

2017-04-29

Modelling and Verification of a Cluster-tree Formation Protocol Implementation for the IEEE 802.15.4 TSCH MAC Operation Mode

Mahmoud Talebi

Department of Mathematics and Computer Science Eindhoven University of Technology, The Netherlands

Jan Friso Groote

Department of Mathematics and Computer Science Eindhoven University of Technology, The Netherlands

Conrad Dandelski

*Nimbus Centre for Embedded Systems Research Cork Institute of Technology, Ireland,
conrad.dandelski@mycit.ie*

Follow this and additional works at: <https://sword.cit.ie/nimbusart>



Part of the [Computer Sciences Commons](#)

Recommended Citation

Talebi, M, Groote, JF & Dandelski, C 2017, Modelling and verification of a cluster-tree formation protocol implementation for the IEEE 802.15.4 TSCH MAC operation mode. in H Hermanns & PH Höfner (eds), Proceedings 2nd Workshop on Models for Formal Analysis of Real Systems (MARS 2017). pp. 117-128, 2nd Workshop on Models for Formal Analysis of Real Systems (MARS 2017), Uppsala, Sweden, 29/04/17. <https://doi.org/10.4204/EPTCS.244.5>

This Article is brought to you for free and open access by the Nimbus Centre at SWORD - South West Open Research Deposit. It has been accepted for inclusion in NIMBUS Articles by an authorized administrator of SWORD - South West Open Research Deposit. For more information, please contact sword@cit.ie.

Modelling and Verification of a Cluster-tree Formation Protocol Implementation for the IEEE 802.15.4 TSCH MAC Operation Mode

Mahmoud Talebi

Jan Friso Groote

Conrad Dandelski

*Department of Mathematics and Computer Science
Eindhoven University of Technology, The Netherlands*

Email: {M.Talebi, J.F.Groote}@tue.nl

*Nimbus Centre for Embedded Systems Research
Cork Institute of Technology, Ireland*

Email: conrad.dandelski@mycit.ie

Correct and efficient initialization of wireless sensor networks can be challenging in the face of many uncertainties present in ad hoc wireless networks. In this paper we examine an implementation for the formation of a cluster-tree topology in a network which operates on top of the TSCH MAC operation mode of the IEEE 802.15.4 standard, and investigate it using formal methods. We show how both the mCRL2 language and toolset help us in identifying scenarios where the implementation does not form a proper topology. More importantly, our analysis leads to the conclusion that the cluster-tree formation algorithm has a super linear time complexity. So, it does not scale to large networks.

1 Introduction

The IEEE 802.15.4 protocol standard is used for low-rate and low-powered Personal Area Networks (PANs). The IEEE 802.15.4e standard was released in 2012 to extend the IEEE 802.15.4 protocol to overcome limitations in performance and reliability [2, 5]. This extension was subsequently incorporated into the IEEE 802.15.4 standard in 2015, and among other changes introduces the Time-Slotted Channel Hopping (TSCH) MAC operation mode, with the goal of ensuring reliability by providing more structure to wireless communication, as well as scalability and robustness by channel hopping [3].

In the TSCH mode, robustness against noise and interference is provided by a mechanism for channel hopping in consecutive units of time called time-slots, and the reliability is provided by communicating according to a deterministic schedule over time-slots. The schedule is created by higher-layer protocols on the PAN coordinator, which is also in charge of initiating and maintaining the PAN.

For multi-hop applications, a logical structure is maintained which is called a topology, and is often a cluster-tree. This structure is used for multi-hop routing and often also for creating schedules. However, as networks become larger and/or denser the problem of designing an algorithm for forming the cluster-tree efficiently becomes non-trivial, and ensuring its correctness becomes difficult.

A problem is that the IEEE 802.15.4 protocol standard is not explicit about many of the operational details of the protocol, including the cluster-tree formation algorithm. Therefore, we used the description of an early implementation [4] and investigated its correctness using the mCRL2 toolset [1, 10]. Our findings were in line with results of simulations by the designers, with substantial saving of time in case of verification compared to simulations. Some, but not all, observations were incorporated in later implementations.

In order to investigate the implementation, we first model the lower layers (including the IEEE 802.15.4 TSCH mode protocol) in mCRL2. This model is presented in Section 3.2. We then model the cluster-tree formation protocol operating on TSCH. This model is quite large and can be found in the model repository. We formally analyse correctness using the mCRL2 model checker. We subsequently present a full argument on the scalability of the algorithm, based on the results derived by verification.

From our analysis we derived a number of guidelines on how to design a better cluster-tree formation protocol (in Section 6). But we do not devise such an alternative protocol, as the purpose of this paper is to show that with relatively little effort formal methods can be of great help in getting insight in the behaviour of protocols used in the realm of wireless networks.

