorial on* UPPAAL. In Marco Bernardo & Flavio Corradini, editors: *Formal Methods for the Design of Real-Time Systems*, *Lecture Notes in Computer Science* 3185, Springer, pp. 200–236, doi:10.1007/978-3-540-30080-9_7.

[3] R. Bellman (1958): *On a Routing Problem*. *Quarterly of Applied Mathematics* 16, pp. 87–90, doi:10.1090/qam/102435.

[4] J. A. Bergstra & J. W. Klop (1986): *Algebra of Communicating Processes*. In J. W. de Bakker, M. Hazewinkel & J. K. Lenstra, editors: *Mathematics and Computer Science*, CWI Monograph 1, North-Holland, pp. 89–138.

[5] J. Borgström, S. Huang, M. Johansson, P. Raabjerg, B. Victor, J. Å Pohjola & J. Parrow (2011): *Broadcast Psi-calculi with an Application to Wireless Protocols*. In G. Barthe, A. Pardo & G. Schneider, editors: *Software Engineering and Formal Methods (SEFM'11)*, *Lecture Notes in Computer Science* 7041, Springer, pp. 74–89, doi:10.1007/978-3-642-24690-6_7.

[6] E. Bres, R. van Glabbeek & P. Höfner (2016): *A Timed Process Algebra for Wireless Networks with an Application in Routing (Extended Abstract)*. In P. Thiemann, editor: *Programming Languages and Systems (ESOP'16)*, *Lecture Notes in Computer Science* 9632, Springer, pp. 95–122, doi:10.1007/978-3-662-49498-1_5.

[7] E. Bres, R. van Glabbeek & R. Höfner (2016): *A Timed Process Algebra for Wireless Networks*. *CoRR* abs/1606.03663. Available at `http://arxiv.org/abs/1606.03663`.

[8] E. Clarke, D. Kroening & F. Lerda (2004): *A Tool for Checking ANSI-C Programs*. In K. Jensen & A. Podelski, editors: *Tools and Algorithms for the Construction and Analysis of Systems (TACAS'04)*, *Lecture Notes in Computer Science* 2988, Springer, pp. 168–176, doi:10.1007/978-3-540-24730-2_15.

[9] R. Coltun, D. Ferguson, J. Moj & A. Lindem (2008): *OSPF for IPv6*. RFC 5340, Network Working Group. Available at `http://www.ietf.org/rfc/rfc5340.txt`.

[10] L. De Moura & N. Bjørner (2008): *Z3: An Efficient SMT Solver*. In C. R. Ramakrishnan & J. Rehof, editors: *Tools and Algorithms for the Construction and Analysis of Systems (TACAS'08)*, *Lecture Notes in Computer Science* 4963, Springer, pp. 337–340, doi:10.1007/978-3-540-78800-3_24.

[11] E. W. Dijkstra (1959): *A Note on Two Problems in Connexion with Graphs*. *Numerische Mathematik* 1(1), pp. 269–271, doi:10.1007/BF01386390.

[12] C. Ene & T. Muntean (2001): *A Broadcast-based Calculus for Communicating Systems*. In: *Parallel & Distributed Processing Symposium (IPDPS '01)*, IEEE Computer Society, pp. 1516–1525, doi:10.1109/IPDPS.2001.925136.

[13] A. Fehnker, R. J. van Glabbeek, P Höfner, A. K. McIver, M. Portmann & W. L. Tan (2012): *Automated Analysis of AODV using UPPAAL*. In C. Flanagan & B. König, editors: *Tools and Algorithms for the Construction and Analysis of Systems (TACAS '12)*, *Lecture Notes in Computer Science* 7214, Springer, pp. 173–187, doi:10.1007/978-3-642-28756-5_13.

[14] A. Fehnker, R. J. van Glabbeek, P. Höfner, A. K. McIver, M. Portmann & W. L. Tan (2012): *A Process Algebra for Wireless Mesh Networks*. In H. Seidl, editor: *European Symposium on Programming* (ESOP '12), *Lecture Notes in Computer Science* 7211, Springer, pp. 295–315, doi:10.1007/978-3-642-28869-2_15.

[15] A. Fehnker, R. J. van Glabbeek, P. Höfner, A. K. McIver, M. Portmann & W. L. Tan (2013): *A Process Algebra for Wireless Mesh Networks used for Modelling, Verifying and Analysing AODV*. Available at `http://arxiv.org/abs/1312.7645`.

[16] L. R. Ford Jr. (1956): *Network Flow Theory*. Technical Report Paper P-923, The RAND Corporation.

[17] F. Ghassemi, W. Fokkink & A. Movaghar (2008): *Restricted Broadcast Process Theory*. In A. Cerone & S. Gruner, editors: *Software Engineering and Formal Methods (SEFM '08)*, IEEE Computer Society, pp. 345–354, doi:10.1109/SEFM.2008.25.

[18] R. J. van Glabbeek & P. Höfner (2017): *Split, Send, Reassemble: A Formal Specification of a CAN Bus Protocol Stack*. In H. Hermanns & P. Höfner, editors: *Models for Formal Analysis of Real Systems (MARS'17)*, *Electronic Proceedings in Theoretical Computer Science* 244, Open Publishing Association, pp. 14–52, doi:10.4204/EPTCS.244.2.

[19] R. J. van Glabbeek, P. Höfner, W. L. Tan & M. Portmann (2013): *Sequence Numbers Do Not Guarantee Loop Freedom —AODV Can Yield Routing Loops—*. In: *Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM '13)*, ACM, pp. 91–100, doi:10.1145/2507924.2507943.

[20] R. J. van Glabbeek, P. Höfner, M. Portmann & W.L. Tan (2016): *Modelling and Verifying the AODV Routing Protocol*. *Distributed Computing* 29(4), pp. 279–315, doi:10.1007/s00446-015-0262-7.

[21] J. C. Godskesen (2007): *A Calculus for Mobile Ad Hoc Networks*. In A. L. Murphy & J. Vitek, editors: *Coordination Models and Languages (COORDINATION '07)*, *Lecture Notes in Computer Science* 4467, Springer, pp. 132–150, doi:10.1007/978-3-540-72794-1_8.

[22] J. C. Godskesen (2010): *Observables for Mobile and Wireless Broadcasting Systems*. In D. Clarke & G. A. Agha, editors: *Coordination Models and Languages (COORDINATION '10)*, *Lecture Notes in Computer Science* 6116, Springer, pp. 1–15, doi:10.1007/978-3-642-13414-2_1.

[23] C. A. R. Hoare (1985): *Communicating Sequential Processes*. Prentice Hall, Englewood Cliffs.

[24] ISO/IEC/IEEE 8802-11 (2018): *Information Technology — Telecommunications and information exchange between systems — Local and metropolitan area networks — Specific requirements — Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications*. Available at `https://www.iso.org/standard/73367.html`.

[25] K. G. Larsen, P. Pettersson & Wang Yi (1997): *UPPAAL in a Nutshell*. *International Journal of Software Tools for Technology Transfer* 1(1-2), pp. 134–152, doi:10.1007/s100090050010.

[26] S.U.R Malik, S. K. Srinivasan & S. U. Khan (2012): *A Methodology for OSPF Routing Protocol Verification*.

[27] M. Merro (2009): *An Observational Theory for Mobile Ad Hoc Networks (full version)*. *Information and Computation* 207(2), pp. 194–208, doi:10.1016/j.ic.2007.11.010.

[28] N. Mezzetti & D. Sangiorgi (2006): *Towards a Calculus For Wireless Systems*. *Electronic Notes in Theoretical Computer Science* 158, pp. 331–353, doi:10.1016/j.entcs.2006.04.017.

[29] R. Milner (1980): *A Calculus of Communicating Systems*. *Lecture Notes in Computer Science* 92, Springer, doi:10.1007/3-540-10235-3.

[30] R. Milner (1989): *Communication and Concurrency*. Prentice Hall, Englewood Cliffs.

[31] S. Miskovic & E. W. Knightly (2010): *Routing Primitives for Wireless Mesh Networks: Design, Analysis and Experiments*. In: *INFOCOM'10*, IEEE, pp. 2793–2801, doi:10.1109/INFCOM.2010.5462111.

[32] K. Möhring (2017): *Transformation of AWN Protocol Specifications to the Uppaal Model Checker*. Bachelor Thesis, Universität Hamburg and Data61, CSIRO.

[33] J. Moj (1998): *OSPF Version 2*. RFC 2328, Network Working Group. Available at `http://www.ietf.org/rfc/rfc2328.txt`.

[34] E. F. Moore (1957): *The Shortest Path Through a Maze*. Technical Report, Bell Telephone System.

[35] G. Nakibly, A. Sosnovich, E. Menahem, A. Waizel & Y. Elovici (2014): *OSPF Vulnerability to Persistent Poisoning Attacks: A Systematic Analysis*. In: *Computer Security Applications Conference*, ACSAC '14, ACM, pp. 336–345, doi:10.1145/2664243.2664278.

[36] S. Nanz & C. Hankin (2006): *A Framework for Security Analysis of Mobile Wireless Networks*. *Theoretical Computer Science* 367, pp. 203–227, doi:10.1016/j.tcs.2006.08.036.

[37] K. V. S. Prasad (1991): *A Calculus of Broadcasting Systems*. In S. Abramsky & T. S. E. Maibaum, editors: *Theory and Practice of Software Development (TAPSOFT '91)*, *Lecture Notes in Computer Science* 493, Springer, pp. 338–358, doi:10.1007/3-540-53982-4_19.

[38] K. V. S. Prasad (1995): *A Calculus of Broadcasting Systems*. *Science of Computer Programming* 25(2-3), pp. 285–327, doi:10.1016/0167-6423(95)00017-8.

[39] A. Shimbel (1955): *Structure in Communication Nets*. In: *Symposium on Information Networks*, Polytechnic Press of the Polytechnic Institute of Brooklyn, pp. 199–203.

[40] A. Singh, C. R. Ramakrishnan & S. A. Smolka (2010): *A process calculus for Mobile Ad Hoc Networks*. *Science of Computer Programming* 75, pp. 440–469, doi:10.1016/j.scico.2009.07.008.

[41] K. Varadhan, R. Govindan & D. Estrin (2000): *Persistent Route Oscillations in Inter-domain Routing*. *Computer Networks* 32(1), pp. 1–16, doi:10.1016/S1389-1286(99)00108-5.

[42] P. Zave (2011): *Experiences with Protocol Description*. In: *Workshop on Rigorous Protocol Engineering (W-RiPE'11)*.

A Informal Description of the Process Expressions of (T-)AWN[5]

In this appendix we describe the *process expressions* of (T-)AWN, given in Table 1.

A process name X comes with a *defining equation*

$$X(\mathtt{var}_1,\ldots,\mathtt{var}_n) \stackrel{\mathit{def}}{=} P\,,$$

where P is a process expression, and the $\mathtt{var}_i$ are data variables maintained by process X. A named process is like a *procedure*; when it is called, data expressions exp_i of the appropriate type are filled in for the variables $\mathtt{var}_i$. Furthermore, φ is a condition, $\mathtt{var}:=exp$ an assignment of a data expression exp to a variable var of the same type, *dest*, *dests*, *data* and *ms* data expressions of types IP, $\mathscr{P}(\mathtt{IP})$, DATA and MSG, respectively, and msg a data variable of type MSG.

Given a valuation of the data variables by concrete data values, the process $[\varphi]P$ acts as P if φ evaluates to true, and deadlocks if φ evaluates to false.[6] In case φ contains free variables that are not yet interpreted as data values, values are assigned to these variables in any way that satisfies φ, if possible. The process $[\![\mathtt{var}:=exp]\!]P$ acts as P, but under an updated valuation of the data variables. The process $P+Q$ may act either as P or as Q, depending on which of the two is able to act at all. In a context where both are able to act, it is not specified how the choice is made. The process **broadcast**$(ms).P$ broadcasts (the data value bound to the expression) *ms* to all connected components (routers), and subsequently acts as P, whereas the process **unicast**$(dest,ms).P \blacktriangleright Q$ tries to unicast the message *ms* to the destination *dest*; if successful it continues to act as P and otherwise as Q. The process **groupcast**$(dests,ms).P$ tries to transmit *ms* to all destinations *dests*, and proceeds as P regardless of whether any of the transmissions is successful. The process **receive**(msg)$.P$ receives any message m (a data value of type MSG) either from another network node, from another process running on the same node or from an application layer process connected to that component. It then proceeds as P, but with the data variable msg bound to the value m. In particular, **receive**$(\mathtt{newpkt}(id,data))$ models the injection of data from the application layer, where the function newpkt generates a message containing the application layer *data* and the identifier *id*, here indicating the message type. Data is delivered to the application layer by **deliver**$(data)$.

The internal state of a sequential process described by an expression P in this language is determined by P, together with a *valuation* ξ associating data values $\xi(\mathtt{var})$ to the data variables var maintained by this process. In case a process maintains no data values, we use the empty valuation ξ_0. A *valuated process* is a pair (ξ,P) of a sequential process P and an initial valuation ξ.

Finally, $V \langle\!\langle W$ denotes a parallel composition of valuated processes V and W, with information piped from right to left; in typical applications (e.g. ours or [20]) W is a message queue.

In the full process algebra [14], *node expressions* $id\!:\!V\!:\!R$ are given by process expressions V, annotated with a (unique) process identifier *id* and a set of nodes R that are connected to *id*.

A partial network is then modelled as a parallel composition of node expressions, using the operator $\|$, and a complete network is obtained by placing this composition in the scope of an encapsulation operator $[_]$. The main purpose of the encapsulation operator is to prevent the receipt of messages that have never been sent by other nodes in the network—with the exception of messages $\mathtt{newpkt}(data,dip)$ stemming from the application layer of a node.

---

[5] This appendix (in nearly identical form) also appeared in [18].

[6] As operators we also allow *partial* functions with the convention that any atomic formula containing an undefined subterm evaluates to false.

B The Simple T-AWN Model

We present all details concerning our simple T-AWN model, which we sketched in Section 5. Although we created the detailed model first, we believe that presenting the simple model in detail first is easier for the reader as she can concentrate on the core aspects of OSPF. When presenting the detailed model in Appendix C we can concentrate on the differences, which makes the description shorter.

B.1 Data Structure

We describe the data structures required for our formal specification of OSPF. First we define a number of types that are used to handle and process information. Based on these types, we then define functions. Finally, we describe the exact intention and effects of the various operations performed by each node.

Our formal specification of OSPF focuses on adjacency establishment between two routers and the LSA flooding mechanism. As a node receives messages from its neighbours it will update its *neighbour list*, which documents its conversations with other nodes. After new neighbours or other topological changes are discovered, information of the local topology will be updated and shared in the form of LSAs. Both the *neighbour structures* and *link state databases* are key components of the data structure.

B.1.1 Mandatory Types of (T-)AWN

Messages are used to send information via the network. In our specification we use the variable msg of type MSG. We distinguish four types of OSPF control messages: HELLO messages, DBD messages, LSR messages, and LSU messages. We also use an internal message for intranode communication.

The type IP describes a set of IP addresses or, more generally, *a set of node identifiers*. We assume that each node has a unique identifier $\mathrm{ip} \in \mathtt{IP}$. Each node maintains a variable ip, which always stores the node's identifier. Every message contains the variable sip that holds the sender's identifier. Furthermore, nip denotes an identifier of an arbitrary neighbour.

Lastly, T-AWN provides the variable now of the type TIME, which is local to a node and increases whenever time passes.

B.1.2 Link State Advertisements (LSAs) & Link State Databases

Link State Advertisements (LSAs) are one of the key components of OSPF. They are used to describe how routers and networks are interconnected and will be advertised through parts of the network. The RFC defines five distinct types of LSAs. For the simple model we assume an area with Point-to-Point networks only; thus only Router-LSAs exist in our model. A *Router*-LSA describes the outgoing 'links' belonging to a given router, which effectively describes the connections that this router provides.

Every router periodically generates and distributes a Router-LSA (for brevity we refer to it as an LSA) to describe its knowledge about the local network topology. An LSA is given by three components:

1. The IP address of the originator, which is an element of IP.
2. A timestamp denoting when the LSA was generated, which is an element of TIME.[7]
3. A set of IP addresses representing all outgoing links of the originator (discovered so far), which is an element of $\mathscr{P}(\mathtt{IP})$.

