

Advanced Models for the OSPF Routing Protocol

Courtney Darville

Data61, CSIRO, Sydney, Australia
University of New South Wales, Sydney, Australia
C.Darville@unsw.edu.au

Peter Höfner

School of Computing
ANU, Canberra, Australia
Peter.Hoefner@anu.edu.au

Franc Ivankovic

University of Trento, Trento, Italy
francekivankovic@gmail.com

Adam Pam

Macquarie University, Sydney, Australia
Adam.Pam@students.mq.edu.au

We present two formal models for the OSPF routing protocol, designed for the model checker UPPAAL. The first one is an optimised model of an existing model that allows to check larger network topologies. The second one is a specialised model for adjacency building, a complex subprocedure of OSPF, which is not part of any existing model and which is known to be vulnerable to cyber attacks. We illustrate how both models can be used to discover vulnerabilities in routing protocols.

1 Introduction

The Open Shortest Path First (OSPF) protocol [22] is a widely used, proactive, link-state routing protocol that distributes routing information throughout a single autonomous system. Similarly to many other routing protocols, it aims to establish shortest paths between network routers while keeping network overhead — messages carrying routing information rather than user data — to a minimum. Although OSPF, as many other routing protocols, is based on ‘simple’ algorithms such as Dijkstra’s shortest path algorithm [9], it seems incredibly hard to ensure that the protocol is functionally correct, see e.g. [21, 14].

To allow formal analysis of OSPF, three detailed formal models have been presented in [10]:¹ two are formalised in the timed process algebra T-AWN [5], 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. The more abstract model is the basis for the third, which presents OSPF as a network of timed automata that can be executed in the model checker UPPAAL [18, 2].

The process-algebraic models do not allow automatic analysis of the protocol, for there is currently no tool support available for T-AWN. UPPAAL, however, provides several tools, from simulation to model checking. Unfortunately, the current model faces a combinatorial blow up of the state space that must be addressed before automatic analysis of the protocol becomes feasible. In fact, the current model is already infeasible when analysing networks consisting of three nodes only.

Based on the existing UPPAAL-model, we advance the modelling efforts for OSPF:

1. We develop a UPPAAL model that is behaviourally equivalent to the model presented in [10] but with a much smaller state space.
2. We develop a new model for adjacency building, a subprocedure of OSPF, which so far has been ignored in formal models.
3. We develop flexible models of adversaries that cover many different attack scenarios. Although the modelling is straightforward, it is extremely powerful and can be used to discover vulnerabilities.
4. We illustrate how this adversarial model can be used in combination with both new models to discover vulnerabilities in OSPF.

¹Some other models of OSPF exist [24, 19]; see Section 8 and [10] for a discussion.

2 The Open Shortest Path First (OSPF) Protocol

The *Open Shortest Path First (OSPF) protocol* [22] is a link-state protocol that falls into the group of interior gateway protocols, operating within a single autonomous system. As with any link-state protocol, OSPF routers (nodes in a network) exchange topological information about one-hop links with their neighbours, i.e. the nodes within transmission range. That information is distributed through the network so that (eventually) every router has a complete picture of all available links. This knowledge is used to calculate the shortest/best route between any two nodes, using a variant of Dijkstra’s algorithm [9].

To discover immediate (one-hop) neighbours, OSPF uses *HELLO messages*, which are broadcast periodically by all nodes. HELLO messages consist of the sender’s unique identification (ID) and a list of all known one-hop neighbours of the sender, i.e. nodes from which the sender has received HELLO messages. Upon receipt of a HELLO message, a node updates its own data structures accordingly. The receiver consequently drops that HELLO message, i.e. HELLO messages are not forwarded.

A node distributes information about its known connections using *Link State Advertisements (LSAs)*.

A node stores data concerning connections between other nodes in their *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.

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

The term *adjacency* describes the relationship between two neighbouring nodes that exchange all of their topological information — the content of their LSDBs. After information has been shared between two adjacent nodes, they have an identical understanding of the entire network. Nodes not forming an adjacency do not share this information and hence network overhead is reduced. When two nodes recognise that they need to form an adjacency, the node with the larger node ID becomes the *master*, and the other becomes the *slave*. It is the master that initiates the exchange of data by sending the necessary information using *Database Description (DBD) messages*. After that, the two nodes exchange a sequence of DBD messages. We describe further details about adjacencies in Section 5.

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, it sends *Link State Request (LSR) messages*, asking for LSAs containing the newest available information. 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 adjacencies to further reduce network traffic. As they are not crucial to the understanding of the main functionality of OSPF, we omit the details.

3 Modelling OSPF in UPPAAL

One of our aims is to analyse OSPF automatically, both with regards to functional correctness and to discovery of vulnerabilities. A standard tools for analysing and verifying systems automatically is model checking. As there are existing models for UPPAAL, we have chosen to stick with that model checker.