2 Related work

When discussing cluster-trees, we are essentially referring to spanning trees, in which each node is a vertex in the tree, and for every pair of node identifiers (A,B) there is a unique path between the node with identifier A and the node with identifier B in the tree. The problem of forming a cluster-tree is directly related to the distributed minimum spanning tree (MST) algorithms in an asynchronous setting [13]. However, several assumptions by distributed MST algorithms are in conflict with the situation in wireless sensor networks: links between nodes should be reliable meaning that no collisions or message losses occur, and processes on nodes should have complete knowledge regarding the neighbours of the node. Moreover, in some versions of the distributed spanning tree algorithm the leader or the root node is elected during the course of execution [7, 8], while in the algorithms we consider the root node is fixed.

The general strategy in distributed spanning tree algorithms is to start with a spanning forest, and consistently merge trees until eventually a single spanning tree is formed [13]. However, in IEEE 802.15.4 due to having a single initiator an essentially different approach is taken which forms a tree by a series of waves (i.e., floodings) sent through the network. Therefore, the algorithm presented here fits the definition of echo algorithms in the theory of distributed algorithms [6].

To our knowledge there have been relatively few attempts to rigorously verify the IEEE 802.15.4 protocol up until now. Recently, in [12] the authors used the UPPAAL framework to analyse inconsistencies in slot allocation in an implementation of the distributed synchronous multi-channel extension (DSME) mode of the IEEE 802.15.4 MAC layer protocol. Earlier, in [11] the authors modelled a part of the IEEE 802.15.4 specification using the MODEST language to investigate the influence of clock precision on energy consumption. In any case, this work is one of the first to formally analyse parts of the TSCH operation mode.

3 Modelling the TSCH MAC operation mode

In this section we introduce the service of the TSCH MAC layer as specified in [3]. We also sketch how we model protocols using this MAC protocol. We leave out the details of the channel hopping procedure from the TSCH MAC layer. The channel hopping procedure is an essential part of TSCH, which provides scalability and reliability in communication. However, as the correctness of channel hopping procedures is not a concern in this paper we abstract away all related concepts.

In what follows we first give an informal specification for the TSCH MAC operation mode in Section 3.1. Then, in Section 3.2 we show how the informal descriptions can be encoded formally.

3.1 The TSCH MAC operation mode

In IEEE 802.15.4, 16 channels are available for communication. The TSCH MAC operation mode is time-slotted, and divides time into equally sized chunks called *time-slots*. For each time-slot on each channel a node can be either a receiver, a transmitter, or not involved (e.g., if it is sleeping).

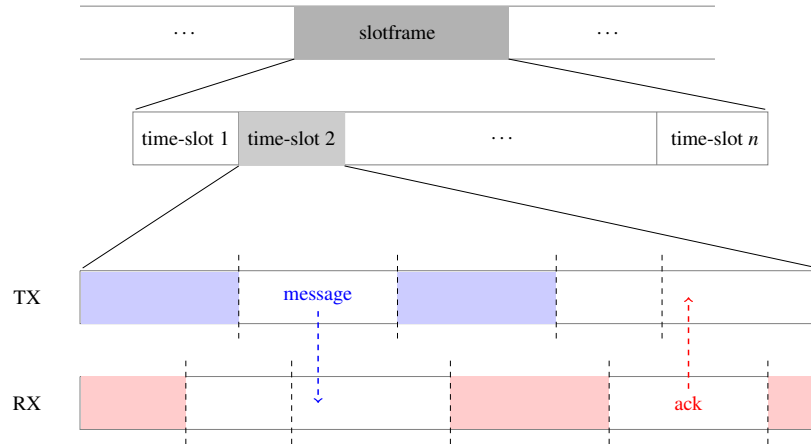


Figure 1: A slotframe with n time-slots and communication between a transmitter and a receiver.

In TSCH, time-slots have two types: they are either *dedicated*, meaning that only one transmitter and one or more receivers are involved, or they are *shared*, meaning that there can be zero, one or more transmitters attempting to transmit a frame and multiple receivers listening. In case of shared time-slots conflicts between transmitters are resolved using CSMA-CA.

For communicating over time-slots, nodes must share a common schedule, which is designed such that their frames do not collide. Otherwise, collisions are a possibility. Let us assume that in a time-slot node A is the transmitter (TX) and node B is the receiver (RX), according to the schedule. The events occurring in the time-slot are as follows: at the beginning of the time-slot both A and B wait for a time period. After waiting, first B starts listening on the channel, then A becomes active and transmits its message. When A is done transmitting the message, it stops and waits for some time. Next, B stops listening and waits. After some time node A starts listening on the channel. Then B becomes active and if it has succeeded in receiving the frame, it transmits an acknowledgement. If not, it does no action. After a while (before the start of the next time-slot) A stops listening. In this way by the end of the time-slot, A and B reliably know the outcome of the communication. Also note that listening starts before the expected transmission time and ends sometime after to allow possible synchronisation imperfections.