We denote the type of LSA by LSA and define a function $(_,_,_) : \mathtt{IP} \times \mathtt{TIME} \times \mathscr{P}(\mathtt{IP}) \to \mathtt{LSA}$ to generate LSAs. An LSA is uniquely identified by its first two components, the LSA *header*. The type of headers

[7] We assume the type TIME to be isomorphic to the integers, which allows us to use operations such as addition or subtraction.

is denoted by LSAHDR. We define a generator function $(_,_) : \mathtt{IP} \times \mathtt{TIME} \to \mathtt{LSAHDR}$ and an extraction function to get the header from an LSA:

$$\begin{array}{rcl} \mathtt{hdr} : \mathtt{LSA} & \to & \mathtt{LSAHDR} \\ \mathtt{hdr}(\mathrm{lsa}) & = & (\pi_1(\mathrm{lsa}), \pi_2(\mathrm{lsa}))\,.^8 \end{array}$$

When the protocol encounters two LSAs with the same originator, it needs to decide which one is more recent. Therefore, we define a partial order on LSAHDR.

$$\mathrm{lhdr}_1 \leq \mathrm{lhdr}_2 \quad :\Longleftrightarrow \quad \pi_1(\mathrm{lhdr}_1) = \pi_1(\mathrm{lhdr}_2) \wedge \pi_2(\mathrm{lhdr}_1) \leq \pi_2(\mathrm{lhdr}_2)$$

As usual, we define a strict partial order by $\mathrm{lhdr}_1 < \mathrm{lhdr}_2 :\Longleftrightarrow \mathrm{lhdr}_1 \leq \mathrm{lhdr}_2 \wedge \mathrm{lhdr}_1 \neq \mathrm{lhdr}_2$. The term $\mathtt{hdr}(lsa_1) < \mathtt{hdr}(lsa_2)$ indicates that lsa_2 is more recent than lsa_1 (lsa_2 was generated later) and should replace lsa_1 in the link state database to keep it up to date.

A *link state database* is a set of LSAs, where each LSA is uniquely identified by its originator. The type, denoted by LSDB, is defined as

$$\mathtt{LSDB} := \{lsdb \mid lsdb \in \mathscr{P}(\mathtt{LSA}) \ \wedge \ \forall lsa_1, lsa_2 \in lsdb : lsa_1 \neq lsa_2 \Rightarrow \mathtt{hdr}(lsa_1) \not\leq \mathtt{hdr}(lsa_2)\}\,.$$

Every router maintains a link state database and uses it to calculate the routing table. models do not consider routing table calculations and only focus on synchronisation of link state databases. Note that this data type is equivalent to

$$\{lsdb \mid lsdb \in \mathscr{P}(\mathtt{LSA}) \ \wedge \ \forall lsa_1, lsa_2 \in lsdb : lsa_1 \neq lsa_2 \Rightarrow \pi_1(lsa_1) \neq \pi_1(lsa_2)\}\,.$$

B.1.3 Data Structure Relating Neighbours

A router stores a set of neighbour structures so that it can keep track of how much and whether it needs to exchange information with each neighbour. Since we assume point-to-point networks it means that routers must exchange information with all of their neighbours. This fact simplifies the structure greatly when compared to the detailed model.

A *neighbour structure* contains two components:

1. The IP address of the neighbour, which is an element of IP.
2. A timestamp of type TIME, which indicates when the neighbour is considered inactive or dead—unless another message is received from this neighbour before that time.

NBR denotes the type of neighbour structures and an element is generated by $(_,_) : \mathtt{IP} \times \mathtt{TIME} \to \mathtt{NBR}$.

A *neighbour list* is a set of neighbour structures, where each IP address has at most one entry. These correspond to the Interface Data Structure in the RFC.

$$\mathtt{NBRS} := \{nbrs \mid nbrs \in \mathscr{P}(\mathtt{NBR}) \ \wedge \ \forall nbr_1, nbr_2 \in nbrs : nbr_1 \neq nbr_2 \Rightarrow \pi_1(nbr_1) \neq \pi_1(nbr_2)\}\,.$$

B.1.4 Messages

The simple model uses four types of OSPF messages:

1. A HELLO message contains two entries: (a) a set of IP addresses of all neighbours discovered by the sender and (b) the identifier of the sender itself. The generation function is
$$\mathtt{hello} : \mathscr{P}(\mathtt{IP}) \times \mathtt{IP} \to \mathtt{MSG}\,.$$
2. A DBD message consists of two fields: (a) a set of LSA headers describing the LSDB of the sender (an element of $\mathscr{P}(\mathtt{LSAHDR})$) and (b) the sender itself. The generation function is
$$\mathtt{dbd} : \mathscr{P}(\mathtt{LSAHDR}) \times \mathtt{IP} \to \mathtt{MSG}\,.$$

[8] We use projection functions π_i to distill the i-th component of a tuple.

3. An LSR message carries two items as payload: (1) a set of LSA headers representing the LSAs required and, again, (2) the identifier of the sender itself. The generation function is
$$\texttt{req} : \mathscr{P}(\texttt{LSAHDR}) \times \texttt{IP} \to \texttt{MSG}\,.$$
4. An LSU message consists of two fields: (a) an LSDB and (b) the identifier of the sender. The generation function is
$$\texttt{upd} : \texttt{LSDB} \times \texttt{IP} \to \texttt{MSG}\,.$$

Sending messages takes time, but should not result in the OSPF process getting stuck. Therefore we assume the existence of queues to buffer incoming messages and a sending process, QSND, to handle the sending of messages to other routers (in some sense this is an output queue).

This assumption reflects modern hardware architecture, but bears a problem in modelling. Messages are unicast to a single destination, groupcast (mulitcast) to a set of destinations or broadcast. This information (the sending method and the destinations) is not contained within the OSPF control messages; it needs to be passed onto the sending process QSND.

We define a new (internal) message type that combines the sending method and the destinations with outgoing messages. This new message type consists of two fields: (a) the real message that needs to be sent, which is an element of MSG, and (b) a set of destination IPs, which is an element of $\mathscr{P}(\texttt{IP})$. As usual we define a generator function
$$\texttt{sndmsg} : \texttt{MSG} \times \mathscr{P}(\texttt{IP}) \to \texttt{MSG}\,.$$
The main process OSPF (and its subprocesses) use this message type to pass messages to the process QSND, which manages the actual sending.

B.1.5 Update Functions

We require a set of manipulation functions to change the local data structures such as the LSDB.

When a router discovers a new neighbour it initiates a process to synchronise its own Link State Databases (LSDB) with the LSDB maintained by the neighbour. It sends a *database description (*DBD*) message* that contains the headers of all the LSAs in its own LSDB to the newly discovered neighbour. When the neighbour receives this message it compares all the LSA headers against the entries in its own LSDB. If one is missing or out-dated, it sends a request for that LSA. The newly discovered neighbour also sends a DBD message to the original router, which triggers a symmetric process.

We generate a new LSA for node ip at time t, using a neighbour structure nbrs by

$$\begin{array}{rcl}\texttt{newLSA} : \texttt{IP} \times \texttt{TIME} \times \texttt{NBRS} & \to & \texttt{LSA}\\ \texttt{newLSA}(\mathrm{ip},\mathrm{t},\mathrm{nbrs}) & := & (\mathrm{ip},\mathrm{t},\Pi_1(\mathrm{nbrs})\})\,,\end{array}$$

where $\Pi_i(S) := \{\pi_i(s) \mid s \in S\}$ is the pointwise lifted function π_i.

To update a node's LSDB with a given set of LSAs we use

$$\begin{array}{rcl}\texttt{install} : \texttt{LSDB} \times \texttt{LSDB} & \to & \texttt{LSDB}\\ \texttt{install}(\mathrm{lsdb},\mathrm{lsas}) & := & \{lsa \mid lsa \in \mathrm{lsdb} \wedge \nexists lsa' \in \mathrm{lsas} :\ \texttt{hdr}(lsa) < \texttt{hdr}(lsa')\}\\ & & \cup\, \{lsa \mid lsa \in \mathrm{lsas} \wedge \nexists lsa' \in \mathrm{lsdb} :\ \texttt{hdr}(lsa) \leq \texttt{hdr}(lsa')\}\,.\end{array}$$

It is easy to see that this definition is well defined, i.e. returns an element of type LSDB
The protocol also requires functions to inspect elements of the neighbour list, or to modify it.

1. To check whether an entry for nip exists in nbrs we use

$$\begin{aligned} \texttt{nbrExist} : \texttt{NBRS} \times \texttt{IP} &\rightarrow \texttt{BOOL} \\ \texttt{nbrExist}(\text{nbrs}, \text{nip}) &:\Longleftrightarrow \text{nip} \in \Pi_1(\text{nbrs})\,. \end{aligned}$$

2. To insert a new neighbour structure for node nip we use the partial function

$$\begin{aligned} \texttt{newNBR} : \texttt{NBRS} \times \texttt{IP} &\rightharpoonup \texttt{NBRS} \\ \texttt{newNBR}(\text{nbrs}, \text{nip}) &:= \text{nbrs} \cup \{(\text{nip}, 0)\} \qquad \text{if } \neg\texttt{nbrExist}(\text{nbrs}, \text{nip})\,. \end{aligned}$$ [9]

3. We update the inactivity timer of an entry of nbrs by

$$\begin{aligned} \texttt{setINACTT} : \texttt{NBRS} \times \texttt{IP} \times \texttt{TIME} &\rightarrow \texttt{NBRS} \\ \texttt{setINACTT}(\text{nbrs}, \text{nip}, \text{t}) &:= \{n \mid n \in \text{nbrs} \wedge \pi_1(n) \neq \text{nip}\} \cup \{(\text{nip}, \text{t})\}\,. \end{aligned}$$

4. Given the current time t, nodes need to know which neighbours from nbrs are inactive.

$$\begin{aligned} \texttt{deadNBRS} : \texttt{NBRS} \times \texttt{TIME} &\rightarrow \texttt{NBRS} \\ \texttt{deadNBRS}(\text{nbrs}, \text{t}) &:= \{n \mid n \in \text{nbrs} \wedge \pi_2(n) < \text{t}\}\,. \end{aligned}$$

B.2 The Full Model

Our simple model, written in the process algebra T-AWN, consists of 7 processes, named OSPF, HELLO, DBD, REQ, UPD, QMSG and QSND.

- The main process OSPF reads a single message from the input queue. Depending on the type of the message it calls other processes to handle further operations. The process also sends HELLO messages at periodic intervals and maintains the node's LSDB by removing dead neighbours.
- The process HELLO describes all the actions performed when a HELLO message is received. This includes updating the relevant inactivity timer and, if the HELLO message is from a previously unknown neighbour, updating the node's neighbour list. The process also generates a DBD message, which is sent back to the originator of the HELLO message, and a new LSA when required.
- The process DBD handles incoming DBD messages. Sending out a request for those LSAs that are described in a DBD message, but do not have a matching partner in the node's LSDB is among its actions.
- The process REQ describes the actions following the receipt of an LSR message, such as finding the requested LSAs in the node's local data, and sending them back to the sender of the message.
- The process UPD manages incoming LSU messages, including the installation of up-to-date LSAs in the node's LSDB, and broadcasting the updated information.
- The process QMSG models a message queue. In our model we have two instances. Whenever a message is received by a node, it is first stored in the first instance, the message input queue. If the corresponding node—or more precisely the process OSPF—is able to handle a message it pops the oldest message from the input queue. Whenever a message is sent by the process OSPF, it is passed to the second instance of QMSG, which holds the message until the process QSND can handle it.
- The process QSND is responsible for sending outgoing messages that were generated by the other processes described above.

In the remainder of this section we present formal specifications of each of these processes and explain them step by step.

[9] For brevity we omit the cases that are undefined.

B.2.1 The Basic Routine

The basic process OSPF (Process 1) either reads a message from its queue, sends a HELLO message or removes dead neighbours from its neighbour list. It maintains four variables: ip, nbrs, lsdb and hello_t, in which it stores its own IP address, information about its immediate neighbours, its link state database, and a timestamp stating when the next hello message should be sent.

Process 1 The Basic Routine

$\texttt{OSPF}(\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t}) \stackrel{def}{=}$

1. **receive**(msg) .
2. /* depending on the message, the node calls different processes */
3. (
4. [$\mathrm{msg} = \texttt{hello}(\mathrm{ips},\mathrm{sip})$] /* HELLO message received */
5. $\texttt{HELLO}(\mathrm{ips},\mathrm{sip},\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$
6. + [$\mathrm{msg} = \texttt{dbd}(\mathrm{lsa_hdrs},\mathrm{sip})$] /* DBD message received */
7. $\texttt{DBD}(\mathrm{lsa_hdrs},\mathrm{sip},\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$
8. + [$\mathrm{msg} = \texttt{req}(\mathrm{lsa_hdrs},\mathrm{sip})$] /* LSR message received */
9. $\texttt{REQ}(\mathrm{lsa_hdrs},\mathrm{sip},\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$
10. + [$\mathrm{msg} = \texttt{upd}(\mathrm{lsas},\mathrm{sip})$] /* LSU message received */
11. $\texttt{UPD}(\mathrm{lsas},\mathrm{sip},\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$
12.)
13. + [$\texttt{now} \geq \mathrm{hello_t}$] /* send HELLO message */
14. $[\![\mathrm{hello_t} := \texttt{now} + \texttt{hello_intvl}]\!]$
15. **send**$(\texttt{sndmsg}(\texttt{hello}(\{\pi_1(n) \mid n \in \mathrm{nbrs}\},\mathrm{ip}),\emptyset))$.
16. $\texttt{OSPF}(\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$
17. + [$\texttt{deadNBRS}(\mathrm{nbrs},\texttt{now}) \neq \emptyset$] /* inactive neighbours, which should be removed, exist */
18. $[\![\mathrm{nbrs} := \mathrm{nbrs} - \texttt{deadNBRS}(\mathrm{nbrs},\texttt{now})]\!]$ /* remove corresponding neighbour structures */
19. $[\![\mathrm{lsa} := \texttt{newLSA}(\mathrm{ip},\texttt{now},\mathrm{nbrs})]\!]$ /* generate a new LSA */
20. $[\![\mathrm{lsdb} := \texttt{install}(\mathrm{lsdb},\{\mathrm{lsa}\})]\!]$ /* install it */
21. **send**$(\texttt{sndmsg}(\texttt{upd}(\{\mathrm{lsa}\},\mathrm{ip}),\{\pi_1(n) \mid n \in \mathrm{nbrs}\}))$. /* and flood it out */
22. $\texttt{OSPF}(\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$

The message handling is described in Lines 1 to 12. The process first receives a message by **receive**(msg); by our specification of the input queue (see Process 6 below) this will be the message currently in the front of the queue (the oldest message). OSPF then checks the type of the message and calls the corresponding process: in case of a HELLO message the process HELLO is called, in case of a DBD message the process DBD, etc.

The second part of OSPF (Lines 13 to 16) ensures that HELLO messages are sent periodically. The parameter hello_t indicates when the next HELLO message should be sent. The process checks whether it is time to send the next HELLO message; if so it resets the timer (Line 14) and broadcasts a HELLO message. In fact, the message is sent to the process QSND (Process 7), which handles the actual sending. The HELLO message contains the IP addresses of all the neighbours that the node ip is aware of. After handing over the HELLO message the node returns to the main process OSPF to handle the next incoming message, or send yet another HELLO message.

The third part of OSPF (Lines 17 to 22) collects all of the neighbours from which no activity has been seen for the last `rtdead_intvl` seconds, removes them from the neighbour structure nbrs and

informs the neighbours about these updates. Line 17 checks whether there exist neighbours from which no activity has been seen for `rtdead_intvl` seconds, using the function `deadNBRS`. If such neighbours exist, the node removes all of the dead IP addresses from the neighbour structure (Line 18). Then the node generates a new LSA with this information (Line 19), installs this information in its own link state database (Line 20) and sends out the information to its neighbours (Line 21).

B.2.2 Message Processing

The process `HELLO` (Process 2) handles the receipt of a HELLO message.

