

Applications of Coloured Petri Nets for Functional Validation of Protocol Designs

Lars M. Kristensen^{1,*} and Kent Inge Fagerland Simonsen^{1,2}

¹ Department of Computer Engineering, Bergen University College, Norway
`{lmk,r,kifs}@hib.no`

² DTU Informatics, Technical University of Denmark, Denmark
`kisi@imm.dtu.dk`

Abstract. Communication protocols constitute central building blocks in most modern IT systems as they define components, rules, and languages that make data communication possible. The development of correct protocols is a challenging engineering discipline, making modelling and validation of protocol design an important application domain for Coloured Petri Nets (CPNs). We illustrate the practical application of CPNs for protocol validation by focusing on selected aspects of four recent projects involving industrial-sized protocols. These projects demonstrate how CPNs can be used to model protocol elements and improve protocol specifications, how state space exploration can be used to verify protocol properties, and how behavioural visualisation in combination with a CPN model provides an effective way of rapidly constructing an executable prototype of a protocol design.

1 Introduction

Communication protocols play an important role in most IT systems. A prominent example is the vast amount of web applications that are in use today for, e.g., online banking, shopping, government administration, and entertainment. The services provided by these applications all rely on the protocols governing the operation of the Internet. Other examples are telecommunication systems, logistic systems with sensors and actuators, and control systems in vehicles. All these systems rely heavily on communication and synchronisation between concurrently executing software components and subsystems. As protocols are to support still more complex services that are critical to both the operation of companies and the everyday life of citizens, it is important that they are working correctly already from the initial deployment.

Protocol engineering [80] typically involves a specification of the *service* that the protocol is to provide. Through a synthesis or design step, a protocol design is developed with the aim of providing the desired service. For protocol design, *functional* and *performance* validation can be conducted to investigate and reason about the properties of the design. Functional validation focuses on the

* Work supported by the Research Council of Norway project 194521 (FORMGRID).

logical correctness of the protocol such as the absence of deadlocks and livelocks, that a request is always followed by a response, or whether the proposed protocol design provides the desired services. Performance validation is concerned with quantitative properties such as delays, throughput, and response time. Eventually, the protocol design is implemented and may then be subject to further testing.

Protocol design is in many cases a challenging task. One reason for this is that the execution of a protocol can proceed in different ways, e.g., depending on which messages are lost in transmission, the scheduling of the protocol entities, the time at which events are received from the environment of the protocol, and the execution path taken by the protocol entities. Another reason is that a protocol by nature involves independently scheduled entities which makes testing and reproduction of executions difficult. All this means that protocols often have a very large number of possible executions. In this process, it is easy for a protocol engineer to overlook important interaction patterns which may in turn lead to gaps or malfunction of the protocol.

The specification of the protocol service and the protocol design is, in many cases, based on natural language descriptions. One example of this is the Request for Comments (RFC) documents published by the Internet Engineering Task Force (IETF) [47]. Natural language specifications of protocols often have many issues that need to be resolved before a properly working implementation can be obtained. One class of issues originates from the fact that such specifications are inherently ambiguous making it difficult to achieve inter-operability between independent implementations. Another source of issues to resolve is that the specifications are often incomplete in that the behaviour of the protocol is not described for all cases.

The challenges outlined above have made protocols a prominent application domain for formal description techniques [46], including Petri Nets [93, 97]. In this paper we concentrate on the use of Coloured Petri Nets (CPNs) [56, 59, 61] for modelling and functional validation of protocol designs. Our purpose is to provide an introduction to, and an overview of, how CPNs have been applied for practical validation of protocol designs. We approach this by presenting selected parts of CPN models and associated results originating from projects conducted in an industrial context with industrial-sized protocols. More specifically, we present in the core of this paper the application of the CPN modelling language, tools, and techniques for functional validation of the following protocols:

The DYMO Routing Protocol. The Dynamic On-Demand Routing Protocol for Mobile Ad-hoc Networks (DYMO) [15] is a routing protocol for mobile ad-hoc networks being developed by the MANET working group of the IETF. The DYMO case study is used to illustrate *protocol modelling* with CPNs and to introduce the basic constructs of the CPN modelling language. Our presentation is based on the CPN model constructed in a project on modelling and validating DYMO [25].

The Generic Access Networks (GAN) Architecture. The GAN protocol architecture [2] is developed by the 3rd Generation Partnership Project (3GPP)

for accessing telephone services via Internet Protocol (IP) networks. The GAN case study is used to introduce the basics of *explicit state space exploration* and show how it can be used in a fully automatic manner as a first step in the *verification* of a protocol design. The presentation is based on the project [30] conducted at TietoEnator A/S to specify the detailed usage of protocol software and services via specialisation of the GAN protocol architecture.

The Routing Interoperability Protocol (RIP). The RIP protocol developed at Ericsson Telebit A/S enables routing of IP packets between core IP networks and mobile ad-hoc networks. The RIP case study is used to illustrate how *application-specific behavioural visualisation* can be applied on top of CPN models. In particular, how it can be used to obtain a first executable prototype of the protocol design allowing for early experiments and for presentation to customers and management with the aim of soliciting protocol design requirements. Our presentation is based on the project [74] conducted in cooperation with Ericsson where CPN modelling was used to specify the operation of the RIP protocol.

The Edge Router Discovery Protocol (ERDP). The ERDP protocol is an IPv6-based protocol allowing edge routers to configure gateways in mobile ad-hoc networks with IP address prefixes. The ERDP case study is used to illustrate how the combined use of CPN modelling, state space exploration, and behavioural visualisation all contributed to identify and resolve design issues and errors during ERDP development. Our presentation is based on the project [67] conducted at Ericsson Telebit A/S on the design of the ERDP protocol.