UPPAAL [2, 18] is an established model checker, which is frequently used for protocol verification, e.g. [3, 26, 25, 13, 12]. UPPAAL analyses *networks of timed automata*, with clocks supporting the




Figure 1: Model Network Message Passing: unicast and multicast

modelling and the analysis of temporal aspects. It provides two synchronisation mechanisms: *binary* and *broadcast channels*. In the setting of routing protocols these usually translate to unicast and broadcast communication; the latter presenting the transmission of messages to *all* neighbours of a node. Hence, when using UPPAAL’s broadcast mechanism one must consider the network topology to determine the exact set of nodes that are able to receive messages (see below). UPPAAL also provides common data structures, such as arrays, and a C-like programming language to define updates on these data structures.

Networks of Timed Automata The state of the system is determined, in part, by the values of data variables that can be either local or shared between automata. We assume a data structure with several types, variables ranging over these types, operators and predicates. Common Boolean and arithmetic expressions are used to denote data values and statements about them.

The automata are extended with clock variables. UPPAAL uses a dense-time model where a clock variable evaluates to a real number. All the clocks progress synchronously.

Each automaton is a graph, with locations, and edges between locations. Every edge has a guard, optionally a synchronisation label, and an update, which allows local and global data structures to be updated. 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 other transitions use one of the two types of synchronisation.

Synchronisation In *binary handshake* synchronisation, an automaton having an edge with a label that has the suffix $!$ synchronises with another automaton with an edge having the same label with a $?$ -suffix. These two transitions may synchronise if and only if both guards are true in the current state of the system. 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 the set of other automata that all have a matching edge with a $?$ -label. 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 system state. The automaton with the $!$ -edge will update the state first, followed by the other automata in some lexicographic order. If more than one automaton can initiate a transition on an $!$ -edge, the choice will be made non-deterministically.

Network Communication To model the network mechanisms of unicast, broadcast and multicast — sending to a dedicated set of neighbours — we combine UPPAAL’s synchronisation with guards. Message contents are exchanged via global data using UPPAAL’s update mechanism.

For communication within routing protocols, unicast is a node-to-node communication with an in-built acknowledgment-of-receipt mechanism, which is implemented at a lower layer in the network stack.² More formally, unicast in networking means that a node sip can only send to another node

²Usually implemented by the link layer of relevant standards such as IEEE 802.11 [16].

dip if (a) both nodes are within transmission range of each other, (b) dip is ready to receive and (c) after attempting to send a message, the node sip is aware of whether transmission was successful. The corresponding model in UPPAAL is depicted on the top of Figure 1. It uses a unicast channel `uc`. The network topology is modelled by a Boolean predicate `connect`. The sender has two successors, one for successful transmission — as we do not model lossy channels, this can be modelled as `connect(sip,dip)` — and one for failed transmission (`!connect`), which models the actions of a failed sending attempt. The receiving edge is straight forward.

Broadcast in networking means sending a message to all other network nodes within transmission range. Similarly to unicast, this is modelled by UPPAAL’s broadcast mechanism — the broadcast channel is named `bc` — in combination with the predicate `connect`. However, the model has to ensure that all nodes within transmission range are in a state where messages can be received.

A restricted version of broadcast is multicast: a node `sip` tries to transmit a message to destinations `dests`, and proceeds regardless of whether any of the transmissions is successful. The corresponding modelling in UPPAAL is shown on the bottom of Figure 1. Here, `D` is a global data structure, maintaining the set of intended destinations. Before the actual message transfer occurs, the sender `sip` has to update `D`. By doing so, nodes within transmission range of `sip` can determine whether they are intended recipients, using the Boolean function `iselem`.

4 Optimising a Model of OSPF

Drury et al. [10] have created a UPPAAL-model that closely follows the OSPF specification [22]; they also take into account the amendments described in [7]. In that model, each network router is modelled by two automata: one describing the core behaviour of OSPF and the other a queue for outgoing messages.

The main automaton (Figure 2) is built around a ‘central’ location `OSPF_pro`. This location models an idle state of the protocol — the protocol can stay there indefinitely. As soon as a message is received, the automaton identifies the type of message and acts accordingly. For example, the two transitions in the upper right of Figure 2 model the receipt of an incoming HELLO message: the automaton checks whether the content of the message shows the node’s ID (see Section 2) and takes the corresponding transition. The automaton uses a clock to send out HELLO messages periodically (upper left of Figure 2). The receipt of messages does not take time; only the transmission of messages does. When the protocol needs to send a message, this automaton creates the message and passes it, via unicast-synchronisation, to another automaton that models a queue for outbound messages. The queue manages the delivery of the message to corresponding recipients. Each node maintains a data structure for LSDBs, two independent arrays of messages to be sent (as part of the queue automaton) and to be received (part of the main automaton), as well as two clocks, one for each automaton. Each router makes use of four different channels, one for internal communication, one for broadcast, one for unicast, and one for multicast.