In TSCH, nodes have a common understanding of time, and agree on schedules. The schedules are repeated over slotframes, which are a collection of time-slots. A summary of these concepts and their relations is given in Figure 1.

3.2 The two layered model

In this section we outline the formal specification of the general behaviour of the network as a two layered model, with one representing the physical and the TSCH operation mode of the MAC layer of the protocol stack, and the other representing the upper layers. At the same time, we also introduce the mCRL2 formal specification language for specifying and verifying the behaviour of concurrent systems [1].

For the specification of the behaviour the following data types are needed. Let *Message_Type* be a sort which at least contains the element `EMPTY`. The inclusion of the element `EMPTY` is due to the fact that in our approach in modelling broadcasting, in each time-slot we require every node to submit a message to the physical layer. Define:

```
sort  Message = struct message(src :  $\mathbb{N}$ , channel :  $\mathbb{N}$ , type : Message_Type, data : List( $\mathbb{N}$ ));
      Traffic  = List(Message);
```

The above **sort** statement defines custom data types. The sort *Message* stands for a structured sort, which can be constructed using the function *message* applied to the elements that follow. The element *src* in *message* is the identifier of the transmitter of the message which belongs to the set of natural numbers, *channel* is the channel over which the message is transmitted, again belonging to the set of natural numbers, *type* is the type of the message, and *data* is a list of data carried by the message (payload), which belongs to the set of lists (sequences) of natural numbers.

The sort *Traffic* is a list of messages. This models all the messages coming from all the sources travelling over all the channels at a certain point in time in the network.

Constants and functions are defined using the **map** statement:

```
map  MAX_ID :  $\mathbb{N}$ ;
      NUM_OF_CHANNELS :  $\mathbb{N}$ ;
```

The above statement defines two constants with values from the set of natural numbers, as mappings with zero arguments. We use the constant *MAX_ID* to indicate that $1, \dots, \text{MAX_ID}$ are identifiers of processes in the network, and the constant *NUM_OF_CHANNELS* indicates that the channels available are $1, \dots, \text{NUM_OF_CHANNELS}$.

Based on the informal specification given in Section 3.1, the behaviour of the network depends on the type of time-slots (dedicated or shared). Here, we only define the behaviour for dedicated time-slots.

We model the physical and MAC layers of the network in a general fashion, such that it applies to any protocol implemented over the dedicated time-slots of TSCH. This is done using the process *TSCH*. The main purpose of defining this process is to model the peculiarities of phenomena such as *broadcasting* and *collisions* using mCRL2.

We model the transmission and receiving of messages using the following **act** statement:

```
act  corecv, recv, received : Message;
      cosend, send, sent    : Message;
```

Actions represent events occurring in the system, and are the building blocks in defining the behaviour of processes. The action *received* indicates that a message has been received in the MAC layer and is handed over to the upper layer. Concretely, the MAC layer performs a *corecv* action and the upper layer performs a *recv* which synchronise into the *received* action. Similarly, the action *sent* means that the upper layer indicates that the MAC layer must send a message. This action is the result of the synchronisation of actions *send* by the upper layer and *cosend* by the MAC layer.

The synchronisations are made explicit by a communication set *C*:

$$C = \{recv|corecv \rightarrow received, send|cosend \rightarrow sent\}.$$

A synchronisation can only take place when both constituent actions are ready and have the same data of type message. The communication operator Γ_C is used in the process to indicate that these communications are possible. The operator ∇_V only allows actions in the set *V* to happen. It can be used to enforce communications by only allowing the result of a synchronisation and not the constituent actions. Concretely, when applying $\nabla_{\{sent, received\}}$ the result is that actions *cosend* and *send* cannot be performed independently and must communicate to *sent*. The same applies to *corecv*, *recv* and *receive*.

Assuming that the process *TSCH* describes the physical and MAC layers and the processes *Node* describe the upper layer, a TSCH network with $n = \text{MAX_ID}$ nodes is specified by the process *Network*:

```
proc Network =  $\nabla_{\{sent, received\}} \Gamma_C(TSCH \parallel Node(1, c_1, d_1) \parallel \dots \parallel Node(n, c_n, d_n));$ 
```

In the following two sections, we give the details on how processes *TSCH* and *Node* are specified.

3.2.1 Specification of the *TSCH* process

In this section, we specify the behaviour of the physical and the MAC layers using the mCRL2 language, as the process *TSCH*. We start out with the *TSCH* process.

We first define the following functions which enable us to manipulate the traffic in the network:

map $findId : \mathbb{N}\#Traffic \rightarrow Message;$
 $findChannel : \mathbb{N}\#Traffic \rightarrow Traffic;$

The function $findId$ accepts a node identifier and a list of messages (*Traffic*), and returns the first message from a node with the mentioned identifier in this list, only if such a message exists. The function $findChannel$ accepts a channel number and a list of messages, and returns a list of messages sent on a certain channel. We specify the behaviour of *TSCH* using two processes $TSCH_1$ and $TSCH_2$.