The rest of this paper is organised as follows. Section 2 provides a high-level overview of CPNs and related techniques used for functional validation of protocol designs. Sections 3-6 then present the application of CPNs on the four protocols introduced above. In Sect. 7 we survey related work where CPNs have been used for protocol validation. Finally, Sect. 8 contains conclusions and outlines directions for future work. The reader is assumed to be familiar with the basic ideas of Petri nets [97] and TCP/IP communication protocols [21]. The reader is referred to [61] for a comprehensive introduction to CPNs, state space exploration, and behavioural visualisation of CPN models.

2 Background: CPNs and Functional Protocol Validation

The CPN modelling language belongs to the family of High-level Petri Nets and combine Petri Nets with the Standard ML (SML) programming language [100]. Petri Nets provide the foundation of the graphical notation and the semantical foundation for modelling concurrency, synchronisation, and communication. The functional programming language SML provides primitives for representing sequential aspects of protocols (such as data manipulation) and for creating compact and parameterisable models. Formal modelling and validation with CPNs

is supported by CPN Tools [95] which provides support for construction, simulation, functional and simulation-based performance analysis of CPN models. The addition of data types and a high-level programming language offered by CPN (in contrast to ordinary Petri nets) is highly important when constructing Petri net models of protocols. As an example, with ordinary Petri nets each message type exchanged between protocol entities need to be present with multiple places, and data manipulation (e.g., comparison of data packet content such as sequence numbers) needs to be modelled relying only on net structure resulting in models that are difficult to comprehend.

The advantage of CPNs (and formal description techniques in general) is that they are based on the construction of executable models that make it possible to observe and experiment with the protocol design prior to implementation and deployment using, e.g., simulation. This typically leads to more complete protocol specifications since the model will not be fully operational until all parts of the protocol have been (at least abstractly) specified. Furthermore, the construction of a formal and executable model helps identify and resolve ambiguities that may be present in a natural language specification. Another advantage is the support for model abstractions that makes it possible to specify the operation of the protocol without being concerned with implementation details such as message layout. A model also makes it possible to explore larger scenarios of a protocol system than what is in many cases practically possible in a laboratory.

2.1 Simulation and Behavioural Visualisation

During a protocol model construction phase it is common to use *interactive simulation* of the CPN protocol model to investigate the operation of the protocol in detail. An interactive simulation is similar to single-step debugging and the execution of the CPN model is viewed directly on its graphical representation and provides a simple way of validating that the model operates as intended. In an interactive simulation, the modeller is in charge and determines the next step by selecting between the enabled events in the current state. Interactive simulation is typically combined with the use of *automatic simulation* which is similar to program execution and the purpose is to execute the CPN model without detailed interaction and inspection. Automatic simulation is typically used for testing purposes, and the modeller typically sets up appropriate breakpoints and stop criteria.

Even though the CPN modelling language supports abstraction and hierarchical modules there can still be a significant amount of detail being presented with this approach, and observing every single step either in an interactive simulation or in a log file based on an automatic simulation is often too detailed a level of observation when investigating the behaviour of a model. Furthermore, even if the CPN model is executable, it still lacks the application- and domain-specific appeal of a conventional software prototype. CPN Tools can use the BRIT-NeY Suite animation framework [111] to create behavioural visualisation [112] and interaction graphics on top of CPN models. The animation framework is a stand-alone application, and CPN Tools invokes the primitives of the animation

framework using remote procedure calls. The animation framework supports a wide range of diagram types via a plug-in architecture that makes it possible to visualise the execution of protocols using both standard diagrams (e.g., message sequence charts) in addition to tailored, application-specific diagrams. In this way it is possible to investigate the behaviour of a protocol design while overcoming the limitations of interactive and automatic simulations. In this paper we give some examples of both standard and application-specific diagrams. The reader is referred to [111] for a comprehensive introduction to the animation framework.

2.2 State Spaces and Verification

Verification of behavioural properties of protocols with CPNs [66] is supported by explicit *state space exploration* [6]. In its simplest form this approach involves computing a directed graph where the nodes corresponds to the set of reachable states of the CPN model and the arcs represent occurrences of events causing state changes. State spaces can be constructed fully automatically by the state space tool in CPN Tools and guarantees complete coverage of all executions. State space hence provides a highly systematic error-detection technique that make it possible to automatically (i.e., algorithmically) check whether a protocol has a formally stated desired property. In addition, state space methods have the advantage that counter examples (error-traces) can be automatically synthesised if the protocol does not satisfy a given property.

The main disadvantage of state space exploration is the inherent state explosion problem [103], and a multitude of advanced state space methods have been developed aimed at alleviating the inherent state explosion problem. Early work on addressing state explosion in the context of CPNs concentrated on computer tool support for, and initial experiments with, the equivalence [57], symmetry [20, 24, 48, 58], and the stubborn set methods [102]. The symmetry and equivalence methods rely on constructing a condensed state space where each node represents an equivalence class of states and each arc represents an equivalence class of events. The symmetry method has, e.g., been applied on a mutual exclusion protocol [62] and an embedded systems protocol [81]. The equivalence method has only been used on a small stop-and-wait protocol [63] due to the obligation of providing a manual soundness proof for the user-provided equivalence relation. The stubborn set method [101, 103] relies on analysing enabling and disabling dependencies between events and use this to explore only a subset of the events in each state encountered during state space exploration. The rich SML-based inscription language which is fundamental building block of the CPN modelling language, however, poses problems for the analysis of transition dependencies in the context of CPNs [72] – unless relying on an unfolding of the CPN model to the equivalent Place/Transition net. Hence, restrictions on the modelling language are required to apply the stubborn set method without relying on unfolding. Another widely used verification approach in the context of CPNs is based is the methodology of [9]. A central component of this approach is an explicit modelling of both the protocol and its service, and the use