Process 2 Handling HELLO messages

$\texttt{HELLO}(\mathrm{ips},\mathrm{sip},\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t}) \stackrel{def}{=}$

1. $[\ \neg\texttt{nbrExist}(\mathrm{nbrs},\mathrm{sip})\ ]$ /* the sender sip is unknown */
2. $[\![\mathrm{nbrs} := \texttt{newNBR}(\mathrm{nbrs},\mathrm{sip})]\!]$
3. $[\![\mathrm{nbrs} := \texttt{setINACTT}(\mathrm{nbrs},\mathrm{sip},\texttt{now}+\texttt{rtdead_intvl})]\!]$
4. $[\![\mathrm{lsa} := \texttt{newLSA}(\mathrm{ip},\texttt{now},\mathrm{nbrs})]\!]$ /* generate new LSA */
5. $[\![\mathrm{lsdb} := \texttt{install}(\mathrm{lsdb},\{\mathrm{lsa}\})]\!]$ /* install LSA */
6. $\textbf{send}(\texttt{sndmsg}(\texttt{upd}(\{\mathrm{lsa}\},\mathrm{ip}),\{\pi_1(n) \mid n \in \mathrm{nbrs}\}))\ .$ /* send LSA */
7. $\textbf{send}(\texttt{sndmsg}(\texttt{dbd}(\{\texttt{hdr}(lsa) \mid lsa \in \mathrm{lsdb}\},\mathrm{ip}),\{\mathrm{sip}\}))\ .$ /* send DBD back to sip */
8. $\texttt{OSPF}(\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$
9. $+\ [\ \texttt{nbrExist}(\mathrm{nbrs},\mathrm{sip})\ ]$ /* sip is a known neighbour */
10. $[\![\mathrm{nbrs} := \texttt{setINACTT}(\mathrm{nbrs},\mathrm{sip},\texttt{now}+\texttt{rtdead_intvl})]\!]$
11. $\texttt{OSPF}(\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$

The first part of the process (Lines 1 to 8) is executed when the originator of the HELLO message is unknown to the node ip. In Line 1, the function `nbrExist` returns a Boolean indicating whether the sender sip of the HELLO message is in the node's neighbour list nbrs (i.e. whether it is known to the node ip). In Lines 2 to 5 the process adds the sender to its neighbour structure (Line 2), resets the inactivity timer for that particular neighbour (Line 3), generates an LSA to reflect the change to the neighbour structure and installs that LSA in its link state database. Line 6 sends out the newly created LSA in the form of an LSU message, which will be sent to all of the known neighbours by the process `QSND`. Line 7 creates and sends a DBD message, which contains all the headers of the LSAs in the link state database. The message is intended for the newly discovered neighbour sip, the sender of the HELLO message.

The second part of the process (Lines 9 to 11) handles the situation where the sender sip of the HELLO message is already known to the node. All that needs to be done is a reset of the inactivity timer (Line 10), using the function `setINACTT`. After this it returns to the main process `OSPF`.

The process `DBD` (Process 3) handles the receipt of a DBD message. It behaves similar to the process `HELLO` as it checks whether the sender of the incoming message is known to the node or not.

The first part of the process (Lines 1 to 8) deals with a message from an unknown sender. Lines 1–7 are identical to the corresponding lines of Process 2: Line 2 adds the sender to the neighbour list nbrs, Line 3 resets the neighbour's inactivity timer, Line 4 generates a new LSA to reflect the new state of the lsdb and Line 5 installs the new LSA in the link state database. Line 6 places the newly created LSA in an update message and sends it out to all known neighbours. Line 7 sends a DBD message back to the sender sip of the original message, informing the node about the contents of the node's lsdb. After this, the process returns to the start of `DBD` again, this time the sender is known, as it was added to the neighbour structure nbrs on Line 2. This time around the second part of the process is executed.

Process 3 Handling DBD messages

$\texttt{DBD}(\text{lsa_hdrs}, \text{sip}, \text{ip}, \text{nbrs}, \text{lsdb}, \text{hello_t}) \stackrel{def}{=}$
1. $[\, \neg\texttt{nbrExist}(\text{nbrs}, \text{sip})\,]$ /* the sender sip is unknown */
2. $[\![\text{nbrs} := \texttt{newNBR}(\text{nbrs}, \text{sip})]\!]$
3. $[\![\text{nbrs} := \texttt{setINACTT}(\text{nbrs}, \text{sip}, \texttt{now} + \texttt{rtdead_intvl})]\!]$
4. $[\![\text{lsa} := \texttt{newLSA}(\text{ip}, \texttt{now}, \text{nbrs})]\!]$ /* generate new LSA */
5. $[\![\text{lsdb} := \texttt{install}(\text{lsdb}, \{\text{lsa}\})]\!]$ /* install LSA */
6. $\textbf{send}(\texttt{sndmsg}(\texttt{upd}(\{\text{lsa}\}, \text{ip}), \{\pi_1(n) \mid n \in \text{nbrs}\}))$. /* send LSA */
7. $\textbf{send}(\texttt{sndmsg}(\texttt{dbd}(\{\texttt{hdr}(lsa) \mid lsa \in \text{lsdb}\}, \text{ip}), \{\text{sip}\}))$. /* send DBD back to sip */
8. $\texttt{DBD}(\text{lsa_hdrs}, \text{sip}, \text{ip}, \text{nbrs}, \text{lsdb}, \text{hello_t})$
9. $+\, [\,\texttt{nbrExist}(\text{nbrs}, \text{sip})\,]$ /* sip is a known neighbour */
10. $\textbf{send}(\texttt{sndmsg}(\texttt{req}(\{lhdr \mid lhdr \in \text{lsa_hdrs} \wedge \nexists lsa \in \text{lsdb} : lhdr \leq \texttt{hdr}(lsa)\}, \text{ip}), \{\text{sip}\}))$.
11. $\texttt{OSPF}(\text{ip}, \text{nbrs}, \text{lsdb}, \text{hello_t})$

The second part (Lines 9 to 11) deals with the situation when the sender of the DBD message is already known. In this case the LSA headers lsa_hdrs from the DBD message are compared with the node's own link state database lsdb. All the headers that are newer than those in `lsdb` or that do not yet exist in `lsdb` are packed into an LSR message and sent back to sip (Line 10).

The process `REQ` (Process 4) handles the receipt of an LSR message.

Process 4 Handling LSR messages

$\texttt{REQ}(\text{lsa_hdrs}, \text{sip}, \text{ip}, \text{nbrs}, \text{lsdb}, \text{hello_t}) \stackrel{def}{=}$
1. $[\, \neg\texttt{nbrExist}(\text{nbrs}, \text{sip})\,]$ /* the sender sip is unknown */
2. $\texttt{OSPF}(\text{ip}, \text{nbrs}, \text{lsdb}, \text{hello_t})$
3. $+\, [\,\texttt{nbrExist}(\text{nbrs}, \text{sip})\,]$ /* sip is a known neighbour */
4. $\textbf{send}(\texttt{sndmsg}(\texttt{upd}(\{lsa \mid lsa \in \text{lsdb} \wedge \exists lhdr \in \text{lsa_hdrs} : lhdr \leq \texttt{hdr}(lsa)\}, \text{ip}), \{\text{sip}\}))$.
5. $\texttt{OSPF}(\text{ip}, \text{nbrs}, \text{lsdb}, \text{hello_t})$

As all other processes it checks whether the sender sip of the message is known. If it is unknown (Lines 1 and 2) the message is simply ignored and the process returns to `OSPF`. If sip is known (Lines 3 to 5) the node checks the LSA headers `lsahdrs` in the received message against the LSAs contained in the node's link state database lsdb. The LSAs that have headers matching those in the message are packaged into a new LSU message and sent back to sip, the sender of the original DBD message.

The process `UPD` (Process 5) handles the receipt of an LSU message.

Process 5 Handling LSU messages

$\texttt{UPD}(\text{lsas}, \text{sip}, \text{ip}, \text{nbrs}, \text{lsdb}, \text{hello_t}) \stackrel{def}{=}$
1. $[\![\text{lsas} := \{lsa \mid lsa \in \text{lsas} \wedge \nexists lsa' \in \text{lsdb} : \texttt{hdr}(lsa) \leq \texttt{hdr}(lsa')\}]\!]$ /* more recent LSAs */
2. (
3. $[\, \text{lsas} = \emptyset\,]$ /* no new LSAs exist */
4. $\texttt{OSPF}(\text{ip}, \text{nbrs}, \text{lsdb}, \text{hello_t})$
5. $+\, [\, \text{lsas} \neq \emptyset\,]$ /* new LSAs exist */
6. $[\![\text{lsdb} := \texttt{install}(\text{lsdb}, \text{lsas})]\!]$ /* install them */
7. $\textbf{send}(\texttt{sndmsg}(\texttt{upd}(\text{lsas}, \text{ip}), \{\pi_1(n) \mid n \in \text{nbrs}\}))$.
8. $\texttt{OSPF}(\text{ip}, \text{nbrs}, \text{lsdb}, \text{hello_t})$
9.)

In Line 1, the process finds those LSAs in the payload of the message that are either more recent than those in the link state database lsdb, or that do not yet exist in lsdb; these are stored in the variable lsas. In case no such LSA exists, i.e. the set lsas is empty, the process does not perform any action and returns to the main process `OSPF` (Lines 3 and 4). In case the incoming message contains new information (Line 5), i.e. lsas is not empty, the process `installs` the newly discovered LSAs into its own database (Line 6). Since the node has received new LSAs, it forwards them to its neighbours: Line 7 packages the set lsas into an update message and sends them to the neighbours. Lastly, the process returns to `OSPF`.

B.2.3 Queues

The process `QMSG` (Process 6), which is identical to the message queue presented in [15], handles the buffering of messages that are to be sent or received. By this we guarantee that no message is lost, regardless of which state the main process `OSPF` is in— `OSPF` can reach a state, such as `HELLO`, `DBD`, `REQ` or `UPD`, in which it is not ready to perform a receive action.

Process 6 Message Queue

```
QMSG(msgs) =def
1.    /* store incoming message at the end of msgs */
2.    receive(msg) . QMSG(append(msg,msgs))
3.  + [ msgs ≠ [] ]        /* the queue is not empty */
4.    (
5.      /* pop top message and send it to another sequential process */
6.      send(head(msgs)) . QMSG(tail(msgs))
7.      /* or receive and store an incoming message */
8.      + receive(msg) . QMSG(append(msg,msgs))
9.    )
```

The process `QMSG` runs in parallel with `OSPF` or any other process that might be called and maintains a queue (list) msgs of messages received. In fact, every node maintains two `QMSG` processes, one handling incoming messages and the other handling outgoing messages.

It is always ready to receive yet another new message msg and append it to msgs.

The process `QSND` (Process 7) handles the sending of messages. Similar to `QMSG`, a copy of `QSND` runs in parallel with `OSPF` or any other process that is called. It receives the next message to be sent in Line 1. Depending on the type of the message, the process either broadcasts the message (Lines 4 and 5), or groupcasts it to the destinations specified in the message received (Lines 6 and 7).

Process 7 Sending Process

```
QSND() =def
1.    receive(msg) .
2.    [ msg = sndmsg(realmsg,dips) ]        /* get content of msg */
3.      (
4.        [ realmsg = hello(ips,sip) ]        /* broadcast HELLO message */
5.          broadcast(realmsg) . QSND()
6.        + [ ⋀_{ips,sip} realmsg ≠ hello(ips,sip) ]        /* groupcast other messages */
7.          groupcast(dips,realmsg) . QSND()
8.      )
```

B.2.4 Initial State

We need to define an initial state to finish our specification. The initial expression of the network is an encapsulated parallel composition of node expressions ip:*P*:*R*, where the finite number of nodes and the range *R* of each node expression is left unspecified. Each node in the parallel composition must have a unique IP address *ip*. The initial process *P* for each *ip* is given by the following expression:

$$(\xi_0, \texttt{QSND}()) \langle\langle (\zeta, \texttt{QMSG}(\text{msgs})) \langle\langle (\chi, \texttt{OSPF}(\text{ip}, \text{nbrs}, \text{lsdb}, \text{hello_t})) \langle\langle (\theta, \texttt{QMSG}(\text{msgs}))$$

with $\zeta(\text{msgs}) = []$, $\chi(\text{ip}) = ip \wedge \chi(\text{nbrs}) = \chi(\text{lsdb}) = \emptyset \wedge \chi(\text{hello_t}) = 0$ and $\theta(\text{msgs}) = []$. As QSND does not maintain any variable its valuation is empty.

Initially both the outgoing and incoming message queues are empty. The process OSPF knows its own IP address, is not aware of any neighbour, has an empty LSDB and the hello timer is set to zero, so it will send a HELLO message as soon as possible.

A node will not lose any messages received even if they arrive while the process OSPF is busy. This is because the rightmost QMSG process is running in parallel and can receive and store any incoming message. As soon as OSPF can handle another message, it synchronises with the rightmost QMSG. The leftmost QMSG performs a similar role for outgoing messages; the processes passes on to QSND, which sends the messages to the incoming QMSG process of other nodes.

C The Detailed AWN Model

In this section we provide an in-depth description of our detailed model of OSPF, which we sketched in Section 4. To simplify the description of this model, many aspects are described by how they differ from the simpler model (Section 5 and Appendix B).

C.1 Data Structure

The types `LSA` and `LSDB` remain unchanged (see Appendix B for their definition), but the data structures dealing with neighbours carry a lot more information and as such they need to be altered.

There are two main reasons why these structures need to be more complex. First, we assume that messages can now be lost, therefore the protocol has to implement message retransmission as well as message acknowledgments. Secondly, it is no longer necessary for a node to form an adjacency with each of its neighbours.

C.1.1 Altered Data Structure Relating to Neighbours

Next to the node's IP address and time stamp, a neighbour structure now contains seven more fields: one for neighbour states, two to enable reliable Database Description exchange, two for LSR message retransmission, and two for LSU message retransmission.

We first introduce the concept of a *neighbour state* (`NS`). It is a string that represents the six phases of OSPF's adjacency establishment.

$$\texttt{NS} := \text{“Init”} \mid \text{“2-Way”} \mid \text{“ExStart”} \mid \text{“Exchange”} \mid \text{“Loading”} \mid \text{“Full”}$$

A node associates a neighbour state with each of its neighbours. The state describes whether the two nodes ip and sip should exchange LSAs and if so, it describes the stage of LSDB synchronisation. We explain the meaning of the strings in the following:

Init: A node ip has recently received a HELLO message from a neighbour. However, bidirectional communication is not confirmed yet—the node's own IP address is not part of the HELLO message.

2-Way: In this state communication between the two nodes has been confirmed to be bidirectional; only neighbours whose adjacency is not needed will enter “2-Way”.

ExStart: This is the first step in the adjacency-establishment procedure. It establishes the master/slave relationship between the nodes.

Exchange: A node in this state exchanges its link state database, using DBD messages.

Loading: The nodes have exchanged descriptions of their LSDBs and now compare the information with their local data. Missing information is requested by LSR messages.

Full: This state indicates full adjacency, i.e. the nodes have exchanged all information and their LSDBs are synchronised.

We define a strict order on `NS`. It reflects the progress in the adjacency-establishment procedure: the greater the value the further the procedure has evolved.

$$\text{“Init”} < \text{“2-Way”} < \text{“ExStart”} < \text{“Exchange”} < \text{“Loading”} < \text{“Full”}$$

Following the RFC, we call neighbours that are in state “ExStart” or greater *adjacencies*.

Before defining the altered neighbour structure we need to introduce the concept of sequence numbers, which were not present in the simple model. For that, we introduce the type `SQN`, which we assume

to be isomorphic to $\mathbb{N}$. Database Description Sequence Numbers (DDSQNs), used in DBD messages, ensure that both nodes correctly communicate. Every neighbour structure stores one sequence number, called a DD sequence number in the RFC. As usual, sequence numbers indicate the freshness of a DBD message. If a node receives a DBD message with a larger than expected sequence number the adjacency-establishment procedure needs to be restarted

We can now formally describe a *neighbour structure*; it is a tuple with nine components:

1. The IP address of a neighbour, which is an element of IP (as in the simple model).
2. A neighbour state, an element of NS.
3. A timestamp of type TIME, which indicates when the neighbour is considered inactive or dead—unless another message is received from this neighbour before that time (as in the simple model).
4. A DD sequence number, which is an element of SQN.
5. A timestamp of type TIME indicating when the DD timer will fire, which triggers resending of a DBD message.
6. A link state request list[10], an element of $\mathscr{P}(\mathtt{LSAHDR})$, containing the headers of all the LSAs that need to be requested from that neighbour.
7. A timestamp indicating when the next LSR message will be sent;
8. A link state retransmission list, an element of LSDB that contains all the LSAs that have been sent to the neighbour and are yet to be acknowledged.
9. A timestamp indicating when the link state retransmission timer will fire.