The process $TSCH_1$ represents the first stage of the events on the physical layer. In this stage, $TSCH_1$ collects from each node i a message to be sent (using the action *cosend*) and adds it to the sent transmissions str such that it will be transmitted in the second stage.

proc $TSCH_1(i : \mathbb{N}, str : Traffic) =$
 $(i \leq MAX_ID)$
 $\rightarrow \sum_{c:\mathbb{N}, d:List(\mathbb{N})} cosend(message(i, c, d)) \cdot TSCH_1(i + 1, str \triangleleft message(i, c, d))$
 $\diamond TSCH_2(1, str, removeCollision(1, str));$

The notation $c \rightarrow p \diamond q$ stands for the if-then-else. The operator \triangleleft is used to put an element at the end of a list. Note that every node should send a message. Nodes which are only listening, communicate EMPTY messages to the *TSCH* in this stage.

While transmitting messages, collisions may occur on the physical layer. A collision occurs over a channel c if there are at least two messages m_1 and m_2 which are on the same channel c . In case of a collision on channel c , none of the messages sent over channel c are readable and therefore they need to be removed when the sent transmissions become the received transmissions in *TSCH*. This behaviour is described using a function *removeCollision*:

map $removeCollision : \mathbb{N}\#Traffic \rightarrow Traffic;$
var $tr_1 : Traffic; chann : \mathbb{N};$
eqn $removeCollision(chann, tr_1) =$
 $if(chann \leq NUM_OF_CHANNELS,$
 $if(\#(findChannel(chann, tr_1)) \approx 1,$
 $removeCollision(chann + 1, tr_1) \triangleleft head(findChannel(chann, tr_1)),$
 $removeCollision(chann + 1, tr_1) \triangleleft message(0, chann, EMPTY, 0)),$
 $[])$;

The function *removeCollision* processes all the messages in the traffic over all channels, and for each channel determines which message is broadcast. The message is then added to the list of results. In case of collision on a channel c , an EMPTY message over c is added to the resulting list.

In the second stage, the process $TSCH_2$ checks the channel each node i is operating on using the expression $channel(findId(i, str))$, and picks a message from that channel in the list rtr and delivers it to the node.

proc $TSCH_2(i : \mathbb{N}, str : Traffic, rtr : Traffic) =$
 $corecv(head(findChannel(channel(findId(i, str)), rtr))) \cdot$
 $((i < MAX_ID) \rightarrow TSCH_2(i + 1, str, rtr) \diamond TSCH_1(1, []));$

Upon finishing this stage, the behaviour changes back to process $TSCH_1$. The process *TSCH* which represents the complete behaviour of the lower layers is then defined as:

proc $TSCH = TSCH_1(1, []);$

3.2.2 General specification of process *Node*

The general structure of the processes *Node* is as follows:

$$\begin{aligned}
\text{proc } Node(id : \mathbb{N}, ch : \mathbb{N}, d : Data) = & \\
& \sum_{i \in I} \gamma_i(d) \rightarrow \left(send(message(id, ch, t_i, g_i(id, d))) \cdot \right. \\
& \quad \left. \sum_{id' : \mathbb{N}, t' : Message_Type, d' : Data} (id' \leq MAX_ID) \rightarrow \right. \\
& \quad \left. recv(message(id', ch, t', d')) \cdot Node(id, \sigma_i(ch, d), \delta_i(ch, d)) \right) \\
+ & \\
& \sum_{j \in J} \gamma_j(d) \rightarrow \left(send(message(id, ch, EMPTY, [])) \cdot \right. \\
& \quad \left. \sum_{id' : \mathbb{N}, t' : Message_Type, d' : Data} (id' \leq MAX_ID) \rightarrow \right. \\
& \quad \left. recv(message(id', ch, t', d')) \cdot Node(id, \sigma'_j(ch, d, id', d'), \delta'_j(ch, d, id', d')) \right)
\end{aligned}$$

In the process *Node*, the number *id* is the identifier of the node. The parameter *ch* is the channel the node operates on. The parameter *d* of sort *Data* is a placeholder for the rest of the state of a node. In the concrete description of *Node* *d* is replaced by 7 concrete parameters.

Each process *Node* represents a wireless node with a single antenna which at any time can only be in a transmit mode or a receive mode, based on which the behaviour of a node is split into two main parts, represented by the two summations.

The set *I* is a finite set of indices *i* of all conditions γ_i under which the custom-defined protocol transmits a message. Similarly, the set *J* is the set of indices *j* of all conditions γ_j under which the custom-defined protocol waits for a message. The first summation (choosing over set *I*) represents the behaviour of nodes in the transmit mode. In this mode, if condition γ_i with data *d* holds, the node sends a message of type t_i over *ch*, containing data $g_i(id, d)$. The antenna is passively exposed to any signal occurring over channel *ch* and this is reflected by the subsequent *recv* action, which records the messages broadcast at the same time over the operating channel. However, being in transmit mode, the node cannot react to this message. Therefore the values of functions σ_i and δ_i which determine the next state of the node are not affected by the received data.