of finite-state automata language comparison as a criteria for checking that the protocol conforms to the specified service. Recent work on addressing the state explosion problem in the context of CPNs has concentrated on making more economical use of memory resources when exploring the state space. Memory is (in many cases) the limiting factor in state space exploration of CPN models due to the large state vectors. This work resulted in the development of the sweep-line method [19, 60] and the comback method [27, 110]. The sweep-line suite of methods [8, 19, 68, 69, 83] is aimed at on-the-fly verification and exploits a notion of progress found in many concurrent systems. Exploiting progress allows for the deletion of states from memory during a progress-first traversal of the state space. This in turn reduces peak memory usage. The sweep-line method has been used [34, 35, 41, 105] for the verification of several industrial-sized protocols specified using the CPN modelling language. The comback method can be viewed as an exploration-order independent storage mechanism based on hash compaction [98, 113]. It allows the usually large state vectors of CPN models to be stored in compact form, and the full state vector of a state is reconstructed when needed for comparison with newly generated states. Unlike the classical hash compaction method, the comback method guarantees full coverage of the state space. The ASAP model checking platform [109] has support for a number of these advanced state space methods – including methods developed outside the context of CPNs.

2.3 Formal Specification Techniques for Protocols

CPNs and Petri Nets represents one approach to the formal specification and verification of protocols. Historically, several non-Petri nets based languages targeting protocol specification have been developed, in particular in relation to telecommunication standardisation efforts [75, 94]. The Language of Temporal Ordering Specification (LOTOS) [1, 14, 50] was developed as part of International Standardisation Organisation (ISO) efforts and linked to the development of the Open Systems Interconnection (OSI) reference model. LOTOS is founded on the Calculus of Communicating (CCS) [86] and add a data type component to CCS based on algebraic specification. The Extended State Transition Language (Estelle) [49] also originated from OSI standardisation efforts and is based on extended finite state machines [13] combined with extensions to the PASCAL programming language. The Specification and Description Language (SDL) [55] has evolved in several generations since 1980 within the International Telecommunication Union - Telecommunication Sector (ITU-T). SDL is based on communicating extended state machines and has in later versions been equipped with a formal semantics [55] making it amendable for formal verification. A Unified Modelling Language (UML) Profile [52] linking SDL and UML also exists. A comparison of these classical specification languages can be found in [5]. Estelle, SDL, and CPNs are all equipped with a language for modelling data manipulation, but have a different theoretical foundations (extended state machines versus Petri Nets). Another difference is that CPNs have very few (but still powerful) modelling constructs in contrast to languages such as Estelle and

SDL which have a large and complex set of language constructs to describe the behaviour of protocol entities and their interaction. From this perspective, CPNs provide a simpler and more lightweight approach to protocol modelling which at the same time less implementation specific than, e.g., typical SDL protocol specifications. In that respect, CPNs are close to languages like LOTOS that focus more on abstract and implementation independent protocol specification. Within ITU-T languages has also been developed related to protocol data representation. The Abstract Syntax Notation One (ASN.1) [53] is a notation of describing data structures carried in messages exchanges between protocol entities. The Encoding Control Notation (ECN) [54] is a language for specifying ASN.1 encoding rules. In terms of specification of data structures, the SML data types for defining colour sets in CPNs provide similar capabilities as ASN.1. The Testing and Test Control Notation 3 (TTCN-3) [26] is a language for writing protocol test specification.

The Process Meta Language (Promela) language [46] providing the modelling foundation of the SPIN tool [45] has been widely used for protocol design and verification. Promela is based on Communication Sequential Processes (CSP) [44] and is in contrast to CPNs, a textual modelling language with a different theoretical foundation. In a UML context, state diagrams (charts) [43] are used for modelling protocol modules (e.g., [84]), and message sequence charts (MSCs) [51] (sequence diagrams in UML) are being used in particular for specifying protocols requirements that can later be used in protocol verification [4, 38]. MSCs have also been used for protocol specification using higher-level control flow constructs. In contrast to MSCs which are action-oriented, then state charts and CPNs are both state and action-oriented modelling formalisms. Timed automata [7] as supported, e.g., by the UppAal tool has also been used for the specification and verification of protocols (e.g., [29, 96]). The UppAal models consists of a network of network of communicating timed automata, and are specifically suited for modelling and verifying protocol where continuous timing constraints are essential. In comparison, the timed concepts provided by CPNs is a discrete time concept of time. An example on the use of CPNs to model protocols with time constraints can be found in [71].

3 The DYMO Protocol

Modelling a protocol involves developing a representation of the *messages* (or packets) exchanged between the *protocol entities*, the *procedure rules* and *internal state* of the protocol entities guarding the processing of messages, and developing a model of the *environment* in which the protocol is being executed. The environment model typically encompass an abstract representation of the communication medium (or channel) over which the protocol operates. The primary purpose of this section is to illustrate how these protocol elements can be represented in the CPN modelling language using the Dynamic On-demand Routing Protocol (DYMO) [15] for mobile ad-hoc networks as an example. This section additionally shows how to construct compact parameterised CPN models

where the number of protocol entities can easily be configured, and how communication networks with a dynamic topology can be modelled.

3.1 MANETs and Operation of the DYMO Protocol

A *mobile ad-hoc network* (MANET) comprises a collection of mobile nodes, such as laptops, personal digital assistants, and mobile phones, capable of establishing a communication infrastructure for their common use. Ad-hoc networking differs from conventional networks in that the nodes operate in a fully self-configuring, autonomous and distributed manner, without any preexisting communication infrastructure such as base stations and routers. Network layer and routing protocols for ad-hoc networking (including the DYMO protocol) are currently under development by the IETF MANET working group.