We denote the type of neighbour structures by NBR and define a generation function:

$$(_,_,_,_,_,_,_,_,_) : \mathtt{IP} \times \mathtt{NS} \times \mathtt{TIME} \times \mathtt{SQN} \times \mathtt{TIME} \times \mathscr{P}(\mathtt{LSAHDR}) \times \mathtt{TIME} \times \mathtt{LSDB} \times \mathtt{TIME} \rightarrow \mathtt{NBR}$$

Identically to the simple model, a *neighbour list* is a set of neighbour structures, where each IP address has at most one entry.

$$\mathtt{NBRS} := \{nbrs \mid nbrs \in \mathscr{P}(\mathtt{NBR}) \ \wedge\ \forall nbr_1, nbr_2 \in nbrs : nbr_1 \neq nbr_2 \Rightarrow \pi_1(nbr_1) \neq \pi_1(nbr_2)\}\,.$$

Of course they now contain the more complicated neighbour structures.

C.1.2 Messages

We also have to modify messages: DBD messages contain more information; LSR messages carry only one LSA header rather than a set of them; and ACK messages are introduced as a new message type. In detail we have the following:

1. A HELLO message remains unchanged and contains two entries: (a) a set of IP addresses of all neighbours discovered by the sender and (b) the identifier of the sender itself. The generation function is
$$\mathtt{hello} : \mathscr{P}(\mathtt{IP}) \times \mathtt{IP} \rightarrow \mathtt{MSG}\,.$$
2. A DBD message consists of four fields: (a) a set of LSA headers describing the LSDB of the sender (as in the simple model); (b) a DD sequence number for keeping the stream of messages reliable and ordered; (c) an init bit, an element of BOOL, indicating whether the message is the first in a sequence when set to true; and (d) the identifier of the sender itself (as in the simple model). The generation function is
$$\mathtt{dbd} : \mathscr{P}(\mathtt{LSAHDR}) \times \mathtt{SQN} \times \mathtt{BOOL} \times \mathtt{IP} \rightarrow \mathtt{MSG}\,.$$

[10]It is called list, but the order does not matter

3. An LSR message still consists of 2 fields, but the first is changed: (a) a single LSA header of type LSAHDR; and, again, (b) the identifier of the sender itself. The generation function is

$$\texttt{req} : \texttt{LSAHDR} \times \texttt{IP} \to \texttt{MSG}\,.$$

4. An LSU message remains unchanged and contains (1) an LSDB and (2) the identifier of the sender. It is generated by $\texttt{upd} : \texttt{LSDB} \times \texttt{IP} \to \texttt{MSG}$.
5. A *Link State Acknowledgement* ACK *message* consists of 2 fields: (a) a set of LSA headers, an element of $\mathscr{P}(\texttt{LSAHDR})$; and, again, (b) the identifier of the sender itself. It is generated by

$$\texttt{ack} : \mathscr{P}(\texttt{LSAHDR}) \times \texttt{IP} \to \texttt{MSG}\,.$$

C.1.3 Update Functions

As we take the adjacency-establishment procedure into account we need further functions. We also have to alter functions that take the new data structure into account. We also introduce auxiliary functions to better structure both the data structure and the processes.

As discussed in Section 4 we assume the existence of a predicate $\texttt{adj}(\mathrm{ip},\mathrm{ip}')$, which evaluates to true if the nodes ip and ip$'$ are supposed to form an adjacency.

$$\begin{array}{rcl} \texttt{adj} : \texttt{IP} \times \texttt{IP} & \to & \texttt{BOOL} \\ \texttt{adj}(\mathrm{ip},\mathrm{ip}') & :\Longleftrightarrow & \text{adjacency should be established between ip and ip}' \end{array}$$

We further assume that $\texttt{adj}(\mathrm{ip},\mathrm{ip}') = \texttt{adj}(\mathrm{ip}',\mathrm{ip})$.
We first consider all the functions that concern LSDBs.

1. The generation of LSAs now takes the neighbour state into account:

$$\begin{array}{rcl} \texttt{newLSA} : \texttt{IP} \times \texttt{TIME} \times \texttt{NBRS} & \to & \texttt{LSA} \\ \texttt{newLSA}(\mathrm{ip},\mathrm{t},\mathrm{nbrs}) & := & (\mathrm{ip},\mathrm{t},\{\pi_1(n) \mid n \in \mathrm{nbrs} \wedge \pi_2(n) \geq \text{“2-Way”}\}) \end{array}$$

2. The function install, which updates a node’s LSDB, remains unchanged.
3. To check whether the LSDB contains up to date information for a given LSA header, we use

$$\begin{array}{rcl} \texttt{lsaExist} : \texttt{LSDB} \times \texttt{LSAHDR} & \to & \texttt{BOOL} \\ \texttt{lsaExist}(\mathrm{lsdb},\mathrm{lhdr}) & :\Longleftrightarrow & \exists\, lsa \in \mathrm{lsdb} : \mathrm{lhdr} \leq \texttt{hdr}(lsa) \end{array}$$

4. A node has to find an entry in its LSDB, given an LSA header.

$$\begin{array}{rcll} \texttt{getLSA} : \texttt{LSDB} \times \texttt{LSAHDR} & \rightharpoonup & \texttt{LSA} & \\ \texttt{getLSA}(\mathrm{lsdb},\mathrm{lhdr}) & := & lsa & \text{if } \exists\, lsa \in \mathrm{lsdb} \wedge \pi_1(lsa) = \pi_1(\mathrm{lhdr}) \end{array}$$

As before, the protocol also requires functions to inspect elements of the neighbour list, or to modify it.

5. The function nbrExist, which checks whether a certain entry exists, remains unchanged.
6. The function that inserts a new neighbour structure needs alteration to reflect the new structure.

$$\begin{array}{rcll} \texttt{newNBR} : \texttt{NBRS} \times \texttt{IP} & \rightharpoonup & \texttt{NBRS} & \\ \texttt{newNBR}(\mathrm{nbrs},\mathrm{nip}) & := & \mathrm{nbrs} \cup \{(\mathrm{nip},\text{“Init”},0,0,0,\emptyset,0,\emptyset,0)\} & \text{if } \neg\texttt{nbrExist}(\mathrm{nbrs},\mathrm{nip}) \end{array}$$

7. Given an IP address we need to find the corresponding neighbour structure in the neighbour list.

$$\begin{array}{rcll} \texttt{getNBR} : \texttt{NBRS} \times \texttt{IP} & \rightharpoonup & \texttt{NBR} & \\ \texttt{getNBR}(\mathrm{nbrs},\mathrm{nip}) & := & n & \text{if } \exists\, n \in \mathrm{nbrs} : \pi_1(n) = \mathrm{nip} \end{array}$$

The function $\texttt{getNBR}(\mathrm{nbrs},\mathrm{nip})$ is defined iff $\texttt{nbrExist}(\mathrm{nbrs},\mathrm{nip})$. By the type definition of NBRS the entry is unique.

We also define functions that return individual components of the corresponding neighbour structure.

8. Return the neighbour state.

$$\begin{array}{rcll} \texttt{getNS} : \texttt{NBRS} \times \texttt{IP} & \rightharpoonup & \texttt{NS} & \\ \texttt{getNS}(\text{nbrs}, \text{nip}) & := & \pi_2(\texttt{getNBR}(\text{nbrs}, \text{nip})) & \text{if } \texttt{nbrExist}(\text{nbrs}, \text{nip}) \end{array}$$

9. Return the DD sequence number.

$$\begin{array}{rcll} \texttt{getDDSQN} : \texttt{NBRS} \times \texttt{IP} & \rightharpoonup & \texttt{SQN} & \\ \texttt{getDDSQN}(\text{nbrs}, \text{nip}) & := & \pi_4(\texttt{getNBR}(\text{nbrs}, \text{nip})) & \text{if } \texttt{nbrExist}(\text{nbrs}, \text{nip}) \end{array}$$

10. Return the link state request list.

$$\begin{array}{rcll} \texttt{getREQS} : \texttt{NBRS} \times \texttt{IP} & \rightharpoonup & \mathscr{P}(\texttt{LSAHDR}) & \\ \texttt{getREQS}(\text{nbrs}, \text{nip}) & := & \pi_6(\texttt{getNBR}(\text{nbrs}, \text{nip})) & \text{if } \texttt{nbrExist}(\text{nbrs}, \text{nip}) \end{array}$$

11. Return the link state retransmission list.

$$\begin{array}{rcll} \texttt{getRXMTS} : \texttt{NBRS} \times \texttt{IP} & \rightharpoonup & \texttt{LSDB} & \\ \texttt{getRXMTS}(\text{nbrs}, \text{nip}) & := & \pi_8(\texttt{getNBR}(\text{nbrs}, \text{nip})) & \text{if } \texttt{nbrExist}(\text{nbrs}, \text{nip}) \end{array}$$

Dual to reading values from a neighbour structure we need to insert and alter entries.

12. Initialise a neighbour state to *ns* for a given IP address *nip*. Both the link state request list and the set of link state retransmissions are set to empty.

$$\begin{array}{rcl} \texttt{initNBR} : \texttt{NBRS} \times \texttt{IP} \times \texttt{NS} & \rightarrow & \texttt{NBRS} \\ \texttt{initNBR}(\text{nbrs}, \text{nip}, \text{ns}) & := & \{n \mid n \in \text{nbrs} \wedge \pi_1(n) \neq \text{nip}\} \\ & & \cup \{(\pi_1(n), \text{ns}, \pi_3(n), \pi_4(n), \pi_5(n), [\,], \pi_7(n), \emptyset, \pi_9(n)) \mid n \in \mathit{nbrs} \wedge \pi_1(n) = \text{nip}\} \end{array}$$

13. Update the neighbour state of a given IP address.

$$\begin{array}{rcl} \texttt{setNS} : \texttt{NBRS} \times \texttt{IP} \times \texttt{NS} & \rightarrow & \texttt{NBRS} \\ \texttt{setNS}(\text{nbrs}, \text{nip}, \text{ns}) & := & \{n \mid n \in \text{nbrs} \wedge \pi_1(n) \neq \text{nip}\} \\ & & \cup \{(\pi_1(n), \text{ns}, \pi_3(n), \ldots, \pi_9(n)) \mid n \in \text{nbrs} \wedge \pi_1(n) = \text{nip}\} \end{array}$$

In case no entry exists for nip, the neighbour structure remains unchanged. The same applies to most of the functions defined in the remainder of the section.

14. Update the activity timer.

$$\begin{array}{rcl} \texttt{setINACTT} : \texttt{NBRS} \times \texttt{IP} \times \texttt{TIME} & \rightarrow & \texttt{NBRS} \\ \texttt{setINACTT}(\text{nbrs}, \text{nip}, \text{t}) & := & \{n \mid n \in \text{nbrs} \wedge \pi_1(n) \neq \text{nip}\} \\ & & \cup \{(\pi_1(n), \pi_2(n), \text{t}, \pi_4(n), \ldots, \pi_9(n)) \mid n \in \text{nbrs} \wedge \pi_1(n) = \text{nip}\} \end{array}$$

15. Update the DD sequence number. We have a function to set the sequence number to an arbitrary value and one for incrementation.

$$\begin{array}{rcl} \texttt{setDDSQN} : \texttt{NBRS} \times \texttt{IP} \times \texttt{SQN} & \rightarrow & \texttt{NBRS} \\ \texttt{setDDSQN}(\text{nbrs}, \text{nip}, \text{sqn}) & := & \{n \mid n \in \text{nbrs} \wedge \pi_1(n) \neq \text{nip}\} \\ & & \cup \{(\pi_1(n), \ldots, \pi_3(n), \text{sqn}, \pi_5(n), \ldots, \pi_9(n)) \mid n \in \text{nbrs} \wedge \pi_1(n) = \text{nip}\} \\ \texttt{incDDSQN} : \texttt{NBRS} \times \texttt{IP} & \rightarrow & \texttt{NBRS} \\ \texttt{incDDSQN}(\text{nbrs}, \text{nip}) & := & \{n \mid n \in \text{nbrs} \wedge \pi_1(n) \neq \text{nip}\} \\ & & \cup \{(\pi_1(n), \ldots, \pi_3(n), \pi_4(n){+}1, \pi_5(n), \ldots, \pi_9(n)) \mid n \in \text{nbrs} \wedge \pi_1(n) = \text{nip}\} \end{array}$$

16. Update the DD timer.

$$\begin{array}{rcl} \texttt{setDDT} : \texttt{NBRS} \times \texttt{IP} \times \texttt{TIME} & \rightarrow & \texttt{NBRS} \\ \texttt{setDDT}(\text{nbrs}, \text{nip}, \text{t}) & := & \{n \mid n \in \text{nbrs} \wedge \pi_1(n) \neq \text{nip}\} \\ & & \cup \{(\pi_1(n), \ldots, \pi_4(n), \text{t}, \pi_6(n), \ldots, \pi_9(n)) \mid n \in \text{nbrs} \wedge \pi_1(n) = \text{nip}\} \end{array}$$

17. We update the link state request list in several ways. First we can replace that list with a given set.

$$\begin{aligned}\texttt{setREQS} : \texttt{NBRS} \times \texttt{IP} \times \mathscr{P}(\texttt{LSAHDR}) &\rightarrow \texttt{NBRS}\\ \texttt{setREQS}(\text{nbrs},\text{nip},\text{reqs}) &:= \{n \mid n \in \text{nbrs} \wedge \pi_1(n) \neq \text{nip}\}\\ &\cup \{(\pi_1(n),\ldots,\pi_5(n),\text{reqs},\pi_7(n),\ldots,\pi_9(n)) \mid n \in \text{nbrs} \wedge \pi_1(n) = \text{nip}\}\end{aligned}$$

We also want to update the list by removing LSA headers that are dominated by elements in lsdb.

$$\begin{aligned}\texttt{cleanREQS} &: \texttt{NBRS} \times \texttt{IP} \times \texttt{LSDB} \rightharpoonup \texttt{NBRS}\\ \texttt{cleanREQS}(\text{nbrs},\text{nip},\text{lsdb}) &:= \texttt{setREQS}(\text{nbrs},\text{nip},\mathit{reqs}) \qquad \text{if } \texttt{nbrExist}(\text{nbrs},\text{nip})\,,\end{aligned}$$

where $\mathit{reqs} = \{\mathit{lhdr} \in \texttt{getREQS}(\text{nbrs},\text{nip}) \mid \nexists \mathit{lsa} \in \text{lsdb} : \mathit{lhdr} \leq \texttt{hdr}(\mathit{lsa})\}$.
Moreover, when receiving a DBD message, the protocol wants to add elements to the link state request list; this is combined with cleaning up the list, i.e. removing elements that are dominated.

$$\begin{aligned}\texttt{addREQS} : \texttt{NBRS} \times \texttt{IP} \times \texttt{LSDB} \times \mathscr{P}(\texttt{LSAHDR}) &\rightharpoonup \texttt{NBRS}\\ \texttt{addREQS}(\text{nbrs},\text{nip},\text{lsdb},\text{lsa_hdrs}) &:= \texttt{cleanREQS}(\mathit{nbrs}',\text{nip},\text{lsdb})\,,\end{aligned}$$

where $\mathit{nbrs}' = \texttt{setREQS}(\text{nbrs},\text{nip},\texttt{getREQS}(\text{nbrs},\text{nip}) \cup \text{lsa_hdrs})$.
The set nbrs' is the neighbour list nbrs where nip's request list is extended by all the LSA headers lsa_hdrs. The function `cleanREQS` removes all those headers that are dominated by lsdb.