As often occurs, this model faces a combinatorial blow up of the state space that must be addressed to analyse the protocol automatically. The state-space explosion is caused by multiple factors:

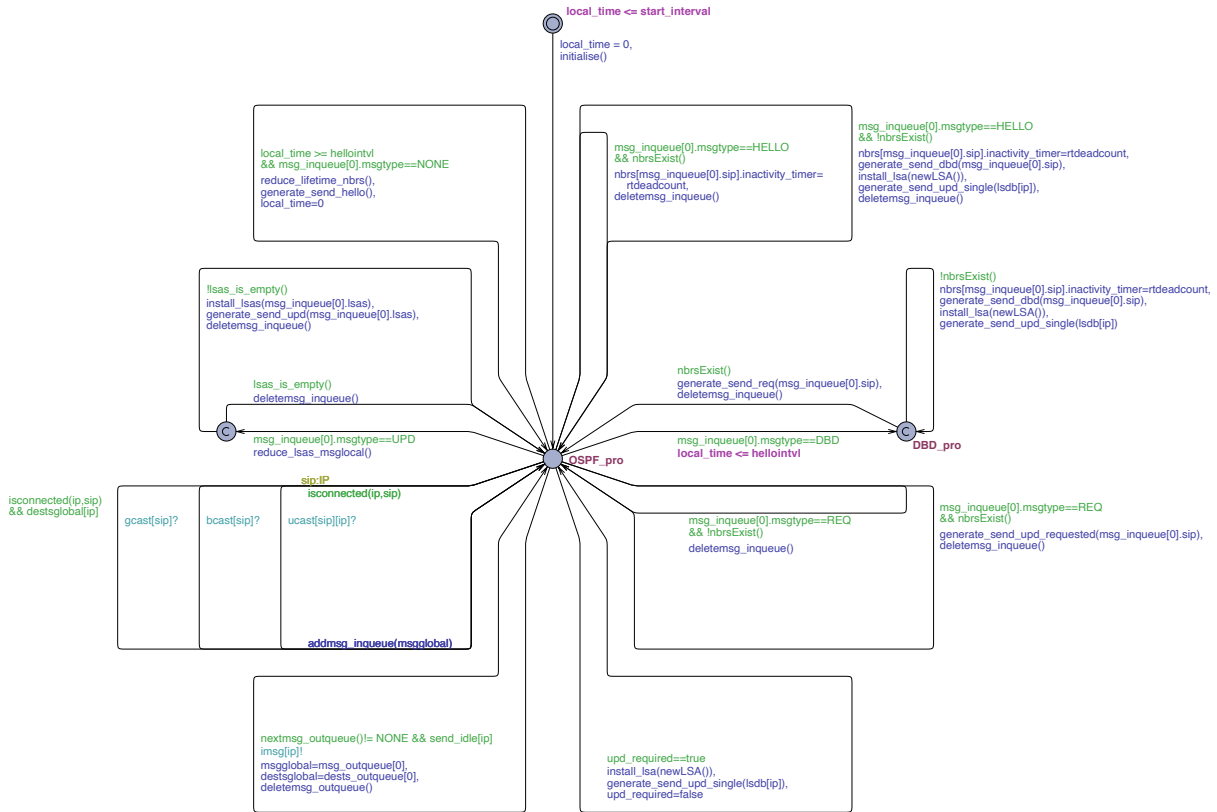
1. Each additional router adds two automata to the network of timed automata.
2. Each additional router adds two clocks to the model, which exponentially increases the complexity.
3. The underlying data structures, messages and message queues significantly add to the state space.

As a result, checking even simple properties, such as deadlock, fails already on topologies of size three.⁴

³The diagrams are intended to show the models’ complexity, not to be ‘read’.

All models are available at <http://mars-workshop.org>.

⁴Deadlock checking failed on a machine with an AMD Ryzen 7 CPU and 64GB RAM.

Figure 2: An OSPF router as a timed automaton³

To overcome (some of) the problems regarding the state space, we create a single automaton describing the entire network. It is depicted in Figure 3.

This single automaton models the nondeterministic interaction between routers, rather than actual routers. Data structures of the individual nodes are organised into multidimensional arrays. This allows a single automaton to gain all knowledge of the network and it is no longer necessary to model the concrete content of OSPF messages. We instead model message transmission as a function that extracts information from the sender's 'local' data structures, such as a the LSDB, and updates the 'local' data structure of the receiver(s) directly.