The operation of the DYMO protocol consists of two parts: *route discovery* and *route maintenance*. Route discovery is used to establish routes between nodes and begins with an *originator node* multi-casting a Route Request (RREQ) message to all nodes in its immediate range. A RREQ message has a *sequence number* to enable other nodes in the network to judge the freshness of the route request. The ad-hoc network is then flooded with RREQs until the request reaches the *target node* (provided that there exists a path from the originating node to the target node). The target node replies with a Route Reply (RREP) message unicasted hop-by-hop back to the originator node. The route discovery procedure is requested by the Internet Protocol (IP) layer on a node when it receives an IP packet for transmission and does not have a route in its routing table to the target node.

Figure 1(left) depicts the topology of a MANET consisting of six nodes numbered 1–6. An edge between two nodes indicates that the nodes are within direct transmission range. In this case, we assume that all communication links are symmetric. Figure 1(right) (to be discussed below) lists for each node the *routing table* entries created as a result of executing a routing discovery procedure with node 1 as the originator node and node 6 as the target node. The routing table entries in Fig. 1(right) are specified as a pair (*target, nexthop*). The second column specifies the entries that are created as a result of a node receiving the RREQ. The third column lists the entries created as a result of receiving the corresponding RREP. When explaining the operation of the DYMO CPN model below, we will use the scenario in Fig. 1 as a running example.

The *message sequence chart* (MSC) in Fig. 2 depicts one possible exchange of messages in the DYMO protocol when the originating node 1 establishes a route to target node 6 in the topology in Fig 1(left). Solid arcs represent multi-cast transmission and dashed arcs represent unicast transmission. In the MSC, node 1 multi-casts a RREQ which is received by nodes 2 and 3. When receiving the RREQ from node 1, nodes 2 and 3 create an entry in their routing table specifying a route back to the originator node 1. Since nodes 2 and 3 are not the target of the RREQ they both multi-cast the received RREQ to their neighbours (nodes 1, 4 and 5, and nodes 1 and 6, respectively). Node 1 discards these messages as it was the originator of the RREQ. When nodes 4 and 5 receive




Fig. 1. Example MANET topology (left) and routing table entries (right)

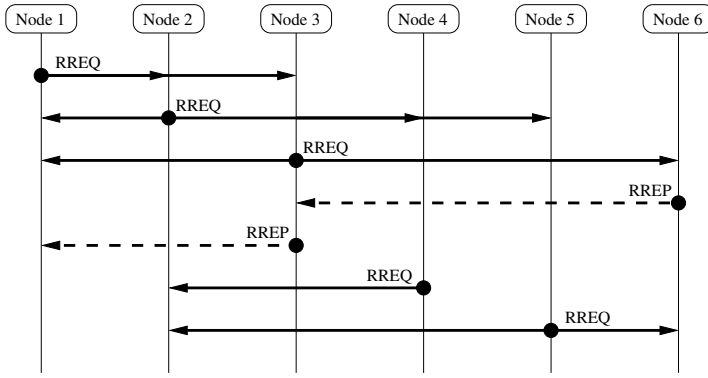


Fig. 2. Message exchange scenario showing DYMO route discovery procedure

the RREQ they add an entry to their routing table specifying that the originator node 1 can be reached via node 2. When node 6 receives the RREQ from node 3, it discovers that it is the target node of the RREQ, adds an entry to its routing table specifying that node 1 can be reached via node 3, and unicasts a RREP back to node 3. When node 3 receives the RREP it adds an entry to its routing table stating that node 6 is within direct range, and use its entry in the routing table that was created when the RREQ was received to unicast the RREP to node 1. Upon receiving the RREP from node 3, node 1 adds an entry to its routing table specifying that node 6 can be reached using node 3 as the next hop. The RREQ is also multi-casted by node 4, but when node 2 receives it again, it will be discarded by node 2 because it has already processed the RREQ message once. Node 5 also multi-casts the RREQ, but nodes 2 and 6 also discard the RREQ message as it has already been received once. From Fig. 1(right) it can be seen that upon completion of the route discovery procedure, a bidirectional route has been discovered and established between node 1 and node 6 using node 3 as an intermediate hop.

The topology of a MANET changes over time because of the mobility of the nodes. DYMO nodes therefore perform *route maintenance* where each node

monitors the links to the nodes it is directly connected to. The DYMO protocol has a mechanism to notify nodes about routes that become broken due to nodes moving out of range of each other. This is done by sending Route Error (RERR) messages which have the effect of informing nodes using the broken route that a new route discovery is needed in order to reestablish a communication path.

3.2 CPN Model Overview and Message Modelling

The DYMO CPN model is a hierarchical model organised in 14 *modules*. Figure 3 shows the *module hierarchy* of the CPN model. Each node in Fig. 3 corresponds to a *module* with **System** representing the top-level module of the CPN model. An arc leading from one module to another indicates that the latter is a *submodule* of the former. The model is organised into two main parts. The **DYMOProtocol** module and its nine submodules model the DYMO protocol entities including the internal state of the protocol entities and the procedure rules for receiving messages, internal processing, and sending of messages. The **MobileWirelessNetwork** module and its two submodules model the environment for the DYMO protocol. This includes the modelling of how messages are transmitted over a wireless link and the modelling of how the mobility of the nodes affects the current topology of the network. The division of the model into submodules reflects the structure of the DYMO specification [16] and hence maintains a close structural relationship between the natural language specification and the formal CPN model. The CPN model does not capture the transmission of payload from the application layer as the focus of the model is on the route establishment and maintenance of the DYMO protocol.