18. Update the link state request timer.

$$\begin{aligned}\texttt{setREQT} : \texttt{NBRS} \times \texttt{IP} \times \texttt{TIME} &\rightarrow \texttt{NBRS}\\ \texttt{setREQT}(\text{nbrs},\text{nip},\text{t}) &:= \{n \mid n \in \text{nbrs} \wedge \pi_1(n) \neq \text{nip}\}\\ &\cup \{(\pi_1(n),\ldots,\pi_6(n),\text{t},\pi_8(n),\pi_9(n)) \mid n \in \text{nbrs} \wedge \pi_1(n) = \text{nip}\}\end{aligned}$$

19. An update to the link state retransmission list, an element of type `LSDB`, happens in different ways. Firstly, we can set the retransmission list directly.

$$\begin{aligned}\texttt{setRXMTS} : \texttt{NBRS} \times \texttt{IP} \times \texttt{LSDB} &\rightarrow \texttt{NBRS}\\ \texttt{setRXMTS}(\text{nbrs},\text{nip},\text{lsas}) &:= \{n \mid n \in \text{nbrs} \wedge \pi_1(n) \neq \text{nip}\}\\ &\cup \{(\pi_1(n),\ldots,\pi_7(n),\text{lsas},\pi_9(n)) \mid n \in \text{nbrs} \wedge \pi_1(n) = \text{nip}\}\end{aligned}$$

Secondly, given a set of LSA headers we remove corresponding LSAs from the retransmission list.

$$\begin{aligned}\texttt{cleanRXMTS} &: \texttt{NBRS} \times \texttt{IP} \times \mathscr{P}(\texttt{LSAHDR}) \rightharpoonup \texttt{NBRS}\\ \texttt{cleanRXMTS}(\text{nbrs},\text{nip},\text{lsa_hdrs}) &:= \texttt{setRXMTS}(\text{nbrs},\text{nip},\mathit{rxmts}) \quad \text{if } \texttt{nbrExist}(\text{nbrs},\text{nip})\end{aligned}$$

where $\mathit{rxmts} = \{\mathit{lsa} \mid \mathit{lsa} \in \texttt{getRXMTS}(\text{nbrs},\text{nip}) \wedge \nexists \mathit{lhdr} \in \text{lsa_hdrs} : \texttt{hdr}(\mathit{lsa}) \leq \mathit{lhdr}\}$.
For a given LSDB the update needs to change all the neighbour structures with a neighbour state that is equal to or greater than "Exchange", only keeping the most recent version of each LSA.

$$\begin{aligned}\texttt{updRXMTS} : \texttt{NBRS} \times \texttt{LSDB} &\rightarrow \texttt{NBRS}\\ \texttt{updRXMTS}(\text{nbrs},\text{lsas}) &:= \{\texttt{updRX}(n,\text{lsas}) \mid n \in \text{nbrs} \wedge \pi_2(n) \geq \text{“Exchange”}\}\\ &\cup \{n \mid n \in \text{nbrs} \wedge \pi_2(n) < \text{“Exchange”}\}\end{aligned}$$

where the function `updRX` installs the LSAs in the given neighbour structure.

$$\begin{aligned}\texttt{updRX} : \texttt{NBR} \times \texttt{LSDB} &\rightarrow \texttt{NBR}\\ \texttt{updRX}(\text{ns},\text{lsas}) &:= (\pi_1(\text{ns}),...,\pi_7(\text{ns}),\texttt{install}(\pi_8(\text{ns}),\text{lsas}),\pi_9(\text{ns}))\end{aligned}$$

20. Update the link state retransmission list timer.

$$\begin{aligned}\texttt{setRXMTT} : \texttt{NBRS} \times \texttt{IP} \times \texttt{TIME} &\rightarrow \texttt{NBRS}\\ \texttt{setRXMTT}(\text{nbrs},\text{nip},\text{t}) &:= \{n \mid n \in \text{nbrs} \wedge \pi_1(n) \neq \text{nip}\}\\ &\cup \{(\pi_1(n),\ldots,\pi_8(n),\text{t}) \mid n \in \text{nbrs} \wedge \pi_1(n) = \text{nip}\}\end{aligned}$$

The following five functions distil information out of a neighbour structure, other than individual values.

21. The function deadNBRS keeps the same functionality as in the simple model, namely identifying inactive neighbours. However, the function needs adaptation to the new neighbour structure.

$$\begin{aligned} \texttt{deadNBRS} : \texttt{NBRS} \times \texttt{TIME} &\rightarrow \texttt{NBRS} \\ \texttt{deadNBRS}(\text{nbrs}, \text{t}) &:= \{n \mid n \in \text{nbrs} \wedge \pi_3(n) < \text{t}\} \end{aligned}$$

For the following selections, we use a function $\texttt{select} : \mathscr{P}(\texttt{IP}) \rightharpoonup \texttt{IP}$ that (deterministically) chooses an element from a set of IP addresses.

22. Given a time t, return the IP of a neighbour whose DD timer has fired and has a neighbour state which is "ExStart" or is the slave in an ongoing exchange process.

$$\begin{aligned} \texttt{ddNIP} : \texttt{NBRS} \times \texttt{TIME} \times \texttt{IP} &\rightharpoonup \texttt{IP} \\ \texttt{ddNIP}(\text{nbrs}, \text{t}, \text{ip}) &:= \texttt{select}(\mathit{ddnips}) \qquad \text{if } \mathit{ddnips} \neq \emptyset\,, \end{aligned}$$

where $\mathit{ddnips} = \{\pi_1(n) \mid n \in \text{nbrs} \wedge \pi_5(n) < \text{t} \wedge ((\pi_2(n) = \text{"ExStart"}) \vee (\pi_2(n) = \text{"Exchange"} \wedge \pi_1(n) \leq \mathit{ip}))\}$. Note that this function requires a total order on IP, which we assume to exist. We use this order in the establishment of the master/slave relationship between two nodes: the one with the larger IP address will be the master.

23. Given a neighbour structure and a time, the protocol picks a neighbour whose link state request timer has fired to start a communication.

$$\begin{aligned} \texttt{reqNIP} : \texttt{NBRS} \times \texttt{TIME} &\rightharpoonup \texttt{IP} \\ \texttt{reqNIP}(\text{nbrs}, \text{t}) &:= \texttt{select}(\mathit{reqnips}) \qquad \text{if } \mathit{reqnips} \neq \emptyset\,, \end{aligned}$$

where $\mathit{reqnips} = \{\pi_1(n) \mid n \in \text{nbrs} \wedge \pi_7(n) < \text{t} \wedge \pi_6(n) \neq \emptyset\}$.
Similarly to this, the protocol also picks a node whose link state retransmission timer has fired, and that has a non-empty link state retransmission list.

$$\begin{aligned} \texttt{rxmtNIP} : \texttt{NBRS} \times \texttt{TIME} &\rightharpoonup \texttt{IP} \\ \texttt{rxmtNIP}(\text{nbrs}, \text{t}) &:= \texttt{select}(\mathit{rxmtnips}) \qquad \text{if } \mathit{rxmtnips} \neq \emptyset\,, \end{aligned}$$

where $\mathit{rxmtnips} = \{\pi_1(n) \mid n \in \text{nbrs} \wedge \pi_9(n) < \text{t} \wedge \pi_8(n) \neq \emptyset\}$.

24. Get the IPs of all of the neighbours whose neighbour state $\geq$ "Exchange".

$$\begin{aligned} \texttt{floodNIPS} : \texttt{NBRS} &\rightarrow \mathscr{P}(\texttt{IP}) \\ \texttt{floodNIPS}(\text{nbrs}) &:= \{\pi_1(n) \mid n \in \text{nbrs} \wedge \pi_2(n) \geq \text{"Exchange"}\} \end{aligned}$$

OSPF only sends messages to those neighbours that have a status equal to or greater than "ExStart".

25. We define a function that generates a DBD message intended for node nip.

$$\texttt{genDBD} : \texttt{NBRS} \times \texttt{LSDB} \times \texttt{IP} \times \texttt{IP} \rightharpoonup \texttt{MSG}$$

$$\texttt{genDBD}(\text{nbrs}, \text{lsdb}, \text{nip}, \text{ip}) := \begin{cases} \texttt{dbd}(\{\texttt{hdr}(\mathit{lsa}) \mid \mathit{lsa} \in \text{lsdb}\}, \texttt{getDDSQN}(\text{nbrs}, \text{nip}), \texttt{true}, \text{ip}) \\ \qquad\qquad \text{if } \texttt{nbrExist}(\text{nbrs}, \text{nip}) \wedge \pi_2(n) = \text{"ExStart"} \\ \texttt{dbd}(\emptyset, \texttt{getDDSQN}(\text{nbrs}, \text{nip}), \texttt{false}, \text{ip}) \\ \qquad\qquad \text{if } \texttt{nbrExist}(\text{nbrs}, \text{nip}) \wedge \pi_2(n) \geq \text{"Exchanging"} \end{cases}$$

C.2 The Full Model

The detailed model consists of the 9 processes OSPF, HELLO, DBD, SNMIS, REQ, UPD, ACK, QMSG and QSND. The main functionality of the processes already present in the simple model do not change. For completeness, we list a short summary nevertheless.

- As before, the process OSPF is the main process of the protocol. It reads messages from the message queue and calls other processes depending on message types; it periodically broadcasts HELLO messages, checks liveness of discovered neighbours, and refreshes the router's own LSA.
- The process HELLO describes all the actions performed when a HELLO message is received. This includes updating the relevant inactivity timer and, if the HELLO message is from a previously unknown neighbour, updating the node's neighbour list. The process also generates a DBD message, which is sent back to the originator of the HELLO message, and a new LSA when required.
- The process DBD handles incoming DBD messages. Sending requests for those LSAs from DBD message that do not have a matching partner in the node's LSDB is among its actions.
- The new process SNMIS describes the actions required when a sequence number mismatch occurs during the adjacency-establishment procedure. The actions include reinitialising neighbour structures and restarting the adjacency-establishment procedure.
- The process REQ describes the actions following the receipt of an LSR message, such as finding the requested LSAs in the node's local data, and sending them back to the sender of the message.
- The process UPD manages incoming LSU messages, including the installation of up-to-date LSAs in the node's LSDB, and broadcasting the updated information.
- The process ACK describes the actions taken to handle the receipt of an LSA. This involves removing acknowledged LSAs from link state retransmission lists.
- The process QMSG models a message queue. As before we have two instances, a queue for incoming messages and one for outgoing processes.
- The process QSND is responsible for sending outgoing messages that were generated by the other processes described above.

The system initialisation is identical to the simple model: each node runs four processes in parallel.

$$(\xi_0, \texttt{QSND}()) \langle\langle (\zeta, \texttt{QMSG}(\text{msgs})) \langle\langle (\chi, \texttt{OSPF}(\text{ip}, \text{nbrs}, \text{lsdb}, \text{hello_t})) \langle\langle (\theta, \texttt{QMSG}(\text{msgs})),$$

where the valuation functions are defined in a straightforward way; details can be found in Section B.2.4.

C.2.1 The Basic Routine

The basic process OSPF (Process 8) consists of seven parts: handling incoming messages and distributing them to subprocesses, sending out HELLO messages periodically, removal of neighbours that are considered inactive, finding neighbours whose DD timer has fired, retransmitting DBD, LSR and LSU messages, as well as the generation of new LSAs.

The message handling is described in Lines 1 to 14. The process first receives a message by **receive**(msg); by our specification of the input queue[11] this will be the message currently in the front of the queue (the oldest message). OSPF then checks the type of the message and calls the corresponding process: in case of a HELLO message the process HELLO is called, etc.

The second part of OSPF (Lines 15 to 18) ensures that HELLO messages are sent periodically. The parameter hello_t indicates when the next HELLO message should be sent. The process checks whether it is time to send the next HELLO message; if so it resets the timer (Line 16) and broadcasts a HELLO message. In fact the message is sent to the process QSND (Process 7), which handles the actual sending.

The third part of OSPF (Lines 19 to 25) collects all of the neighbours from which no activity has been seen for the last `rtdead_intvl` seconds, removes them from nbrs and informs the neighbours about these updates. The only difference to Process 1 is that the link state retransmission list needs to be updated (Line 23), and that the destinations of the LSU message are determined by `floodNIPS`.

[11] The input queue is the same for both models; see Process 6.

Process 8 The Basic Routine

```
OSPF(ip,nbrs,lsdb,hello_t) ≝
 1.  receive(msg) .
 2.  /* depending on the message, the node calls different processes */
 3.  (
 4.     [ msg = hello(ips,sip) ]      /* HELLO message received */
 5.       HELLO(ips,sip,ip,nbrs,lsdb,hello_t)
 6.     + [ msg = dbd(lsa_hdrs,sqn,ibit,sip) ]      /* DBD message received */
 7.       DBD(lsa_hdrs,sqn,ibit,sip,ip,nbrs,lsdb,hello_t)
 8.     + [ msg = req(lhdr,sip) ]      /* LSR message received */
 9.       REQ(lhdr,sip,ip,nbrs,lsdb,hello_t)
10.     + [ msg = upd(lsas,sip) ]      /* LSU message received */
11.       UPD(lsas,sip,ip,nbrs,lsdb,hello_t)
12.     + [ msg = ack(lsa_hdrs,sip) ]      /* ACK message received */
13.       ACK(lsa_hdrs,sip,ip,nbrs,lsdb,hello_t)
14.  )
15. + [ now ≥ hello_t ]      /* send HELLO message */
16.  [[hello_t := now + hello_intvl]]
17.  send(sndmsg(hello({π1(n) | n ∈ nbrs},ip),∅)) .
18.  OSPF(ip,nbrs,lsdb,hello_t)
19. + [ deadNBRS(nbrs,now) ≠ ∅ ]      /* inactive neighbours, which should be removed, exist */
20.  [[nbrs := nbrs − deadNBRS(nbrs,now)]]      /* remove corresponding neighbour structures */
21.  [[lsa := newLSA(ip,now,nbrs)]]      /* generate a new LSA */
22.  [[lsdb := install(lsdb,{lsa})]]      /* install it */
23.  [[nbrs := updRXMTS(nbrs,{lsa})]]      /* update the link state retransmission list */
24.  send(sndmsg(upd({lsa},ip),floodNIPS(nbrs))) .      /* and flood it out */
25.  OSPF(ip,nbrs,lsdb,hello_t)
26. + [ nip = ddNIP(nbrs,now,ip) ]      /* find neighbour whose DD timer has fired */
27.  [[nbrs := setDDT(nbrs,nip,now+rxmt_intvl)]]
28.  send(sndmsg(genDBD(nbrs,lsdb,nip,ip),{nip})) .
29.  OSPF(ip,nbrs,lsdb,hello_t)
30. + [ nip = reqNIP(nbrs,now) ]      /* find a neighbour whose request timer has fired */
31.  [[nbrs := setREQT(nbrs,nip,now+rxmt_intvl)]]
32.  send(sndmsg(req(select(getREQS(nbrs,nip)),ip),{nip})) .
33.  OSPF(ip,nbrs,lsdb,hello_t)
34. + [ nip = rxmtNIP(nbrs,now) ]      /* find a neighbour whose retransmission timer has fired */
35.  [[nbrs := setRXMTS(nbrs,nip,now+rxmt_intvl)]]
36.  send(sndmsg(upd(getRXMTS(nbrs,nip),ip),{nip})) .
37.  OSPF(ip,nbrs,lsdb,hello_t)
38. + [ ¬lsaExist(lsdb,(ip,now − refresh_intvl)) ]      /* self-originated LSA too old */
39.  [[lsa := newLSA(ip,now,nbrs)]]      /* generate new one */
40.  [[lsdb := install(lsdb,{lsa})]]
41.  [[nbrs := updRXMTS(nbrs,{lsa})]]
42.  send(sndmsg(upd({lsa},ip),floodNIPS(nbrs))) .
43.  OSPF(ip,nbrs,lsdb,hello_t)
```

Lines 26 to 29 handle the retransmission of DBD messages every `rxmt_intvl` seconds. The function `ddNIP`(nbrs, `now`, ip) returns the IP address of a neighbour whose DD timer has fired. The DD timer for that neighbour is reset (Line 27) and a DBD message is sent to it.

Similarly, Lines 30 to 33 handle the retransmission of LSR messages every *rxmtintvl* seconds. Instead of a DBD message, a LSR message is sent out requesting a new LSA.

In the same spirit, Lines 37 to 36 handle the retransmission of LSU messages every *rxmtintvl* seconds.