As the entire model is assembled into a single automaton. We summarise some advantages of our newly developed model, compared to the original model:

1. The content of messages are not modelled; hence a significant reduction in the complexity of the required data structures.
2. Only one clock is required rather than $2 \cdot N$, where N is the number of routers in a network; the model requires an additional small data structure to manage the periodic sending of HELLO messages.
3. Message queues and message propagation delays are removed; the model provides similar behaviour through a choice of nondeterministic transitions.

A simple automaton is added alongside the main automaton to monitor the model's clock and notify the main automaton when HELLO messages need to be sent for each router. Although we did not prove the relationship to the original model of OSPF, the model is carefully designed that it is behaviourally equivalent to the original one.

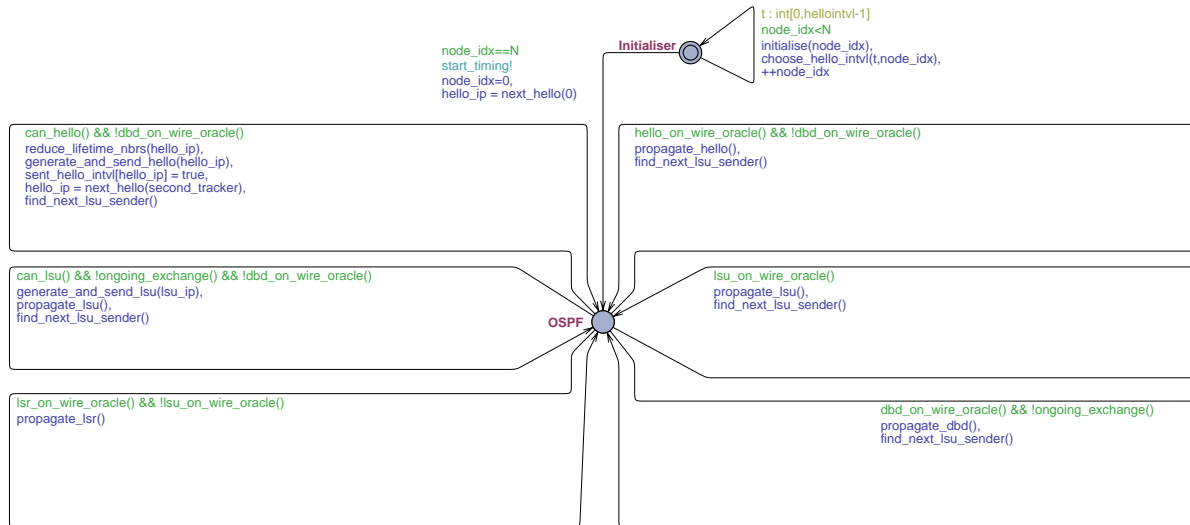


Figure 3: A single automaton for OSPF

In a further abstraction, we can replace the (dense-time) clock by a bounded integer that takes over the periodic sending of HELLO messages. We assume that this model is behaviourally equivalent with regards to message sending and data structures, but obviously not regarding timing properties.

5 Adjacency Building

The original model of OSPF, as well as the model described above, capture all but one core functionality of OSPF. The missing functionality, adjacency building, is an independent subroutine with no immediate effect on other functionality and hence is usually abstracted away. As discussed in Section 2, *adjacency building* describes the activities involved when two neighbouring nodes exchange all of their topological information – the content of their LSDBs. After the exchange, the nodes have an identical understanding of the entire network.

In this section, we present a new, detailed model of adjacency building. When nodes have been assigned to form an adjacency, they progress through the following six *neighbour states*:

Down: No information has been exchanged between the two nodes.

Init: A participating node has recently received a HELLO message from a neighbour. However, bidirectional communication is not confirmed yet – the node’s own IP is not part of the HELLO message.

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

Exchange: The master/slave relationship has been agreed upon, and nodes send DBD messages summarising their LSDBs.

Loading: Nodes compare the full description of their neighbour’s LSDB with their local data. Missing information is requested via LSR messages.

Full: All requested LSAs have been received.

A node’s neighbour state determines which messages it will send to its neighbour, as well as how it responds to messages received.

To validate the correctness of our model, we run some sanity checks. Among others we check that adjacencies are always completely established, i.e. both nodes reach the state Full, and that LSDBs are identical when the adjacency is established. Using UPPAAL’s CTL syntax, these properties correspond to

$$A\langle\rangle(n1.NeighbourState == FULL \ \&\& \ n2.NeighbourState == FULL) \ \text{and} \\ A[]((n1.NeighbourState == FULL \ \&\& \ n2.NeighbourState == FULL) \ \text{imply} \ LSDB_Sync())$$

Here $n1$ and $n2$ are the two nodes forming an adjacency, and $LSDB_Sync()$ is a function that checks LSDBs for synchronicity. The CTL formula $A\langle\rangle\varphi$ is satisfied if φ holds on some state along all paths, and $A[]\varphi$ is satisfied if φ holds on all states along all paths (e.g. [11]). We are able to verify these properties for a number of different starting configurations of the node’s LSDBs.