The top-level module **System** is shown in Fig. 4 and is used to connect the two main parts of the model. It corresponds to the **System** node in Fig. 3. The

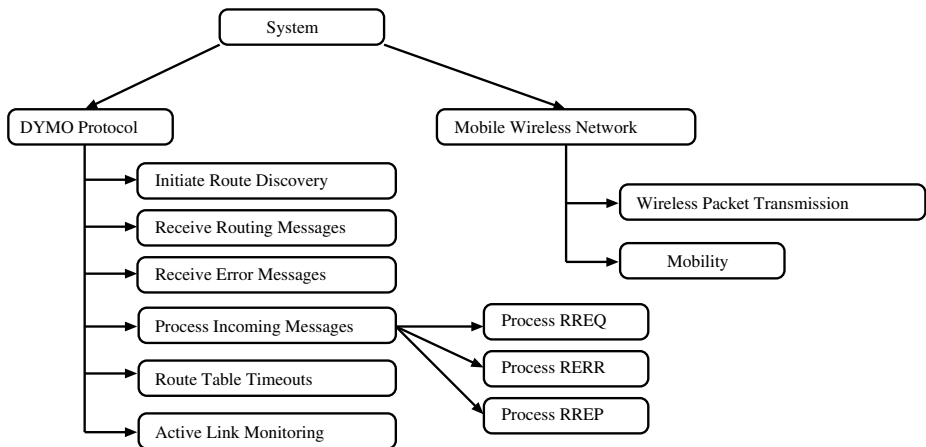


Fig. 3. Module hierarchy for the DYMO CPN model

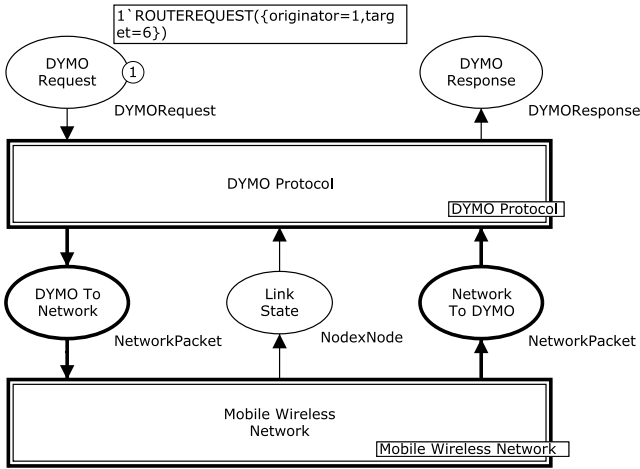


Fig. 4. Top-level System module of the DYMO CPN model

module has two *substitution transitions* drawn as rectangles with double-line borders. Each of the substitution transitions have an associated tag positioned next to it specifying the name of the associated submodule. The `DYMOProtocol` substitution transition has the `DYMOProtocol` module as its associated submodule, and the `MobileWirelessNetwork` substitution transition has the module `MobileWirelessNetwork` as its associated submodule. In this model, the substitution transition has the same name as its associated submodule (but this is not generally required).

The two *socket places* `DYMOToNetwork` and `NetworkToDYMO` connected to the substitution transition `DYMOProtocol` are used to model the interaction between the DYMO protocol and the MANET environment as represented by the submodules of the `MobileWirelessNetwork` substitution transition. The socket place `LinkState` is used to model the *active link monitoring* that nodes perform to check which neighbour nodes are still reachable. When the DYMO protocol module sends a message, it will appear as a token representing a network packet on the socket place `DYMOToNetwork`. Similarly, a network packet to be received by the DYMO protocol module will appear as a token on the `NetworkToDYMO` socket place. Each of the socket places in Fig. 4 (places connected to a substitution transition) is associated with a *port place* in the submodule associated with the substitution transition that the socket place is connected to. The association between a socket and a port place has the effect that the port and the socket places will always have identical markings (tokens). An arc leading to a socket place from a substitution transition means that transitions on the submodule associated with the substitution transitions will add tokens on this place. Analogously, an arc leading from a socket place to a substitution transition means that transitions on the submodule will remove tokens from this place.

The colour set (data types) of each place determining the kind of tokens that can reside on the place is written below each place. The colour set declarations used in Fig. 4 is provided in Fig. 5. A record colour set is used for representing the packets transmitted over the wireless links. A `NetworkPacket` consists of a source (field `src`), a destination (field `dest`), and some `data` (payload). The DYMO messages are designed to be carried in User Datagram Protocol (UDP) datagrams. This means that the network packets are abstract representations of IP/UDP datagrams. The model abstracts from all fields in the IP and UDP datagrams (except source and destination fields) as only these impact the DYMO protocol logic. The source and destination of a network packet are modelled by the `IPAddr` colour set. There are two kinds of IP addresses in the model: `UNICAST` addresses and the `LL_MANET_ROUTERS` multi-cast address. The multi-cast address is used, e.g., in route discovery when a node is sending a `RREQ` to all its neighbouring nodes. Unicast addresses are used as source of network packets and, e.g., as destinations in `RREP` messages. A unicast address is represented as an integer from the colour set `Node`. Hence, the model abstracts from real IP addresses and identify nodes (communication interfaces) using integers in the interval $[1; N]$ where N is a model parameter specifying the number of nodes in the MANET.

```
(* --- Nodes and abstract IP/UDP messages --- *)
colset Node          = int with 0 .. N;
colset IPAddr       = union UNICAST : Node + LL_MANET_ROUTERS;

colset NetworkPacket = record src  : IPAddr * dest : IPAddr *
                          data : DYMOMessage;

(* --- DYMO service --- *)
colset RouteRequest = record originator : Node * target : Node;

colset DYMORequest  = union ROUTEREQUEST : RouteRequest;

colset RouteResponse = record originator : Node * target : Node *
                          status       : BOOL;

colset DYMOResponse = union ROUTERESPONSE : RouteResponse;
```

Fig. 5. Colour set declarations for nodes, network packets, and DYMO service

The two places `DYMORequest` and `DYMOResponse` in Fig. 4 are used to interact with the service provided by the DYMO protocol. A route discovery for a specific destination is requested by putting a token on the `DYMORequest` place and a DYMO response to a route discovery request is then provided by DYMO

as a token via the `DYMOResponse` place. The colour sets `DYMORequest` (see Fig. 5) specifies the identity of the `originator` node requesting the route and the identity of the `target` node to which a route is to be discovered. Similarly, a `DYMOResponse` message contains a specification of the `originator`, the `target`, and a boolean `status` specifying whether the route discovery was successful. The colour sets `DYMORequest` and `DYMOResponse` are defined as union types to make it easy to later extend the model with additional requests and responses. By setting the *initial marking* of the place `DYMORequest`, it can be controlled which route discovery requests are to be made.