The last part of OSPF, Lines 38 to 43, is executed if there is an LSA in the link state database that is too old. In that case the process generates a new LSA in Line 39 (which happens every *refreshintvl* seconds), installs that LSA into its lsdb (Line 40), updates the neighbour list accordingly (Line 41), and floods the LSA out (Line 42).

C.2.2 Message Processing

The process HELLO (Process 9) handles the receipt of a HELLO message.

Process 9 Handling HELLO messages

$\mathtt{HELLO}(\mathrm{ips},\mathrm{sip},\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t}) \stackrel{def}{=}$

1. $[\ \neg\mathtt{nbrExist}(\mathrm{nbrs},\mathrm{sip})\ ]$ /* the sender sip is unknown */
2. $[\![\mathrm{nbrs} := \mathtt{newNBR}(\mathrm{nbrs},\mathrm{sip})]\!]$
3. $\mathtt{HELLO}(\mathrm{ips},\mathrm{sip},\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$
4. $+\ [\ \mathtt{nbrExist}(\mathrm{nbrs},\mathrm{sip})\ ]$ /* sip is a known neighbour */
5. $[\![\mathrm{nbrs} := \mathtt{setINACTT}(\mathrm{nbrs},\mathrm{sip},\mathtt{now}+\mathtt{rtdead_intvl})]\!]$
6. $[\![\mathrm{ns} := \mathtt{getNS}(\mathrm{nbrs},\mathrm{sip})]\!]$
7. (
8. $[\ \mathrm{ip}\in\mathrm{ips}\wedge\mathtt{adj}(\mathrm{ip},\mathrm{sip})\wedge\mathrm{ns}=\text{“Init”}\ ]$ /* 2-WayReceived, start adjacency-forming */
9. $[\![\mathrm{nbrs} := \mathtt{setNS}(\mathrm{nbrs},\mathrm{sip},\text{“ExStart”})]\!]$
10. $[\![\mathrm{nbrs} := \mathtt{incDDSQN}(\mathrm{nbrs},\mathrm{sip})]\!]$
11. $[\![\mathrm{nbrs} := \mathtt{setDDT}(\mathrm{nbrs},\mathrm{sip},\mathtt{now}+\mathtt{rxmt_intvl})]\!]$
12. $\textbf{send}(\mathtt{sndmsg}(\mathtt{genDBD}(\mathrm{nbrs},\mathrm{lsdb},\mathrm{sip},\mathrm{ip}),\{\mathrm{sip}\}))\ .$
13. $\mathtt{OSPF}(\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$
14. $+\ [\ \mathrm{ip}\in\mathrm{ips}\wedge\mathtt{adj}(\mathrm{ip},\mathrm{sip})\wedge\mathrm{ns}\geq\text{“ExStart”}\ ]$ /* Adjacency-forming already underway */
15. $\mathtt{OSPF}(\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$
16. $+\ [\ \mathrm{ip}\in\mathrm{ips}\wedge\neg\mathtt{adj}(\mathrm{ip},\mathrm{sip})\ ]$ /* 2-WayReceived, adjacency-forming not needed */
17. $[\![\mathrm{nbrs} := \mathtt{setNS}(\mathrm{nbrs},\mathrm{sip},\text{“2-Way”})]\!]$
18. $\mathtt{OSPF}(\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$
19. $+\ [\ \mathrm{ip}\notin\mathrm{ips}\ ]$ /* 1-WayReceived */
20. $[\![\mathrm{nbrs} := \mathtt{initNBR}(\mathrm{nbrs},\mathrm{sip},\text{“Init”})]\!]$
21. $\mathtt{OSPF}(\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$
22.)

The first part of the process (Lines 1 to 3) is executed when the originator of the HELLO message is unknown to the node ip. The neighbour list is updated in Line 2 and the process calls itself. As the node is now aware of its neighbour sip, Lines 4 to 22 are executed.

In the second part, Lines 4 ff., the sender sip of the HELLO message is already known. As the node communicates with sip the corresponding inactivity timer is reset in Line 5. The subsequent actions

depend on the state of the node. If the sender has discovered the node ip (ip $\in$ ips) and the two nodes are intended to form an adjacency (adj(sip,ip)) but the establishment process has not yet started (the status is "Init") then the node initialises the exchange of the LSDBs. For that it sets the neighbour status to "ExStart" (Line 9), increments the DD sequence number (Line 10) to identify the outgoing DBD message (Line 12) in a unique way. It also resets the DD timer in Line 11. If the adjacency-establishment procedure has already begun (ns $\geq$ "ExStart")—of course the node sip should still be aware of ip—no actions are required as the other processes, such as DBD, handle the necessary actions (Lines 14 and 15). In case the node was discovered by sip, but adjacency is prohibited ($\neg$adj(ip, sip)) then the neighbour status is changed to "2-Way" (Line 17), indicating that a bidirectional link has been established; no other actions are necessary. If ip has not yet been discovered by sip, or if the connectivity between the two nodes broke down in the past, then the corresponding neighbour structure will be reset (or initialised) in Line 20.

The process DBD (Process 10) handles the receipt of a DBD message. It is the process that differs the most compared to the simple model for it handles most of the adjacency-establishment procedure. To increase readability we use predicates in the guards—we describe these predicates along with a detailed explanation of the steps of the process.

Line 2 ignores the DBD message in case the sender sip is unknown, which is checked in Line 1.

Lines 3– 49 handle the case where the sender sip is known. Depending on the status of the adjacency-establishment procedure the process performs different actions. The decision of which actions are taken depend on six conditions: (a) the intended adjacency relationship between ip and sip (adj(ip, sip)); (b) ip's neighbour state for sip; (c) the DD timer for sip; (d) the DD sequence number; (d) the IP addresses of ip and sip; and (f) the bit ibit in the received DBD message. Lines 4 and 5 describe the case where the nodes do not form an adjacency, and the nodes communicated before (ns = "Init"). By the incoming DBD message bidirectional communication is established; this fact is registered in the neighbour structure by setting the neighbour status for sip to "2-Way". In case the status is already set to "2-Way" (Line 6), indicating that full communication has been established and also that adjacency is not sought, the message can be ignored (Line 7). The remaining cases concern the establishment of adjacency. If the nodes aim for adjacency, but the procedure has not been started (ns = "Init" $\wedge$ adj(ip, sip)), checked in 7, then the adjacency establishment begins by sending out a DBD message in Line 11 and the process DBD is called again. Before that the node's data structures need updates: the neighbour state for sip is set to "ExStart" (Line 8) showing that the adjacency procedure has indeed begun, the DD sequence number is incremented and the DD timer reset (Lines 8 and 9). Calling the DBD process again allows a slave to immediately send another DBD message (executing Lines 13 to 15). Determining the slave/master relationship is the first proper step in the adjacency-establishment procedure. Although the RFC allows other methods as well, we use the total order on IP addresses as underlying criterion. We define abbreviations to check the slave/master status of a node ip with regards to sip:

$$\texttt{is_slave} := \texttt{adj}(\text{ip}, \text{sip}) \wedge \text{ip} < \text{sip} \qquad \texttt{is_master} := \texttt{adj}(\text{ip}, \text{sip}) \wedge \text{ip} > \text{sip}$$

As we assume IP addresses to be unique for each node in a network, we can ignore the case ip = sip.

In case an adjacency establishment has begun (ns = "ExStart"), the node itself identified itself as slave, i.e. it has a smaller sequence number than the sender, and it receives the first message from the master during the establishment (the bit ibit is set),[12] then the node updates its neighbour structure, indicating that an "Exchange" of messages is happening with sip, and the stored sequence number of that node (both Line 14). It does not reset the DD timer as it is up to the master to investigate in case

[12] This phase is called "NegotiationDone" in the RFC.

Process 10 Handling DBD messages

```
DBD(lsa_hdrs,sqn,ibit,sip,ip,nbrs,lsdb,hello_t) ≝
1.  [ ¬nbrExist(nbrs,sip) ]      /* the sender sip is unknown */
2.    OSPF(ip,nbrs,lsdb,hello_t)
3.  + [getNBR(nbrs,sip) = (sip,ns,inact_t,ddsqn,dd_t,reqs,req_t,rxmts,rxmt_t) ]     /* sip is known */
4.    ( [ns = "Init" ∧ ¬adj(ip,sip)]      /* no need for adjacency establishment */
5.      [[nbrs := setNS(nbrs,sip,"2-Way")]] OSPF(ip,nbrs,lsdb,hello_t)
6.    + [ ns = "2-Way" ]      /* ignore the message */
7.      OSPF(ip,nbrs,lsdb,hello_t)
8.    + [ ns = "Init" ∧ adj(ip,sip) ]      /* start adjacency-forming */
9.      [[nbrs := setNS(nbrs,sip,"ExStart")]] [[nbrs := incDDSQN(nbrs,sip)]]
10.     [[nbrs := setDDT(nbrs,sip,now+rxmt_intvl)]]
11.     send(sndmsg(genDBD(nbrs,lsdb,sip,ip),{sip})) .
12.     DBD(lsa_hdrs,sqn,ibit,sip,ip,nbrs,lsdb,hello_t)
13.   + [ negotiate_slave ]      /* negotiation done */
14.     [[nbrs := setNS(nbrs,sip,"Exchange")]] [[nbrs := setDDSQN(nbrs,sip,sqn)]]
15.     send(sndmsg(genDBD(nbrs,lsdb,sip,ip),{sip})) . OSPF(ip,nbrs,lsdb,hello_t)
16.   + [ negotiate_master ]      /* negotiation done */
17.     [[nbrs := setNS(nbrs,sip,"Exchange")]] [[nbrs := incDDSQN(nbrs,sip)]]
18.     [[nbrs := setDDT(nbrs,sip,now+rxmt_intvl)]] [[nbrs := addREQS(nbrs,sip,lsdb,lsa_hdrs)]]
19.     send(sndmsg(genDBD(nbrs,lsdb,sip,ip),{sip})) . OSPF(ip,nbrs,lsdb,hello_t)
20.   + [ negotiate_others ]      /* ignore the message */
21.     OSPF(ip,nbrs,lsdb,hello_t)
22.   + [ exchange_duplicate_slave ]      /* duplicate message */
23.     send(sndmsg(genDBD(nbrs,lsdb,sip,ip),{sip})) . OSPF(ip,nbrs,lsdb,hello_t)
24.   + [ exchange_duplicate_master ]      /* duplicate message */
25.     OSPF(ip,nbrs,lsdb,hello_t)
26.   + [ exchange_slave ]      /* exchange done */
27.     [[nbrs := incDDSQN(nbrs,sip)]] [[nbrs := setDDT(nbrs,sip,now+rtdead_intvl)]]
28.     [[nbrs := addREQS(nbrs,sip,lsdb,lsa_hdrs)]]
29.     send(sndmsg(genDBD(nbrs,lsdb,sip,ip),{sip})) .
30.     ( [getREQS(nbrs,sip) ≠ ∅]
31.       [[nbrs := setNS(nbrs,sip,"Loading")]] . OSPF(ip,nbrs,lsdb,hello_t)
32.     + [ getREQS(nbrs,sip) = ∅ ]
33.       [[nbrs := setNS(nbrs,sip,"Full")]] [[lsa := newLSA(ip,now,nbrs)]]
34.       [[lsdb := install(lsdb,{lsa})]] [[nbrs := updRXMTS(nbrs,{lsa})]]
35.       send(sndmsg(upd({lsa},ip),floodNIPS(nbrs))) . OSPF(ip,nbrs,lsdb,hello_t) )
36.   + [ exchange_master ]      /* exchange done */
37.     [[nbrs := incDDSQN(nbrs,sip)]] [[nbrs := addREQS(nbrs,sip,lsdb,lsa_hdrs)]]
38.     ( [getREQS(nbrs,sip) ≠ ∅]
39.       [[nbrs := setNS(nbrs,sip,"Loading")]] . OSPF(ip,nbrs,lsdb,hello_t)
40.     + [ getREQS(nbrs,sip) = ∅ ]
41.       [[nbrs := setNS(nbrs,sip,"Full")]] [[lsa := newLSA(ip,now,nbrs)]]
42.       [[lsdb := install(lsdb,{lsa})]] [[nbrs := updRXMTS(nbrs,{lsa})]]
43.       send(sndmsg(upd({lsa},ip),floodNIPS(nbrs))) . OSPF(ip,nbrs,lsdb,hello_t) )
44.   + [ exchange_others ∨ load_others ]      /* sqn mismatch sqn */
45.     SNMIS(sip,ip,nbrs,lsdb,hello_t)
46.   + [ load_duplicate_slave ]      /* duplicate message */
47.     send(sndmsg(genDBD(nbrs,lsdb,sip,ip),{sip})) . OSPF(ip,nbrs,lsdb,hello_t)
48.   + [ load_duplicate_master ]      /* duplicate message */
49.     OSPF(ip,nbrs,lsdb,hello_t) )
```

some DBD messages are not acknowledged. Eventually, in Line 15, the node ip replies with another DBD message, which contains the information of its own lsdb. The check of 13, described above, is formally defined as

$$\texttt{negotiate_slave} := (\text{ns} = \text{“ExStart”}) \land \texttt{is_slave} \land \text{ibit}\,.$$

Lines 16– 19 is the counterpart of Lines 13– 15 in case the node identifies itself as master. The check in Line 16 is defined by

$$\texttt{negotiate_master} := (\text{ns} = \text{“ExStart”}) \land \texttt{is_master} \land (\text{sqn} = \text{ddsqn}) \land \neg\text{ibit}\,.$$

The node does not only check whether the adjacency establishment has started and whether it is the master in the procedure, it also checks that the sequence number of the incoming message matches its expectation and that its bit is set to false. These two conditions intuitively mean that the slave acknowledges the current node as master. In case the guard of Line 16 evaluates to true, the neighbour state for sip is set to “Exchange” and the DD sequence number is incremented (Line 17). Moreover, the DD timer is reset and any LSA header listed in the incoming message that is more recent than the corresponding one in the node’s own LSDB are added to the link state request list (both Line 18). Lastly, the node sends a DBD message back to the sender sip (Line 19). All other cases where the neighbour state is “Init” are ignored (21), the guard of Line 20) is

$$\texttt{negotiate_others} := (\text{ns} = \text{“ExStart”}) \land \neg\texttt{negotiate_slave} \land \neg\texttt{negotiate_master}\,.$$

Lines 22 –25 handle situations where a DBD message is received a second time, which can happen due to loss of messages acknowledging other messages. If the slave receives a DBD message a second time ($\text{sqn} \leq \text{ddsqn}$) from the master, it assumes that its reply, sent in Line 15, was lost and resends the message (Line 23). Formally the guard is

$$\texttt{exchange_duplicate_slave} := (\text{ns} = \text{“Exchange”}) \land \texttt{is_slave} \land (\text{sqn} \leq \text{ddsqn})\,.$$

In case the master receives a duplicate of a message, checked in Line 24 using

$$\texttt{exchange_duplicate_master} := (\text{ns} = \text{“Exchange”}) \land \texttt{is_master} \land (\text{sqn} < \text{ddsqn})\,,$$

no actions are required and the message is disregarded. In Line 26 the node, which partakes in a adjacency-establishment procedure as slave, handles a message from the master as answer to the DBD message sent earlier.

$$\texttt{exchange_slave} := (\text{ns} = \text{“Exchange”}) \land \texttt{is_slave} \land (\text{sqn} = \text{ddsqn}{+}1) \land \neg\text{ibit}$$

It enters the next phase of the adjacency-forming procedure, called “ExchangeDone” in the RFC. In this phase the slave receives a DBD message presenting the LSDB of the master. As a consequence, DD sequence number is incremented and the DD timer is reset (Line 27). Moreover, any LSA header that is more recent than the corresponding on in the node’s own database is added to the link state request list (Line 28). After the node updated its local data, it sends another DBD message back to the master sip (Line 29).

After that, the node checks its request list, i.e. it checks whether there are some LSAs in the master’s LSDB that are needed by ip. In case there are such LSAs (Line 30) the node enters the next stage of the adjacency establishment by setting the neighbour state to “Loading”, and returns to main process OSPF, waiting for sip to send the necessary data. In case there are no LSAs in the master’s LSDB that are needed by the node itself (Line 32), the slave sip considers the establishment procedure finished and sets the status to “Full” (Line 33). It also updates its own data structures by generating a new LSA (Line 32), installing it in its own lsdb (Line 33) and flooding it out (Line 35).