6 Modelling Adversaries

Over the last decades, model checking has become one of the standard tools for analysing and verifying systems, with the main focus on functional correctness, i.e. checking whether a system acts as expected.

Although correctness proofs are important, they are only half of the story. With systems connected more than ever, detecting and analysing vulnerabilities of systems are of equal importance. In the general area of security protocols, including cryptographic protocols, modelling adversaries and finding vulnerabilities using model checking techniques are common [20, 4]. With very few exceptions, e.g. [15], finding vulnerabilities in systems, such as routing protocols, using model checkers is uncommon.

In this section, we present first attempts to discover vulnerabilities, using model checking. When modelling routing attacks we distinguish two different concepts: *attack capabilities* and *attack goals*. Attack capabilities define the behaviour of the adversary. They include properties such as the number of malicious nodes and connectivities (e.g. [23]). In contrast, an attack goal describes the specific aim of an attack. This includes general properties such as delivery failure and routing loops, as well as protocol-specific targets such as the destruction of adjacencies. The consequences of (successful) attacks are sometimes also analysed. For example, a routing-loop attack yields packet loss.

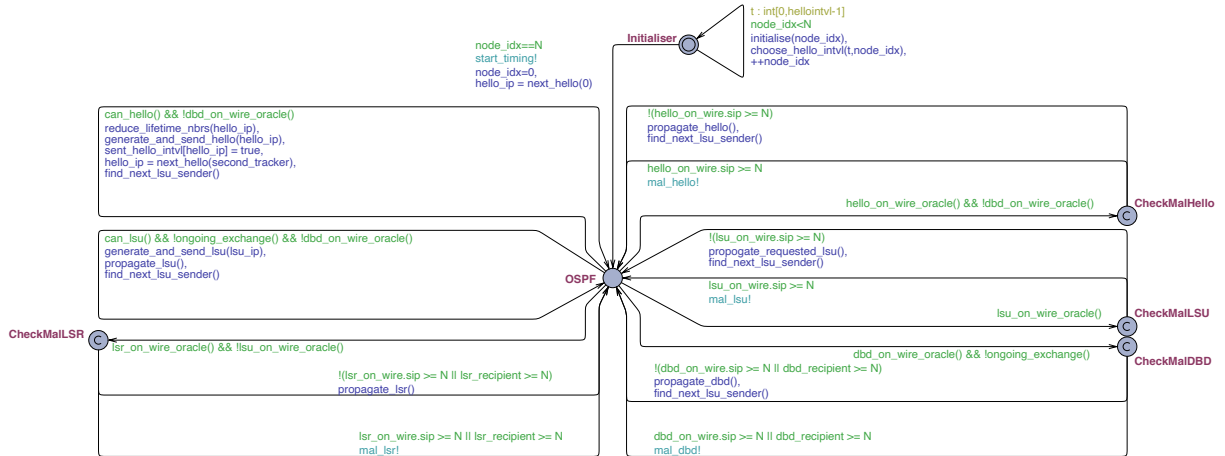
6.1 Attack Capabilities

We model capabilities as one or more separate timed automata. This modular approach allows us to combine the same attacker model with different formalisations of OSPF, or even other routing protocols.

Number of Malicious Nodes The *number of malicious nodes* in the network is one of the capabilities. In our setting, each malicious node (attacker) is modelled by a separate automaton executed in parallel to the model of the routing protocol. For the moment, we restrict ourselves to one malicious node only.

Connectivity There exist two reasonable scenarios: (a) a malicious node can only send to its immediate neighbours, similar to a node running an uncorrupted version of OSPF, and (b) a node can inject messages to all nodes in the network, which can be achieved in real networks (e.g. [23]). In both models of OSPF, we characterise connectivity as a Boolean matrix, or equivalently with a Boolean predicate `connect`, see Section 3. When modelling capability (a), no change is required. For option (b), a minor modification of the predicate `connect` suffices. Technically, the relation is not symmetric any longer.

Time of the Attack The third capability is the *time point an attack begins*. It determines at what point during the lifetime of the network a malicious actor performs the take over of a router and begins an attack. Although this point in time is arbitrary, we only consider two scenarios: (a) the node is already

Figure 6: A single automaton for OSPF, modified for attacks³

malicious when the protocol is initiated, and (b) the attacker gains access to a node after OSPF reaches a steady state, i.e. all adjacencies are established and all nodes have the same view of the network topology.

Other possible time points, e.g. the attack starts during initialisation of the protocol, are not considered since the likelihood of exploitation during these time periods is basically zero.