The small circles and associated boxes in Fig. 4 show the *current marking* of the CPN model. The small circle positioned inside a place indicates the number of tokens on the place in the current marking. In Fig. 4, there is a single token on the place `DYMORequest` with colour `ROUTEREQUEST({originator=1,target=6})` as specified in the box positioned next to the small circle. This marking corresponds to the DYMO protocol being requested to establish a route from node 1 to node 6 as considered in the scenario in Fig. 1.

3.3 Modelling the DYMO Protocol Entities

The top-level module for the DYMO protocol part of the CPN model is the `DYMOProtocol` module shown in Fig. 6. The module has five substitution transitions modelling initiating route requests (substitution transition `InitiateRouteDiscovery`), reception of `RREQ` and `RREP` messages (substitution transition `ReceiveRoutingMessages`), the reception of `RERRs` (substitution transition `ReceiveErrorMessage`), processing of incoming messages (substitution transition `ProcessIncomingMessages`), and timer management associated with the routing table entries (substitution transition `RouteTableEntryTimeouts`). The places `DYMORequest`, `LinkState`, and `NetworkToDYMO` are *input port places* of the module as indicated by the `In` tag positioned next to them. Each of these places are associated with the accordingly named socket places in Fig. 4. Similarly, the places `DYMOToNetwork` and `DYMOResponse` are *output port places* as indicated by the `Out` tag positioned next to them, and they are associated to the accordingly named socket places in Fig. 4.

All submodules of the substitution transitions in Fig. 6 create and manipulate DYMO messages which are represented by the colour sets defined in Fig. 7. The definition of the colour sets used for modelling the DYMO messages is based on a direct translation of the description of DYMO messages as found in the DYMO specification [16]. In particular, the same names of message fields as in [16] have been used. The model abstracts from the compact packet layout defined for the DYMO protocol. This is done to ease the readability of the CPN model, and since the packet layout is not important when considering only the functional operation of the DYMO protocol.

The place `RoutingTable` and the place `OwnSeqNum` are used to model the routing table and the sequence number of nodes, respectively, that are maintained as part of the internal state of each mobile node. In the marking depicted in Fig. 6 both of these places contain a multi-set containing six tokens. Within the

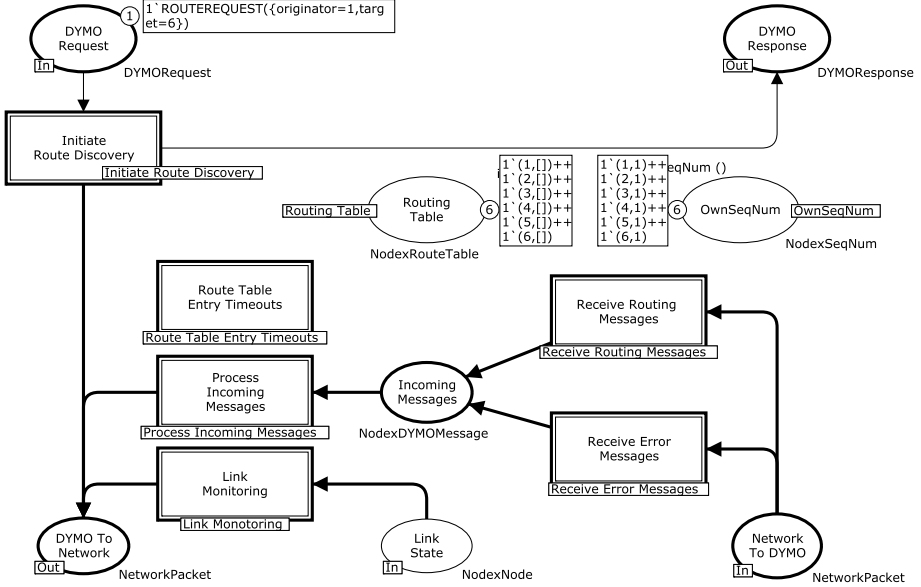


Fig. 6. The DYMOProtocol module

boxes specifying the colours of the individual tokens, ++ (pronounced and) is used to denote union of multi-sets and ' (pronounced of) is used to specify the coefficients (i.e., the number of occurrences of tokens with a given colour). The colour set `SeqNum` used to represent the sequence number of a node was defined above, and the colour set `RouteTable` is defined in Fig. 8. To allow each node to have its own sequence number, we use the colour set `NodexSeqNum`. The marking in Fig. 6 corresponds to a MANET with six mobile nodes. The first component of each token on the place `OwnSeqNum` specifies the identity of a node and the second component specifies the sequence number of the node. Initially, the sequence number of all nodes is set to one. Similarly, it can be seen that the routing table of each mobile node is empty as represented by the empty list (`[]`) specified for each node in the marking of `RoutingTable`.