The second summation (choosing over set *J*) shows the behaviour of nodes in the receive mode. In this mode, a node does not transmit any messages. This particular behaviour is shown by a *send* action of a message of type *EMPTY* over channel *ch*. The *send* action is only to inform the *TSCH* process of the value of *ch*. The main part of the behaviour is then expressed by the *recv* action. Functions σ'_j and δ'_j express the change to the state of the node based on the received data and the previous state.

Note that despite the fact that two communications occur in each time-slot in TSCH, namely transmission of a frame and its acknowledgement, we do not explicitly enforce them in our model. Therefore the user defined protocol may choose to include the acknowledgement in its behaviour by specifying two rounds of communication per time-slot, or to leave the acknowledgements out.

4 The cluster-tree formation protocol

In this section, we specify the cluster-tree formation protocol using mCRL2, and then we present the results of the verification of this protocol which shows the flaws in its design.

4.1 Informal specification

A slotframe in the cluster-tree formation protocol consists of 12 time-slots, out of which two are reserved for the operation of the cluster-tree formation protocol. The idea is that the protocol establishes and maintains a cluster-tree topology in these two time-slots while the network is doing its normal operation in the rest of the time-slots.

Nodes can have one of 4 types. They can be either a cluster head (CH), a cluster slave (CS), a tentative cluster head (TCH), or a free node (FN). In the scenarios that we are interested in, the network starts with one cluster head node and the rest are free nodes. The original cluster head is a gateway node receiving commands from an external source, and it initiates the cluster-tree formation protocol.

The cluster-tree formation protocol takes a distance based approach. For each node, nodes in its broadcast range fall into two categories: close nodes, and other nodes in range. A node estimates the distance of another node based on its received signal strength (RSSI). Upon receipt of a message, a node measures the power of the signal, and compares it to a power threshold and determines the category of the transmitter.

Cluster heads in the network are connected to each other (they form a tree) in a setting in which every cluster head has one parent and a number of children. The original cluster head is an exception and has no parent. Based on this a cluster head is assigned a *tier*, which shows its distance from the original cluster head, with the original cluster head being assigned tier 0. In addition, each cluster head operates on only two channels: one to communicate with its parent cluster head, and one to communicate with its children. We call these channels the *parent channel* and the *assigned channel* respectively.

A cluster head broadcasts a BEACON message on its assigned channel to advertise its presence in the network, according to the scheme that follows. Cluster heads with even-numbered tiers advertise on the first time-slot reserved for cluster-tree formation in a slotframe, and cluster heads with odd-numbered tiers advertise on the second time-slot reserved for cluster-tree formation. Therefore every cluster head gets the chance to advertise once in each slotframe. The BEACON message includes the tier of the transmitter node.

The message is received by all the nodes which are within transmission range, and are listening on the assigned channel of the cluster head. When a free node receives the message, it estimates its distance from the transmitter, and if it is close then it sends an ASSOCIATE message in response and requests to become a cluster slave. A cluster slave operates on the same channel as its parent. Otherwise if the node is not close to the cluster head, it responds with a BEACON_ACK to the request, and the cluster head then sends an ACK_RESPONSE message which contains a `wait_time`.

When an ACK_RESPONSE message is received by a free node, it becomes a tentative cluster head. A tentative cluster head scans other channels one by one (staying in each for 2 time-slots) until either `wait_time` expires, or it receives a BEACON from a cluster head which is close and has the same tier. In such a case the node becomes a cluster slave of the close cluster head, as previously explained. This is to achieve approximately the smallest tree in the network which covers all of the nodes.

If a tentative cluster head fails to find a close cluster head and `wait_time` expires, it moves back to the channel of its parent and listens for a BEACON message. This time it sends an ASSOCIATE message in response and requests to become a child cluster head. The parent then forwards this message to the original cluster head. The original cluster head finds a suitable and possibly unoccupied channel, and sends an ASSOCIATE_ACK message containing the assigned channel to the sender of the request, which is then forwarded to the tentative node, officially granting it the status of a cluster head.

We assume that the network is well connected, meaning that for each two nodes A and B in the network it is possible to deliver a message from A to be B in one or more hops. The main requirements

that the above protocol should satisfy in such a network is the following:

- Eventually every node in the network becomes either a cluster head or a cluster slave.

Below we present our method to check this property.