Adversaries' Activities The final, and arguably most important, capability describes the *activities* an adversary can perform. They are manifold, sometimes protocol-specific, and include the following:

- type of OSPF messages (DBD, LSR, LSU, ...) the attacker is allowed to send;
- components of the message which can be manipulated; e.g. the attacker can 'only' change the content of a message or it can pretend to send the message on behalf of some other node (false identity);
- the order of messages sent; ...

For our initial experiments, described below, we use an 'all mighty' adversary that can inject any type of message with arbitrary content at any time. We model this behaviour as a nondeterministic automaton.

Such an adversary is trivial to run in parallel with the original model of OSPF, featuring individual automata for each router (Figure 2), since concrete messages exist as first-class citizens. An 'all mighty' attacker in our optimised model requires some minor modifications as message contents are abstracted away. Figure 6 demonstrates the modified automaton; each transition for delivering a message is split into two branches. One is used for the exchange of standard OSPF messages, the other allows for the exchange of messages involving a malicious party.

It is noteworthy that the new, modified automaton decouples the OSPF protocol logic from the attacker logic and allows the same automaton of OSPF to be used for any number of different attack models without modification. It also allows the possibility of executing standard OSPF without any malicious nodes – in this case, our new automaton behaves identical to the original model.

6.2 Attack Goals

We model an attack goal as a state in the system. The reachability of this state determines if the attack is successful. Of course, it is not required that the attack is successful under all possible scenarios.

The formula $lsdp[sip][oip]$ describes sip 's current view of the network topology around node oip , i.e. it lists, among other information, all neighbours of oip that are known to sip .

Blackhole Attack In this type of attack, the attacker tries to convince one or more nodes to send (data) packets via the attacker so that the packets can be dropped. Any packet routed through the malicious node will suffer from partial or total data loss. These routes can only be established if either nonexistent links are installed in a node’s LSDB or if some links are removed. Both existence and removal is characterised by a malicious/fake topology:

$$E \leftrightarrow \left(\bigwedge_{\substack{\text{sip} \in S \\ \text{dip} \in T}} \text{lsdb}[\text{sip}][\text{dip}].\text{malicious} \right) \quad (1)$$

for some sets S and T . In case one does not want to encode non-optimal routes via concrete LSDB entries, one can calculate the shortest routes from a given LSDB.

Sub-optimal Routes Establishing sub-optimal routes means that the number of hops of the established route is greater than the actual shortest path. This attack does not cause packet loss but wastes network resources. As packet routing is based on shortest paths, we can use the same formulas (Eq. (1)) to encode this attack.

Adjacency Interruption The first two attacks are concerned with OSPF. An attacker can also inject or spoof OSPF messages that prevent nodes which should form an adjacency from completing the synchronisation process

7 Initial Experiments

Using our optimised model, we analyse the above attack goals on very small topologies. These experiments are not intended to be a systematic analysis of vulnerabilities, but are merely intended to show the power of our approach.

7.1 Blackhole Attack and Sub-optimal Routes

We report on a number of examples of networks up to size 4, a full, systematic analysis on all topologies up to a much larger size is part of future work. Some topologies used for the experiments are depicted in Figure 7. The solid lines indicate the real topology, the dashed lines stand for the malicious topology. In all cases, the malicious router is node 3, and its (main) victim is node 0.

Attacks whose effects get subsequently corrected by the mechanisms of the protocol, in our case OSPF, are called *nonpersistent*. They can, nevertheless, temporarily affect the way the traffic is routed and can slow down the network traffic.

The attacks we discover rely on placing malicious entries in the victims’ LSDBs. Each experiment related to an attack goal is run in two instances (cf. Section 6.1): (a) the attack starts before the network is up and running, i.e. all nodes in the network start with an empty LSDB and the OSPF runs through its initialisation, and (b) the attack starts after OSPF has been initialised, i.e. all LSDBs of all honest nodes contain all truthful information about the topology. In both cases, the attacker starts with an LSDB having malicious version of the topology as content. This LSDB will not change during an attack. The adversary then tries to propagate this false information through the network.

In the blackhole attack, the adversary convinces its victim to redirect the traffic through the attacker node. In the simple topology of Figure 7(a) an adversary can be successful, i.e. the property $E \leftrightarrow (\text{lsdb}[0][2].\text{malicious} \ \&\& \ \text{lsdb}[0][3].\text{malicious})$ holds. This is because most of the nodes are one-hop neighbours and proper calculation of shortest paths are not needed. In this situation the attack is



Figure 7: (a) Blackhole

(b) Sub-optimal routes

successful if node 0 believes that there is a connection between nodes 2 and 3 (and no connection between 1 and 2). Instead of sending the traffic through node 1, node 0 will then forward the traffic for node 2 through node 3, which will drop it.