The submodules of the DYMOProtocol module all need to access the routing table and the sequence number maintained by each node. To reduce the number of arcs in the modules, the routing table and the sequence numbers have been modelled using *fusion sets*. A fusion set allows a set of places in different modules to be linked together into one compound place across the hierarchical structure of the model. In this case, we have a fusion set `OwnSeqNum` (for linking together places modelling the sequence number of each node) and a fusion set `RoutingTable` (for linking the places modelling the routing table of each node). The name of the fusion set which a place belongs to (if any) is written in a tag positioned next to the place.

```

colset SeqNum          = int with 0 .. 65535;
colset NodexSeqNum    = product Node * SeqNum;
colset NodexSeqNumList = list NodexSeqNum;

colset RERRMessage = record HopLimit          : INT *
                          UnreachableNodes : NodexSeqNumList;

colset RoutingMessage = record TargetAddr : Node   * OrigAddr   : Nodes *
                          OrigSeqNum  : SeqNum * HopLimit   : INT   *
                          Dist         : INT;

colset DYMOMessage = union RREQ : RoutingMessage + RREP : RoutingMessage+
                          RERR : RERRMessage;

```

Fig. 7. Colour set declarations for DYMO messages

```

colset RouteTableEntry = record
                          Address      : IPAddr * SeqNum : SeqNum *
                          NextHopAddress : IPAddr * Broken : BOOL   *
                          Dist         : INT;

colset RouteTable      = list RouteTableEntry;
colset NodexRouteTable = product Node * RouteTable;

```

Fig. 8. Colour set declarations for routing table entries

Initiate Route Discovery Module. We consider the `InitiateRouteDiscovery` module shown in Fig. 9 as a representative example of a submodule at the most detailed level of the CPN model. This module specifies how the route discovery procedure is initiated when a request for a route discovery arrives via the `DYMOREquest` input port. The rectangles in Fig. 9 are ordinary transitions (i.e., non substitution transitions) which means that they can become *enabled* and *occur*. In the marking shown in Fig. 9, a token corresponding to a request for a route discovery originating at node 1 and targeting node 6 is present on the `DYMOREquest` place. In this marking, the transition `ProcessRouteRequest` is enabled in the following *binding*:

$$\langle rreq = \{ \text{originator} = 1, \text{target} = 1 \} \rangle$$

which binds the *variable* `rreq` of colour set `RouteRequest` to the value in the `ROUTERREQUEST`. Evaluating the input arc expression on the arc from `DYMOREquest` to `ProcessRouteRequest` results in a multi-set consisting of the single token present on place `DYMOREquest`. The effect of an occurrence of `ProcessRequest`

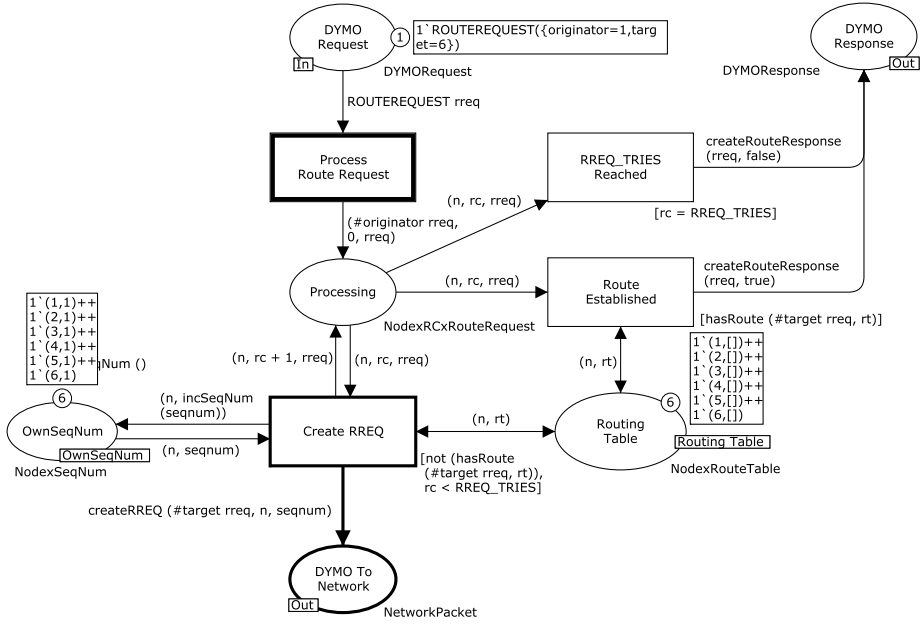


Fig. 9. The Initiate Route Discovery module - initial marking

with the binding above in the marking in Fig. 9 is that the token on **DYMOResponse** is removed and a token is added to place **Processing**. The colour of the token added to **Processing** is obtained by evaluating the *arc expression* on the arc from **ProcessRouteRequest** to **Processing** in the binding from above:

$(\#originator\ rreq, 0, rreq)$

The SML operator **#originator** extracts the originator field in the value bound to **rreq**. The marking resulting from the occurrence of **ProcessRouteRequest** is shown in Fig. 10. A route request being processed is represented by a token on **Processing** over the colour set **NodeXCxRouteRequest** which is a product type. The first component of the token on **Processing** specifies the node processing the route request (i.e., the originator), the second component specifies how many times the RREQ has been retransmitted, and the third component specifies the route request.

In the marking shown shown in Fig. 10, the transition **CreateRREQ** is enabled with the binding:

$(rc = [], rreq = \{originator = 1, target = 1\}, rc = 0, n = 1, seqnum = 1)$

The expression in square brackets positioned next to the **CreateRREQ** transition is a *guard* specifying an additional boolean conditions (beyond the presence of required tokens on input places) for the **CreateRREQ** transition to be enabled. In this case, the guard specifies that for the transition to be enabled, a route