4.2 Formal specification of the cluster-tree formation protocol in mCRL2

In this section we explain how the cluster-tree formation protocol introduced in Section 4.1 can be formally specified. Based on the informal specification, we define data types *Node_Type* and *Message_Type* as follows:

```
sort  Node_Type = struct CLUSTER_HEAD | TENTATIVE | CLUSTER_SLAVE | FREE;
      Message_Type = struct EMPTY | BEACON | BEACON_ACK | ACK_RESPONSE |
                    ASSOCIATE | ASSOCIATE_ACK;
```

Next, we model the physical topology of the network using the distance of nodes with respect to each other and by declaring the functions *inRange* and *inClose*.

```
map  inRange :  $\mathbb{N}\#\mathbb{N} \rightarrow \mathbb{B}$ ;
      inClose :  $\mathbb{N}\#\mathbb{N} \rightarrow \mathbb{B}$ ;
```

For two node identifiers *A* and *B* the value of the term *inClose(A, B)* is true if and only if node *A* is within range of node *B*, and is closer than the threshold. The term *inRange(A, B)* is true, if and only if node *A* is within range of node *B* but farther away than the threshold.

Using these data types, and the general definition of a TSCH node, the basic declaration of the process modelling a node is:

```
proc  Node(id :  $\mathbb{N}$ , ch :  $\mathbb{N}$ ,
          type : Node_Type, pid :  $\mathbb{N}$ , pc :  $\mathbb{N}$ , children : List(Pair), tier :  $\mathbb{N}$ , state :  $\mathbb{N}$ , timeslot :  $\mathbb{N}$ )
```

In this process, if a node is of type cluster head or cluster slave, *pid* is the identifier of the parent of the node and *pc* is the channel over which the parent of the node operates. Given type *Pair* which defines any pair of natural numbers, a cluster head keeps a record of all its children and their assigned channels as a list of pairs of natural numbers. The parameter *tier* is the tier of the cluster head or cluster slave and *state* is the internal state of the node which in general holds information such as the node's history and what the node should do next. Finally, *timeslot* is the information regarding the clock of the node, which is used by the node to keep its actions in sync with the rest of the network.

We model a TSCH network in which the original cluster head is initialised to the state:

$$Node(1, 1, CLUSTER_HEAD, 0, 0, [], 0, 0, 0),$$

and every other node with identifier $n \in \{2, \dots, MAX_ID\}$ is initialised to: $Node(n, c_n, FREE, 0, 0, [], 0, 0, 0)$, where c_n is a channel number, chosen based on a scenario of interest. The process *Network* given in Section 3.2 then models a network of *n* nodes running the cluster-tree formation protocol.

4.3 Verification of the cluster-tree formation protocol in mCRL2

We want to verify the property: “eventually every node in the network is either a cluster head or a cluster slave”. In order to check the state of the whole network throughout its operation, we introduce a new action *increase*. In the specification of the process *Node*, every time the type of the node changes to CLUSTER_HEAD or CLUSTER_SLAVE, the action *increase* is performed. Then the property “eventually


every node in the network is either a cluster head or a cluster slave” can be expressed as “for every possible computation at some finite point in the future the action *increase* has occurred $\text{MAX_ID} - 1$ times”. This statement formulated in the modal μ -calculus looks as follows:

$$\mu X(n : \mathbb{N} := 1). ([\textit{increase}]X(n+1) \wedge [\overline{\textit{increase}}]X(n) \wedge \langle \textit{true} \rangle \textit{true}) \vee (n \approx \text{MAX_ID}) \quad (1)$$


For a network of 3 nodes, we manually generated all the topologies in which the network is well connected. For all those cases, we verified the *Network* process against property (1) using the mCRL2 model checker. The toolset generated a number of *counterexamples* for the property.

Below we present three basic scenarios in which the network does not behave as desired. In our representation of topologies that follow, we use dashed lines to show two nodes being in range and full lines to show two nodes being close. Moreover, we give the type of the nodes by annotating them by CH, CS, F and T which stand for cluster head, cluster slave, free and tentative respectively.


Collision of acknowledgements. In the cluster-tree formation protocol whenever a free node receives a BEACON message, it responds with an acknowledgement right away. However, in the physical topology given to the right, since the BEACON is sent by broadcast there could be more than one node responding on the same channel, and this leads to collision.



Collision of associate messages. Once a tentative cluster head reaches the timeout without finding a closer cluster head, it returns to the parent channel and sends an associate message to the parent. However, in the physical topology given here it is likely that two nodes send associate messages at the same time to the same cluster head. This leads to a collision, and in such a case the specification does not state how the tentative cluster heads should proceed.




Narrow bridge problem. Consider the physical topology to the right, in which the connection of the original cluster head to the rest of the nodes in the network is only possible through a close node. In this case, the direct child of the cluster head becomes a cluster slave. Subsequently since the cluster slave does not send out BEACON messages in its vicinity, the rest of the network never joins the tree. This problem arises due to the fact that the condition for becoming a cluster slave does not take the importance of the position of a node into account.