Sub-optimal routes can be installed in many topologies, such as the one depicted in Figure 7(b). As for the previous attack we can simplify the attack goal and use the formula $E \langle \rangle (\text{lsdb}[0][1].\text{malicious} \ \&\& \ \text{lsdb}[0][3].\text{malicious})$. The attacker convinces node 0 that there is a connection between nodes 1 and 3. Node 0 then sends traffic for node 3 through node 1 (instead of sending it through node 2). Node 1 correctly believes that there is no link between itself and node 3, so it passes traffic through node 2.

The network easily recovers from these attacks by receiving further HELLO messages, or by receiving further LSU messages sent to correct faulty information. Remember that these simple attacks are designed to test our setup, not to do a full-scale analysis of vulnerabilities.

7.2 Tear Down Adjacencies

Similarly to our use of the abstract model of OSPF, we perform initial experiments with the detailed model for adjacency building. As before we use a nondeterministic automaton, which can inject arbitrary messages. We also use more limiting attack capabilities that can only inject a certain type of OSPF message. The one for HELLO messages is depicted in Figure 8.

Our analysis reveals that adjacencies can be torn down. We sketch this attack; here $n1$ and $n2$ are the two routers involved.

- Sending a HELLO message on behalf of $n1$ ($n2$) which does not carry the IP address of $n2$ ($n1$) in its payload resets the adjacency building process by reverting the neighbour state back to `Init`.
- Sending an LSR message on behalf of $n1$ or $n2$ which requests an LSA that does not exist in the recipient's LSDB reduces the neighbour state to `ExStart` if it is in a higher state at the time of receiving. In case the neighbour state is `Down`, `Init` or `ExStart` this injection does not have any effect.
- Sending a DBD message on behalf of $n1$ or $n2$ which meets one of the following criteria will regress to the neighbour state to `ExStart` if it is in a higher state at the time of receiving:
 - The DBD sequence number, a unique number which is part of every DBD message, does not match the expected sequence number.
 - The DBD's master flag is inconsistent with the current connection.
 - The message has the `init` flag set – this flag indicates that this DBD is an attempt to *initiate* a neighbour relationship.

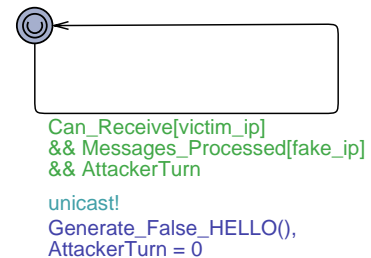


Figure 8: Injecting HELLOs

Once a node's neighbour state has been reverted to `Init` or `ExStart`, it is no longer able to share information with its neighbour. As a consequence, the adjacency needs to be re-established. One of these attacks can delay the establishment of adjacencies; repeating one of them indefinitely avoids the establishment of

adjacencies entirely. When checking $A \langle \rangle (n1.NeighbourState == FULL \ \&\& \ n2.NeighbourState == FULL)$ UPPAAL generates a trace that falsifies the property, characterising a specific attack.

The above trace demonstrates that adjacency building can be interrupted, but it does not show that it is possible to tear down full adjacencies. In order to demonstrate this, we make very minor adaptations to our model which enable the nodes to start in state Full with their LSDBs synchronised. Using the same attacker automata, we use UPPAAL to show that it is possible to regress our node's neighbour state to Init or ExStart. Since UPPAAL does not allow for nested CTL expressions, it is not possible to verify that the nodes are unable to re-establish a state Full. However, once the nodes have regressed to Init or ExStart, we are able to apply our earlier findings, as this is the same configuration as in our first experiment.

8 Related Work

Some of the attacks presented in Section 7.2 are already mentioned in [17]. That analysis is based on a manual inspection of the specification of OSPF. While it is an impressive piece of work, our analysis is more systematic and reveals more vulnerabilities w.r.t. adjacency building.

As for OSPF in general, we are aware of only three other formal approaches that analyse the OSPF routing protocol, using formal methods. We believe that our analysis is built on the most detailed, open-source formal model of OSPF [10].

Nakibly et al. [24] have created their own model of OSPF for the model checker CBMC [6], a bounded model checker accepting a simplified C-language as input. Their model uses a fixed topology — similar to our model-checking models — and abstracts away from many details, including HELLO and DBD messages. The model that is available online analyses a topology with three nodes, but not more. Moreover, adjacency building is not modelled. With their level of abstraction, it is not possible to model and find some of the attacks described in this paper.

Another model, which can be fed into the Z3-Solver [8], is described in [19]. The authors claim that it is a detailed model covering concepts such as designated routers —a concept we abstract from— and that topologies up to 30 nodes can be analysed. As the model is not presented in full nor available online, we could neither verify these claims nor compare our analysis to theirs.

9 Conclusion and Future Work

We have presented two advanced models of OSPF: the first is an optimised model of an existing UPPAAL-model; the second models adjacency building, a subprocedure of OSPF, in great detail. Furthermore, we have developed several attack automata that can run in parallel to the models of OSPF. This means our method is compositional: both the model and the attacker can be exchanged to variant versions.