If the master receives a DBD message from the slave containing information it is waiting for, i.e.

$$\texttt{exchange_master} := (\text{ns} = \text{“Exchange”}) \wedge \texttt{is_master} \wedge (\text{sqn} = \text{ddsqn}) \wedge \neg\text{ibit}\,,$$

it enters the next phase of the adjacency-forming procedure, which again is called “ExchangeDone” in the RFC. In that case, the master increments the corresponding sequence number and updates its link state request list, as usual (Line 37). Similar to the previous case the node checks whether there are some LSAs in the slave's LSDB that the master requires. If there are LSAs missing, i.e. the request list is not empty (Line 38), the neighbour state is changed to “Loading” (Line 39), indicating that the slave will send the missing information in the future. In case there are no LSAs missing (Line 40), the master considers the adjacency-establishment procedure to be done, indicates this by setting the neighbour status to “Full” (Line 41) updates its data structures in the same way as the slave did (see above), and, identical to the slave, informs its neighbours about the new available information, by sending out an LSA (Line 43). Afterwards, the node returns to the main process `OSPF`.

In case that during the establishment procedure either the slave or the master receives a message out of order, which means that the sequence number in the DBD message does not fit expectations, the process `SNMIS` is called, which handles this mismatch. To check whether the incoming message is out of order we use, in Line 44,

$$\begin{aligned}\texttt{exchange_others} := (\text{ns} = \text{“Exchange”}) &\wedge \neg\texttt{exchange_slave} \wedge \neg\texttt{exchange_master}\\ &\wedge \neg\texttt{exchange_duplicate_master} \wedge \neg\texttt{exchange_duplicate_slave}\,.\end{aligned}$$

Lines 46–49 handle the cases where a slave (master) receives a message it handled before while being in phase “Loading” or “Full” of the adjacency-establishment procedure. For the slave and the master this is checked by

$$\begin{aligned}\texttt{load_duplicate_slave} &:= (\text{ns} \geq \text{“Loading”}) \wedge \texttt{is_slave} \wedge (\text{sqn} \leq \text{ddsqn}) \wedge \neg\text{ibit} \wedge \text{dd_t} \geq \texttt{now} \text{ and}\\ \texttt{load_duplicate_master} &:= (\text{ns} \geq \text{“Loading”}) \wedge \texttt{is_master} \wedge (\text{sqn} < \text{ddsqn}) \wedge \neg\text{ibit}\,,\end{aligned}$$

respectively. If the slave receives a duplicate (Line 46), it responds with the same DBD message it had sent before in Line 29, assuming that its message was lost; this happens in Line 47. In case the master handles a duplicate (Line 48) then no actions are required, the node returns to the main process `OSPF` (Line 49).

The only remaining case is the neighbour status is “Loading” or “Full”, but the message is neither expected nor a duplicate. In short that means that something went wrong during the adjacency-establishment procedure. We check this situation using the guard

$$\texttt{load_others} := (\text{ns} \geq \text{“Loading”}) \wedge \neg\texttt{load_duplicate_slave} \wedge \neg\texttt{load_duplicate_slave}\,.$$

As in the case of an out-of-order message during the message-exchange phase (Lines 44 and 45), the process calls `SNMIS` to handle this situation. As these two scenarios are handled in the same way, we use the combined guard $\texttt{load_others} \vee \texttt{exchange_others}$ in Line 44.

Figure 4 summarises the (somewhat complicated) situation of sending messages forth and back between master and slave. It illustrates a possible scenario in which DBD messages are exchanged by two routers. Time increases from the top to the bottom, each arrow indicates a message sent, the words above each arrow represent the message type and contents and the words on the left and the right of the figure represent the node's neighbour status of the other node.

The process `SNMIS` (Process 11) contains all actions that are taken when something goes wrong during the adjacency-forming procedure for a node and its neighbouring router sip. Such an error, called “SeqNumberMismatch” in the RFC, effectively restarts the adjacency-establishment procedure.

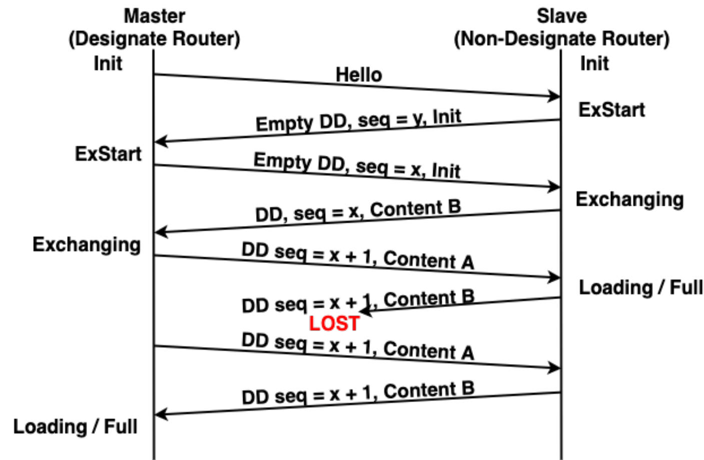

Figure 4: Communication between master and slave during adjacency establishment

Process 11 Sequence Numbers Mismatches

$\mathrm{SNMIS}(\mathrm{sip},\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t}) \stackrel{def}{=}$

1. $[\![\mathrm{nbrs} := \mathtt{initNBR}(\mathrm{nbrs},\mathrm{sip},\text{“ExStart”})]\!]$
2. $[\![\mathrm{nbrs} := \mathtt{incDDSQN}(\mathrm{nbrs},\mathrm{sip})]\!]$
3. $[\![\mathrm{nbrs} := \mathtt{setDDT}(\mathrm{nbrs},\mathrm{sip},\mathtt{now}{+}\mathtt{rxmt_intvl})]\!]$
4. $\textbf{send}(\mathtt{sndmsg}(\mathtt{genDBD}(\mathrm{nbrs},\mathrm{lsdb},\mathrm{sip},\mathrm{ip}),\{\mathrm{sip}\}))$.
5. $\mathtt{OSPF}(\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$

The process resets the status of the adjacency-establishment to “ExStart” (Line 1), increments the DD sequence number to indicate a newer forthcoming message (Line 2) and, as usual, resets the DD timer (Line 1). It then restarts the establishment procedure by sending out a new DBD message, similar to Line 11 of Process 10.

The process REQ (Process 12) involves all actions that are taken upon the receipt of a LSR message.

Process 12 Handling LSR messages

$\mathtt{REQ}(\mathrm{lhdr},\mathrm{sip},\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t}) \stackrel{def}{=}$

1. /* the sender sip is “unknown” */
2. $[\ \neg\mathtt{nbrExist}(\mathrm{nbrs},\mathrm{sip}) \lor \mathtt{getNS}(\mathrm{nbrs},\mathrm{sip}) < \text{“Exchange”} \lor \neg\mathtt{lsaExist}(\mathrm{lsdb},\mathrm{lhdr})\ ]$
3. $\quad\mathtt{OSPF}(\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$
4. /* the sender sip is “known” */
5. $+\ [\ \mathtt{nbrExist}(\mathrm{nbrs},\mathrm{sip}) \land \mathtt{getNS}(\mathrm{nbrs},\mathrm{sip}) \geq \text{“Exchange”} \land \mathtt{lsaExist}(\mathrm{lsdb},\mathrm{lhdr})\ ]$
6. $\quad\textbf{send}(\mathtt{sndmsg}(\mathtt{upd}(\{\mathtt{getLSA}(\mathrm{lsdb},\mathrm{lhdr})\},\mathrm{ip}),\{\mathrm{sip}\}))\ .\ \mathtt{OSPF}(\mathrm{ip},\mathrm{nbrs},\mathrm{lsdb},\mathrm{hello_t})$

It is nearly identical to the simplified version, depicted in Process 4, only the guards are more complex as they have to reflect the status of the adjacency-establishment procedure. Line 2 checks whether the sender of sip of the LSR message is unknown, is in a premature state (ns < “Exchange”) where is should not send a request, or whether more recent information was already received. In all these cases the message is dropped (Line 3). Line 5 is the complimentary guard: the node handles the request, looks up information in its own link state database and replies with an LSU message to sip.

The process UPD (Process 13) includes all actions that should be taken upon receipt of a LSU message.

Process 13 Handling LSU messages

```
UPD(lsas,sip,ip,nbrs,lsdb,hello_t) ≝
1.  [ ¬nbrExist(nbrs,sip) ]        /* sender sip unknown */
2.    OSPF(ip,nbrs,lsdb,hello_t)
3.  + [ nbrExist(nbrs,sip) ]        /* sender sip known */
4.    send(sndmsg(ack({hdr(lsa) | lsa ∈ lsas},ip),{sip})) .
5.    [[lsas := {lsa | lsa ∈ lsas ∧ ∄lsa′ ∈ lsdb : hdr(lsa) ≤ hdr(lsa′)}]]        /* LSAs to update LSDB */
6.    (
7.      [ lsas = ∅ ]
8.        OSPF(ip,nbrs,lsdb,hello_t)
9.      + [ lsas ≠ ∅ ]
10.       [[lsdb := install(lsdb,lsas)]]
11.       [[nbrs := cleanREQS(nbrs,sip,lsdb)]]
12.       [[nbrs := updRXMTS(nbrs,lsas)]]
13.       send(sndmsg(upd(lsas,ip),floodNIPS(nbrs))) .
14.       (
15.         [ getREQS(nbrs,sip) ≠ ∅ ]
16.           OSPF(ip,nbrs,lsdb,hello_t)
17.         + [ getREQS(nbrs,sip) = ∅ ]
18.           [[nbrs := setNS(nbrs,sip,“Full”)]]
19.           [[lsa := newLSA(ip,now,nbrs)]]
20.           [[lsdb := install(lsdb,{lsa})]]
21.           [[nbrs := updRXMTS(nbrs,{lsa})]]
22.           send(sndmsg(upd({lsa},ip),floodNIPS(nbrs))) .
23.           OSPF(ip,nbrs,lsdb,hello_t)
24.       )
25.   )
```

As in many other processes we first handle the case that the sender of the message is unknown (Line 1); as usual the message is ignored.

Lines 3 to 25 handle the situation where sip is known. First, in Line 4 the node replies to sip by sending an ACK message. This messages contains all headers of the LSAs received. The node then compares the LSAs received with the ones in its own link state database lsdb; it stores those LSA (in the variable lsas) that are newer than the one in the lsdb (or non-existent in lsdb); these are the LSA that need to be handled. In case there is none (Line 7), no action is required and the protocol returns to the main process OSPF in Line 8. If there are LSAs (Line 9), the node installs these newly discovered LSAs in its own database (Line 10), removes old data (LSAs that are dominated by others) from the data structure, and informs its neighbours about this new information by sending out an LSU message (Line 13).

After the updates are performed the node ip checks whether there are entries in the request list that need its attention. In case there are some (Line 15), the process returns to OSPF to await a new incoming LSU message that address some of these LSAs. Otherwise the list is empty (Line 17). That means that the node requires nothing more from its neighbour; it sets the status to “Full” (Line 18), updates its own data structures (Lines 19 to 21), and sends a new update to the neighbours interested in this information (Line 22).

The process ACK handles incoming Link State Acknowledgement (LSA) messages.

Process 14 Handling ACK messages

$\texttt{ACK}(\text{lsa_hdrs}, \text{sip}, \text{ip}, \text{nbrs}, \text{lsdb}, \text{hello_t}) \stackrel{def}{=}$

1. $[\ \neg\texttt{nbrExist}(\text{nbrs}, \text{sip})\ ]$ /* the sender sip is unknown */
2. $\quad \texttt{OSPF}(\text{ip}, \text{nbrs}, \text{lsdb}, \text{hello_t})$
3. $+\ [\ \texttt{nbrExist}(\text{nbrs}, \text{sip})\ ]$ /* sip is a known neighbour */
4. $\quad [\![\text{nbrs} := \texttt{cleanRXMTS}(\text{nbrs}, \text{sip}, \text{lsa_hdrs})]\!]$
5. $\quad \texttt{OSPF}(\text{ip}, \text{nbrs}, \text{lsdb}, \text{hello_t})$

As in most of the other processes, the message is dropped in case the node ip does not know the originator and sender of the ACK message (Lines 1 and 2). In case the sender in known (Line 3) the only action that remains is to remove sip from the link state transmission list, which is done by `cleanRXMTS` in Line 4.

C.2.3 Message Queues

Processes QMSG and QSND are the same as those in the simple model, see Processes 6 and 7.

D The Uppaal Model

In this section we provide a description of our Uppaal model that was presented in Section 6. The main component of the Uppaal model is a timed automaton named `ospf`. This automaton corresponds to the process `OSPF` of Section 5, including all of its subprocesses (`HELLO`, `DBD`, etc.). The outgoing messages queue (Process `QSND`), however, has its own corresponding automaton `qsnd`. The purpose of creating a separate automaton for `qsnd` is so that the sending of a message does not block a node performing other actions while it waits for another node to be ready for synchronisation. Each automaton `ospf` has an associated automaton `qsnd` that handles its outgoing messages. As a slight optimisation we are able to avoid another automaton for the input queue. A network of n nodes is modelled by n copies of the two automata; each tagged with a unique identifier. We use the numbers 1 to n as identifiers, we avoid the use of 0 as it is the default initialisation of integers in Uppaal.

The Uppaal file is split into six sections: **(i)** The first section declares global variables, data structures and defines the topology of the network being modelled. **(ii)** The second defines and declares what is required locally for each automaton `ospf`; this includes function definitions and data structure declarations. **(iii)** A short part that defines and declares what is required for each automaton `qsnd`. **(iv)** A section for declaring and initialising multiple (independent) instances of all automata involved. **(v)** The declaration of the automaton `ospf`, using functions and data structures defined in Part (ii). **(vi)** The automaton for the sending queue, using data structures and functions from Part (iii).