For certain topologies, the narrow bridge problem is inevitable for every possible execution. Moreover, given some fixed topologies and the initial conditions in which all nodes start on the same channel, collision of acknowledgements is also inevitable for every possible execution. This means that in these cases, there is no execution of the protocol in which a full cluster-tree is achieved.

4.4 Scalability of the cluster-tree formation protocol

The main problem with the cluster-tree formation protocol however is its performance. In this section we give an argument for how the performance scales when the number of nodes in the network grows, by finding lower bounds on the termination time of the cluster-tree formation protocol. In order to proceed with our investigation we consider topologies in which the narrow bridge problem does not arise, and we solve the issue of collision of acknowledgements by changing the protocol as follows. After a free node receives a BEACON message, it does not acknowledge it, but instead it generates a random waiting



(a) A simple network of 3 nodes, (b) A network with a balanced binary tree topology with height $h = 3$, with node 1 the original cluster head. The number of nodes on level i of the tree is 2^i , except for $i = 3$. The head and node 2 and 3 free nodes. Triples used in our scalability argument are encircled by dashed lines.

Figure 2: Scalability analysis of a network with 3 nodes and a network corresponding to a binary tree.

time. If the cluster head is close, the free node waits for the random time and then it associates with the cluster head and becomes a cluster slave. Otherwise, it directly changes its type to tentative and follows the procedure we earlier specified.

In the next step we take some of the witnesses of property (1), i.e., paths in the state space over which the property holds, and we calculate the amount of time it takes for the procedure to complete in a network of three nodes, with the physical topology given in Figure 2a. In order to find witnesses we use the `lpsxsim` tool, which given a linearised mCRL2 process specification allows the manual generation of traces. Using this tool and our knowledge of the protocol we look for the smallest sequence of actions for which property (1) is satisfied.

Let `MAX_ID = 3`, `NUM_OF_CHANNELS = 3` and the minimum time for nodes to remain tentative be 2. We find an initial configuration for which the network model satisfies property (1), and has the smallest trace for a witness. In the witness we find the minimum time for the formation of the cluster is 8 time-slots. In the setting of a slotframe which reserves 2 out of 12 time-slots for cluster-tree formation and maintenance, this means that the cluster-tree formation takes at least 3 slotframes to complete. Suppose that a time-slot is 120 milliseconds long, then it takes 5760 milliseconds for the algorithm to form a tree.

We now give our argument on how performance scales for larger networks. First, consider a tentative node with $id = A$ which wants to join a cluster head in tier t . Since every association needs to be handled by the original cluster head, the cluster head on tier t has to forward the association request to its parent on tier $t - 1$, and so forth until it reaches the original cluster head. In the same manner the association response needs to travel back to node A . However, in case other requests arrive at a cluster head which is involved in this process at the same time, the performance of the protocol quickly deteriorates. Therefore for a node at tier t the association procedure takes at least $2t$ to complete.

For simplicity consider the physical topology to be according to a balanced binary tree of height h . An example of such a network is given in Figure 2b, in which node 1 is the original cluster head. The edges in this tree indicate the pairs of nodes that are in range of each other. The tree structure resulting from the protocol in such a case is identical to the physical topology.

We decompose the balanced binary tree into triples of nodes, as shown in Figure 2b. Let $0 < i < h$, each triple consists of a parent with node height $i - 1$ and its children with node height i . Consider a point in the execution of the cluster-tree formation protocol in which the parent node has already joined the network but its two children are still free. Using the witness we derived earlier, we see that excluding

the messages ASSOCIATE and ASSOCIATE_ACK the minimum time required for message exchange inside the triple (BEACON message and tentative waiting time) is 5 time-slots. Therefore, following our earlier observation for a single node with node height i , the time for the completion of the procedure is $5 + 2i$, and for two nodes this time is $5 + 4i$.

Consider the ideal case in which: all messages avoid collisions, every cluster head is involved in at most one association at a time and every free node listens on the channel of its future parent exactly as soon as the parent sends a beacon. Since in a balanced binary tree the number of nodes with node height $i < h$ is exactly 2^i the association of all the nodes with height i takes $5 + 2^{i+1}i$ time-slots to complete.

There is one exception for simultaneous handling of associations which happen in the two separate subtrees of the original cluster head. In this case two associations which arrive at the original cluster head only within 2 time-slots from each other can be processed by the original cluster head. By using this insight the minimum time for association of all nodes of height i becomes $7 + 2^i i$.

And lastly, we know that a balanced binary tree of height h has at least 1 node with node height h , therefore the lower bound on the number of time-slots needed for the formation of the cluster-tree in such a topology is given by the formula:

$$5 + 2h + \sum_{i=1}^{h-1} (7 + 2^i i) \quad (2)$$

As an example, consider a network with 500 nodes. A balanced binary tree with 500 nodes has height $h = 8$. This means that the formation of the cluster-tree in such a network takes at least 27 minutes to complete in the extremely ideal case. With the introduction of collisions, random back-offs, and the increase of the number of channels this time increases substantially.