While this paper concentrates on our modelling efforts, future work will aim at a vulnerability analysis of OSPF at scale. We hope to not only rediscover well-known attacks such as the remote false adjacency attack (a.k.a. phantom router) [23], which allows a malicious router to successfully establish a *permanent* link to a victim, but also to discover new attacks. We envision that we can transfer the knowledge acquired in UPPAAL to (automatically) generate attacks for a real implementation. In particular we believe that traces produced by UPPAAL can be used to generate attacks in test-bed implementations of OSPF (e.g. the Quagga implementation running on the network emulator CORE [1]) or in real networks.

Acknowledgements: This work was conducted in partnership with the Defence Science & Technology Group and Data61, CSIRO, through the *Next Generation Technologies Fund*. We thank Jack Drury for fruitful discussions.

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] Karthikeyan Bhargavan, Davor Obradovic & Carl A. Gunter (2002): *Formal Verification of Standards for Distance Vector Routing Protocols*. *J. ACM* 49(4), pp. 538–576.
- [4] B. Blanchet (2016): *Modeling and Verifying Security Protocols with the Applied Pi Calculus and ProVerif*. *Foundations and Trends in Privacy and Security* 1(1-2), pp. 1–135, doi:10.1561/33000000004.
- [5] 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.
- [6] 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.
- [7] R. Coltun, D. Ferguson, J. Moy & A. Lindem (2008): *OSPF for IPv6*. RFC 5340, Network Working Group. Available at <http://www.ietf.org/rfc/rfc5340.txt>.
- [8] 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.
- [9] E. W. Dijkstra (1959): *A Note on Two Problems in Connexion with Graphs*. *Numerische Mathematik* 1(1), pp. 269–271, doi:10.1007/BF01386390.
- [10] J. Drury, P. Höfner & W. Wang (2020): *Formal Models of the OSPF Routing Protocol*. In A. Fehnker & H. Garavel, editors: *Models for Formal Analysis of Real Systems (MARS'20), Electronic Proceedings in Theoretical Computer Science* 316, Open Publishing Association, pp. 72–120, doi:10.4204/EPTCS.316.4.
- [11] E. A. Emerson (1990): *Temporal and Modal Logic*. In: *Handbook of Theoretical Computer Science (vol. B): Formal Models and Semantics*, MIT Press, pp. 995–1072.
- [12] 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.
- [13] A. Fehnker, L. van Hoesel & A. Mader (2007): *Modelling and Verification of the LMAC Protocol for Wireless Sensor Networks*. In: *Integrated Formal Methods, IFM 2007, Lecture Notes in Computer Science* 4591, Springer.
- [14] 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.
- [15] K. Havelund, K. G. Larsen & A. Skou (1999): *Formal Verification of a Power Controller Using the Real-Time Model Checker UPPAAL*. In J.-P. Katoen, editor: *Formal Methods for Real-Time and Probabilistic Systems (ARTS'99), Lecture Notes in Computer Science* 1601, Springer, pp. 277–298, doi:10.1007/3-540-48778-6_17.
- [16] 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>.

- [17] E. Jones & O. Le Moigne (2006): *OSPF Security Vulnerabilities Analysis*. Routing Protocol Security. Available at <https://tools.ietf.org/id/draft-ietf-rpsec-ospf-vuln-02.txt>.
- [18] 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.
- [19] S.U.R Malik, S. K. Srinivasan & S. U. Khan (2012): *A Methodology for OSPF Routing Protocol Verification*.
- [20] S. Meier, B. Schmidt, C. Cremers & D. Basin (2013): *The TAMARIN Prover for the Symbolic Analysis of Security Protocols*. In N. Sharygina & H. Veith, editors: *Computer Aided Verification (CAV'13)*, *Lecture Notes in Computer Science* 8044, Springer, doi:10.1007/978-3-642-39799-8_48.
- [21] 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.
- [22] J. Moy (1998): *OSPF Version 2*. RFC 2328, Network Working Group. Available at <http://www.ietf.org/rfc/rfc2328.txt>.
- [23] G. Nakibly, D. Gonikman, A. Kirshon & D. Boneh (2012): *Persistent OSPF Attacks*.
- [24] 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.
- [25] R. de Renesse & A.H. Aghvami (2004): *Formal Verification of Ad Hoc Routing Protocols Using SPIN Model Checker*. In: *Proceedings of IEEE MELECON'04*, IEEE, pp. 1177 – 1182.
- [26] O. Wibling, J. Parrow & A. N. Pears (2004): *Automatized Verification of Ad Hoc Routing Protocols*. In: *Formal Techniques for Networked and Distributed Systems – FORTE*, *Lecture Notes in Computer Science* 3235, Springer, pp. 343–358.