D.1 Global Declarations

This section, copied verbatim from the Uppaal model, declares the global data structure. This includes constants (e.g. the size of the network), type definitions (e.g. message types), the node's own data structure (e.g. the neighbour structure), and other concepts such as LSA headers and LSAs. It also defines a connectivity matrix to define the topology of the network, as well as channels for sending messages.

```
1  // GLOBAL DECLARATIONS
2  // -------------
3  // CONSTANTS
4  const int N=3+1;           // number of nodes +1
5  const int M=10;            // max sequence number +1
6  const int time_sending=1; // minimum time taken to send a message
7  const int time_spread=0;  // max time to send message is time_sending + time_spread
8  const int hellointvl=10;  // time between sending hello messages
9  const int rtdeadcount=5;  // number of hellos used to decide neighbour liveness
10 const int start_interval = 10; // max wait time before a node initialises the
11                                // ospf automaton.  All routers will be booted-up
12                                // nondeterministically within the start_interval
13 const int age_bound = 2 * N;    // upper bound for age of an LSA.
14
15 //TYPEDEFs
16 typedef int[1,N-1] IP;     // range of IP addresses in the given topology
17
18 typedef int[0,4] MSGTYPE; // definition of message types
19     const MSGTYPE NONE=0;
20     const MSGTYPE HELLO=1;
21     const MSGTYPE DBD=2;
22     const MSGTYPE REQ=3;
23     const MSGTYPE UPD=4;
24
```

```
25 typedef struct // neighbour structure contains a bounded integer (inactivity timer)
26 {
27   int[0,rtdeadcount] inactivity_timer;
28 } NBR;
29 typedef struct // LSA header contains IP of the originating router and LSA age.
30 {
31   int[0,N-1] ip;
32   int[0,age_bound+1] age; // age acts as a wrap-around sequence number
33 } LSAHDR;
34
35
36 typedef struct // LSA is identical to LSA header since the extra info is unused.
37 {              // We distinguish them as we wish to model the DBD exchange procedure
38   int[0,N-1] ip;
39   int[0,age_bound+1] age;
40 } LSA;
41
42 typedef struct {   // definition of messages
43   MSGTYPE msgtype; // the type of message
44   int[0,N-1] sip;  // sender IP address, used for all message types
45   bool ips[N];     // discovered neighbours, used for HELLO
46   LSAHDR hdrs[N];  // used for DBD and REQ messages
47   LSA lsas[N];     // used for UPD messages
48 } MSG;
49
50 // meta variables for copying messages.  Sender copies from xlocal to xglobal
51 // receiver copies from xglobal to xlocal.
52 meta MSG msgglobal;          // this variable is for the message contents
53 meta bool destsglobal[N];    // this contains the intended destinations of a message
54
55 // CHANNELS
56 urgent chan imsg[N];                    // internal communication
57 broadcast chan bcast[N], gcast[N];      // channels for broadcasting and groupcasting
58 chan ucast[N][N];                       // channel for sending a unicast
59 urgent broadcast chan tau[N];           // used to prioritise internal transitions
60
61 bool send_idle[N]={1,1,1,1}; // indicates if a qsnd automaton is busy sending
62                              // or else is ready to receive a message from ospf
63
64 const bool topology[N][N]={    // topology described in terms of connectivity
65   {0,0,0,0},
66   {0,0,1,0},
67   {0,1,0,1},
68   {0,0,1,0}
69 };
70
71 bool isconnected(IP i, IP j){
72   return(topology[i][j]==1);
73 }
```

D.2 Template `ospf`

The declaration of the automaton `ospf` is presented below. It first declares the data structures that are used (neighbour structures, LSDB, input queue, output queue, etc.). Afterwards it defines functions used by `ospf` for manipulating these data structures.

```
1  // ospf DECLARATIONS
2  // ---------------
3  // local data structure of ospf automaton
4  NBR nbrs[N];          // array of neighbours that have been discovered
5  LSA lsdb[N];          // array of LSAs (which make up the LSDB)
6  bool upd_required; // if a neighbouring router has died we must update the LSDB
7  int[1,age_bound+1] age = 1; // age of the last self-originated LSA
8
9  // internal datastructure for Uppaal
10 clock local_time; //local time (used to trigger hello messages)
11
12 // in-queue of messages (received from other nodes)
13 MSG msg_inqueue[M];      // last M local copies of incoming messages,
14 int buffersize_inqueue; // number of messages currently in the in queue
15
16 // out-queue of messages to be sent
17 MSG msg_outqueue[M];          // Last M local copies of outgoing messages
18 bool dests_outqueue[M][N];   // boolean matrix of destinations for outgoing messages
19 int buffersize_outqueue;     // number of messages currently in the out-queue
20
21 // OUT QUEUE FUNCTIONS
22 // -------------
23 // add message (and its destinations) to the out queue
24 // doesn't check if there is space in the queue; cause overflow
25 void addmsg_outqueue(MSG msg, bool dests[N]){
26   dests_outqueue[buffersize_outqueue]=dests;
27   msg_outqueue[buffersize_outqueue]=msg;
28   buffersize_outqueue++;
29 }
30
31 // returns type of next message to be sent.  If NONE, the out queue is empty
32 MSGTYPE nextmsg_outqueue(){
33   return msg_outqueue[0].msgtype;
34 }
35
36 // delete message from the out-queue and shift all other messages forward one place.
37 void deletemsg_outqueue(){
38   MSG empty;
39   bool emptyd[N];
40   for(i: int[1,M-1]){                   // move all messages forward one position
41     msg_outqueue[i-1]=msg_outqueue[i];
42     dests_outqueue[i-1]=dests_outqueue[i];
43   }
44   msg_outqueue[M-1]=empty;             // write empty message end of the queue
45   dests_outqueue[M-1]=emptyd;
46   buffersize_outqueue--;
47 }
48
49 // IN QUEUE FUNCTIONS
50 // ------------
51 // add a message to the in queue
52 // no check for space, may cause overflow
53 void addmsg_inqueue(MSG msg){
54   msg_inqueue[buffersize_inqueue]=msg;
55   buffersize_inqueue++;
56 }
```

```
57 // returns type of next message, if it returns NONE, we know the in queue is empty
58 MSGTYPE nextmsg_inqueue(){
59   return msg_inqueue[0].msgtype;
60 }
61
62 // delete message from the in-queue and shift all other messages forward one place.
63 void deletemsg_inqueue(){
64   MSG empty;
65   for(i: int[1,M-1]){                  // move all messages by one position
66     msg_inqueue[i-1]=msg_inqueue[i];
67   }
68   msg_inqueue[M-1]=empty;             // write an empty message to end of the queue
69   buffersize_inqueue--;
70 }
71
72 // MISC
73 // ---
74 void initialise(){ // when a node boots-up it will its initial LSA to the LSDB
75   lsdb[ip].ip = ip;
76   lsdb[ip].age = 1;
77 }
78
79
80 // OSPF SPECIFIC
81 // ---------
82 // create new LSA for own IP
83 LSA newLSA(){
84   LSA lsa;
85   lsa.ip = ip;
86   age==age_bound?age=1:age++; // wrap around age if exceeds age_bound
87   lsa.age = age;
88   return lsa;
89 }
90
91 // returns true if age1 is newer than age2 (handles wrap around)
92 bool newer_age(int age1, int age2){
93   if (age1==0) { return false; }
94   else if (age2==0) { return true; }
95   if ((age2 > age1 && 2*(age2 - age1) < age_bound) // wrap around handling
96       ||
97      (age1 > age2 && 2*(age1 - age2) > age_bound))
98       return false;
99   return true;
100 }
101
102 void install_lsa(LSA lsa){ // install an LSA in LSDB if newer than current copy
103   if(newer_age(lsa.age,lsdb[lsa.ip].age))
104       lsdb[lsa.ip]=lsa;
105 }
106
107 void install_lsas(LSA lsas[N]){ // install multiple LSAs into the LSDB if newer
108   for(i : IP){
109     if(newer_age(lsas[i].age,lsdb[i].age))
110       lsdb[i]=lsas[i];
111   }
112 }
```

```
113 void reduce_lsas_msglocal(){ // check array of LSAs against LSDB, return newer LSAs
114   LSA fresh;
115   for(i : IP)
116     if(newer_age(lsdb[i].age, msg_inqueue[0].lsas[i].age))
117       msg_inqueue[0].lsas[i] = fresh;
118 }
119
120 bool lsas_is_empty(){ // after reduce_lsas_msglocal, check if any LSAs remain
121   LSA fresh[N];
122   return(fresh==msg_inqueue[0].lsas);
123 }
124
125 bool nbrsExist(){ // check if sender of current message is active
126   return(nbrs[msg_inqueue[0].sip].inactivity_timer>0);
127 }
128
129 void reduce_lifetime_nbrs(){ // after a hello is sent, reduce all inactivity timers
130   for (i : IP){
131     if(nbrs[i].inactivity_timer > 0){
132       nbrs[i].inactivity_timer--;
133       if(nbrs[i].inactivity_timer == 0){ // if timer is zero, set upd_required flag
134         upd_required=true;
135       }
136     }
137   }
138 }
139
140
141 // CREATE MESSAGES
142 // ----------
143 // create hello message and add it to out queue
144 void generate_send_hello(){
145   MSG msg;
146   bool dests[N];
147   msg.msgtype=HELLO;
148   msg.sip=ip;
149   for(i:IP)
150     if(nbrs[i].inactivity_timer!=0)
151       msg.ips[i]=true;
152   addmsg_outqueue(msg,dests);
153 }
154
155 // create dbd message, add to out queue
156 void generate_send_dbd(IP dip){
157   MSG msg;
158   bool dests[N];
159   msg.msgtype=DBD;
160   msg.sip=ip;
161   for(i:IP){
162     msg.hdrs[i].ip = lsdb[i].ip;
163     msg.hdrs[i].age = lsdb[i].age;
164   }
165   dests[dip]=true;
166   addmsg_outqueue(msg,dests);
167 }
168
```

```
169 void generate_send_upd_single(LSA lsa){//create single-LSA upd msg, put in out queue
170   MSG msg;
171   bool dests[N];
172   msg.msgtype=UPD;
173   msg.sip=ip;
174   msg.lsas[lsa.ip] = lsa;
175   for(i : IP)
176     dests[i] = (nbrs[i].inactivity_timer!=0);
177   addmsg_outqueue(msg,dests);
178 }
179
180 void generate_send_upd(LSA lsas[N]){ // create (multi-LSA) upd msg, add to out queue
181   MSG msg;
182   bool dests[N];
183   msg.msgtype=UPD;
184   msg.sip=ip;
185   msg.lsas = lsas;
186   for(i : IP)
187     dests[i] = (nbrs[i].inactivity_timer!=0);
188   addmsg_outqueue(msg,dests);
189 }
190
191 // create and process an upd message in response to a request message
192 void generate_send_upd_requested(IP dip){
193   MSG msg;
194   bool dests[N];
195   msg.msgtype=UPD;
196   msg.sip=ip;
197   for(i:IP)
198     if(msg_inqueue[0].hdrs[i].ip == i &&
      newer_age(lsdb[i].age,msg_inqueue[0].hdrs[i].age))
199       msg.lsas[i] = lsdb[i];
200   dests[dip]=true;
201   addmsg_outqueue(msg,dests);
202 }
203
204 // after receiving a dbd message create a req message to request required LSAs
205 void generate_send_req(IP dip){
206   MSG msg;
207   bool dests[N];
208   LSAHDR fresh[N];
209   msg.msgtype=REQ;
210   msg.sip=ip;
211   for(i : IP)
212     if(newer_age(msg_inqueue[0].hdrs[i].age, lsdb[i].age))
213       msg.hdrs[i] = msg_inqueue[0].hdrs[i];
214
215   if(msg.hdrs!=fresh){   // only send if request is non-empty
216     dests[dip]=true;
217     addmsg_outqueue(msg,dests);
218   }
219 }
```

D.3 Template `qsnd`

The declaration of the automaton `qsnd` is much shorter, but structured in the same way as `ospf`.

```
// qsnd DECLARATIONS
// -----------
MSG msglocal;
bool dests[N];
clock clk;

int getDip(){ // retrieve destination of message to be sent
  for(i : IP)
    if(dests[i]) return i;
  return 0;
}

void clean_dests(){ // after sending clear the destination from the global store
  bool empty[N];
  destsglobal=empty;
}
```

D.4 System Declarations

The system declarations are used to instantiate the required number automata. In our example, the instantiation of a network consists of three nodes each running both `ospf` and `qsnd`.

```
// system DECLARATIONS
// -------------
// Place template instantiations here.
a1ospf=ospf(1); // instantiate a ospf automaton called a1ospf
a1send=qsnd(1); // instantiate a qsnd automaton called a1send

a2ospf=ospf(2);
a2send=qsnd(2);

a3ospf=ospf(3);
a3send=qsnd(3);

system a1ospf,a1send,a2ospf,a2send,a3ospf,a3send;
```

D.5 The Automaton `ospf`

The automaton `ospf` is presented in Figure 5. The initial state is represented as the node at the top of the image with only a single, downward, outgoing transition. The automaton leaves this node after a non-deterministic amount of time, bounded by `start_interval`.

The node at the centre represents a state where the automaton is ready to handle any message or action required. There are eight loops consisting of multiple transitions each, that surround the centre. We address them in a clockwise fashion beginning from one o'clock. The first three loops handle HELLO, DBD and LSR messages (here called `REQ`). It is easy to see how the synchronisation and updates performed correspond to the actions of our simple model, explained in Appendix B. The fourth loop is executed when a neighbouring node dies and therefore a new LSA needs to be generated. The fifth pushes a message from the output queue to the relevant automaton `qsnd`. The sixth deals with receiving

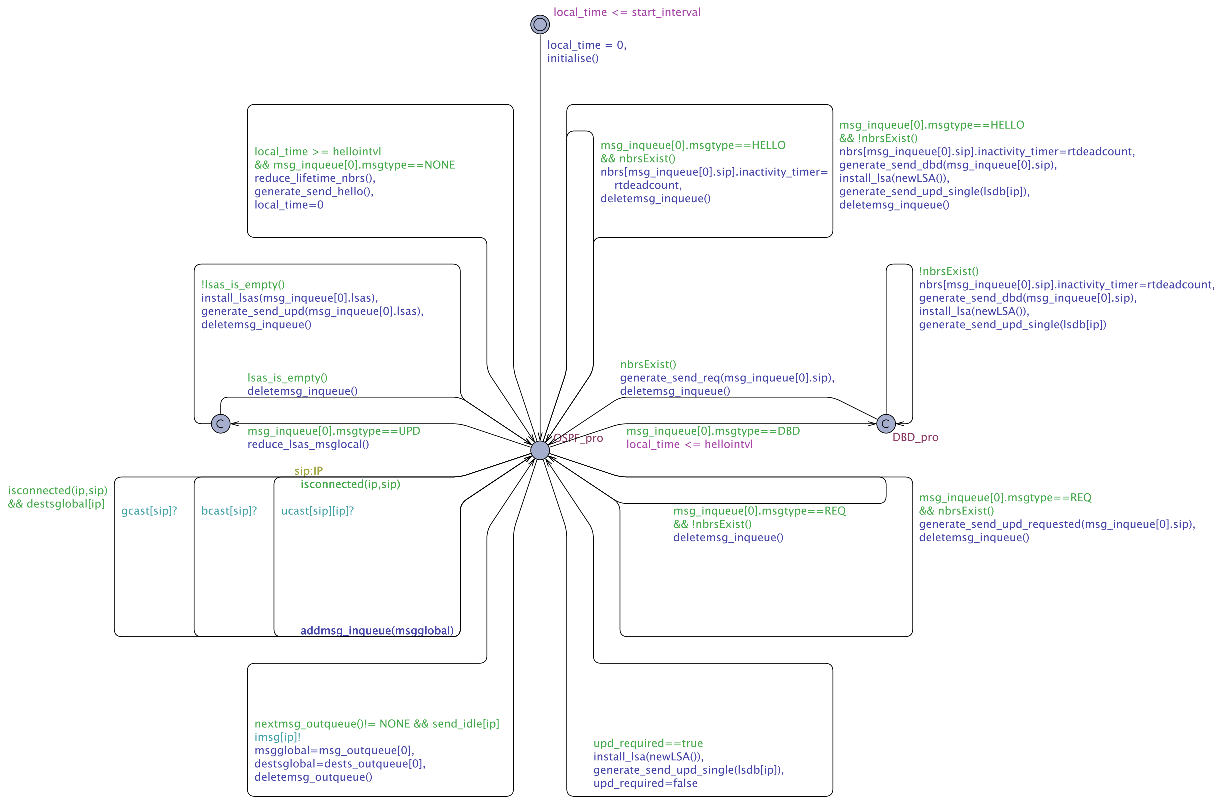

Figure 5: The timed automaton `ospf`

a message and placing it into the input queue. The seventh handles a LSU message (here called UPD). The eighth and final transition loop handles the generation of a HELLO message at the correct time.

D.6 The Automaton `qsnd`

The `qsnd` automaton is presented in Figure 6. Its purpose is to hold messages until other automata are ready to receive them. Its initial state is at the centre. The top right quadrant broadcasts `HELLO` messages. The top left quadrant unicasts `DBD` and `REQ` messages. The bottom left quadrant groupcasts `UPD` messages. The bottom right quadrant receives a message from the corresponding `ospf` automaton.

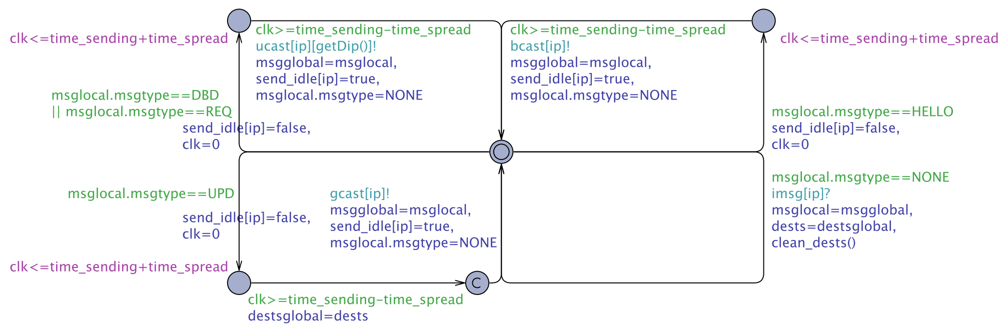

Figure 6: The timed automaton `qsnd`