Based on the lower bound given by formula (2), we conclude that for a network of n nodes the optimal time for the execution of the cluster-tree formation protocol is $\Omega(n \log(n))$, and hence super linear.

5 Further details on modelling and verification

The activities of modelling and verification of the cluster-tree formation protocol was done in a period of five weeks, with approximately 100 hours of effort. This included becoming familiar with the domain and consultation with the designers and domain experts. The generated state spaces are so small that tool effort for verification is negligible.

Our models and the modal formula are included in the model repository:

- `ClusterFormation.mcr12`: the mCRL2 specification corresponding to Sections 4.1 and 3.2.
- `ClusterFormation.NoAcks.mcr12`: the mCRL2 specification of the protocol with no BEACON acknowledgements, as described in Section 4.4.
- `property1.mcf`: the correctness property of the protocol (formula (1)).

6 Conclusion

In this paper we showed how specification and verification of a network protocol by the mCRL2 language and toolset can give insights to the operation of wireless network protocols. Remarkably, by only modelling a few number of nodes (3 nodes in our case) we were able to detect issues in the protocol design such as scalability. The key message here is that by modelling on a suitable abstraction level and with a sufficient number of components, one can effectively reason about the performance and correctness of wireless network protocols.

Based on our findings, we formulate a number of guidelines which can help in the general design of protocols using the IEEE 802.15.4 TSCH MAC layer operation mode:

- Broadcast messages should not be acknowledged (which is already common knowledge).
- Any other sort of communication triggered by single shared events (like a broadcast), should use back-off schemes or channel sensing strategies to promote robustness.

The following guidelines are specific to the design of an improved cluster-tree formation protocol:

- A node should collect information regarding its neighbours before becoming a cluster slave.
- To achieve better performance, the cluster-tree formation should be done during a dedicated, intensive period with no slotframe, in which only cluster-tree formation messages are exchanged.
- The association requests should not be forwarded one by one to the original cluster head.

References

- [1] *mCRL2 Homepage*. <http://www.mcr12.org/>. Accessed: 2017-01-13.
- [2] (2012): *IEEE Standard for Local and Metropolitan Area Networks—Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs), Amendment 1: MAC Layer*. IEEE Computer Society 2012, doi:10.1109/IEEESTD.2012.6185525.
- [3] (2016): *IEEE 802.15.4-2015 - IEEE Standard for Low-Rate Wireless Networks*. IEEE Computer Society, doi:10.1109/IEEESTD.2016.7460875.
- [4] Conrad Dandelski, Bernd-Ludwig Wenning, Michael Kuhn & Dirk Pesch (2016): *RLL-reliable low latency broadcast data dissemination in dense wireless lighting control networks*. In: *Emerging Technologies and Factory Automation (ETFA), 2016 IEEE 21st International Conference on*, IEEE, pp. 1–8, doi:10.1109/ETFA.2016.7733751.
- [5] Domenico De Guglielmo, Simone Brienza & Giuseppe Anastasi (2016): *IEEE 802.15. 4e: A survey*. *Computer Communications* 88, pp. 1–24, doi:10.1016/j.comcom.2016.05.004.
- [6] Wan Fokkink (2013): *Distributed Algorithms: An Intuitive Approach*. The MIT Press.
- [7] Robert G. Gallager, Pierre A. Humblet & Philip M. Spira (1983): *A distributed algorithm for minimum-weight spanning trees*. *ACM Transactions on Programming Languages and systems (TOPLAS)* 5(1), pp. 66–77, doi:10.1145/357195.357200.
- [8] Juan A Garay, Shay Kutten & David Peleg (1998): *A sublinear time distributed algorithm for minimum-weight spanning trees*. *SIAM Journal on Computing* 27(1), pp. 302–316, doi:10.1137/S0097539794261118.
- [9] L Grieco (2015): *Using IEEE 802.15. 4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement*. Technical Report, doi:10.1109/IEEESTD.2012.6185525.
- [10] Jan Friso Groote & Mohammad Reza Mousavi (2014): *Modeling and analysis of communicating systems*. The MIT press.
- [11] Christian Groß, Holger Hermanns & Reza Pulungan (2007): *Does clock precision influence ZigBees energy consumptions?* In: *International Conference on Principles of Distributed Systems*, Springer, pp. 174–188, doi:10.1007/978-3-540-77096-1_13.
- [12] Florian Kauer, Maximilian Köstler, Tobias Lübker & Volker Turau (2016): *Formal Analysis and Verification of the IEEE 802.15. 4 DSME Slot Allocation*. In: *Proceedings of the 19th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, ACM, pp. 140–147, doi:10.1145/2988287.2989148.
- [13] Nancy A Lynch (1996): *Distributed algorithms*. Morgan Kaufmann.