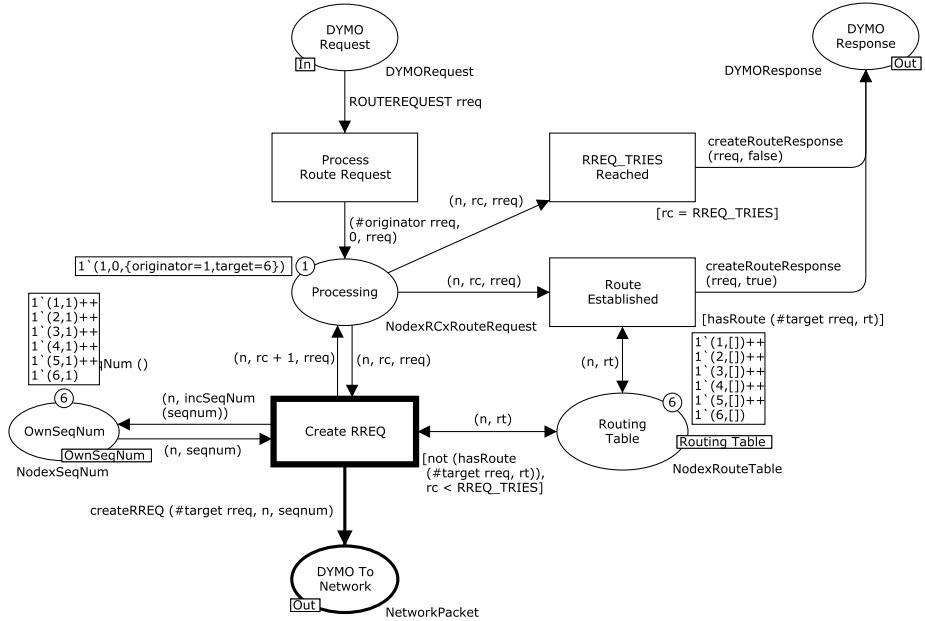


Fig. 10. The Initiate Route Discovery module - after ProcessRouteRequest occurrence

must not already exist in the route table rt to the target node $\#target$, and the number of times rc the current route request has been retransmitted must be less than the retransmission limit $RREQ_TRIES$ for RREQs. The SML function $hasRoute$ used in the guard is implemented as follows:

```
fun hasRoute (target, rt:RouteTable) =
  List.exists (fn {Address, ...} => UNICAST(target) = Address) rt
```

and uses the predefined SML function $List.exists$ to check whether an entry in the route table rt leading to the $target$ node already exists. This is a typical example of how SML is used to represent (sequential) data manipulation.

The marking resulting from the occurrence of $CreateRREQ$ is shown in Fig. 11. When sending a RREQ, the sequence number of node 1 sending the request is incremented by 1 and so is the counter specifying how many times the RREQ has been transmitted. Furthermore, a token corresponding to a network packet containing a RREQ message is produced on place $DYMOToNetwork$. The destination of the packet is set to $LL_MANET_ROUTERS$ since it must be sent to all nodes within reach of node 1.

If a route becomes established (i.e., the originator receives a RREP for the RREQ), the $RouteEstablished$ transition becomes enabled and a token can be put on place $DYMOResponse$ indicating that the requested route has been successfully established. If the retransmission limit for RREQs is reached (before a RREP is received), the $RREQ_TRIES$ Reached transition becomes enabled and a

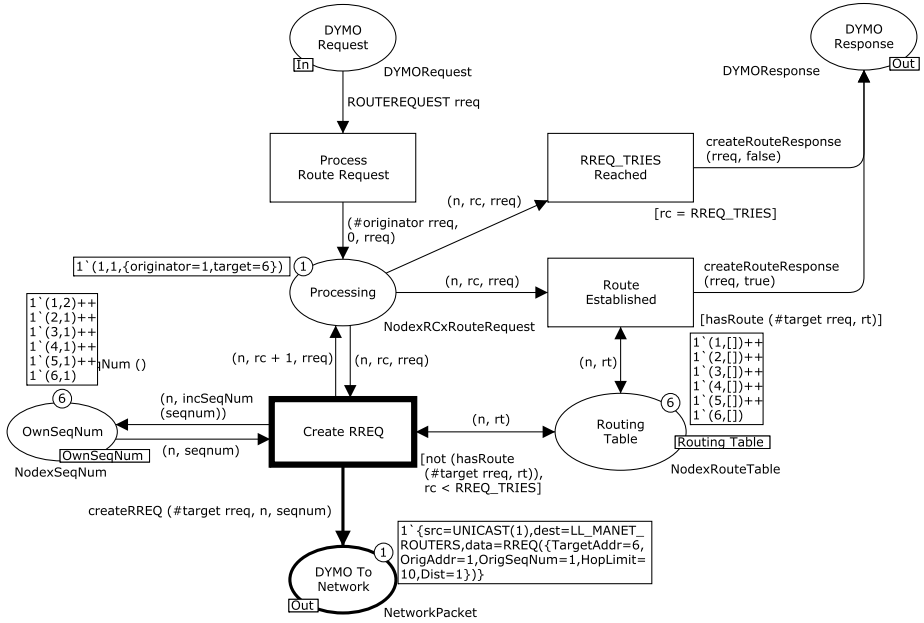


Fig. 11. The Initiate Route Discovery module - after CreateRREQ occurrence

token can be put on place `DYMOResponse` indicating that the requested route could not be established.

3.4 Modelling the DYMO Protocol Environment

The `MobileWirelessNetwork` module shown in Fig. 12 captures the mobile wireless network that DYMO is designed to operate over. It consists of two parts: one part modelling the transmission of network packets represented by the substitution transition `WirelessPacketTransmission`, and one part representing the mobility of the nodes represented by the `Mobility` substitution transition. The places `DYMOToNetwork`, `NetworkToDYMO`, and `LinkState` are associated to the similarly named socket places in Fig. 4. The transmission of network packets is done relative to the current topology of the MANET which is explicitly represented via the current marking of the `Topology` place. The topology is represented using the colour set `Topology` defined in Fig. 13.

The idea is that each node has an adjacency list of nodes that it can reach in one hop, i.e., its neighbouring nodes. The marking of place `Topology` in Fig. 12 corresponds to the topology in Fig. 1(left). This adjacency list is then consulted when a network packet is being transmitted from a node to determine the set of nodes that can receive the network packet. In this way, the dynamic topology is modelled by the addition and removal of nodes from the adjacency lists. The

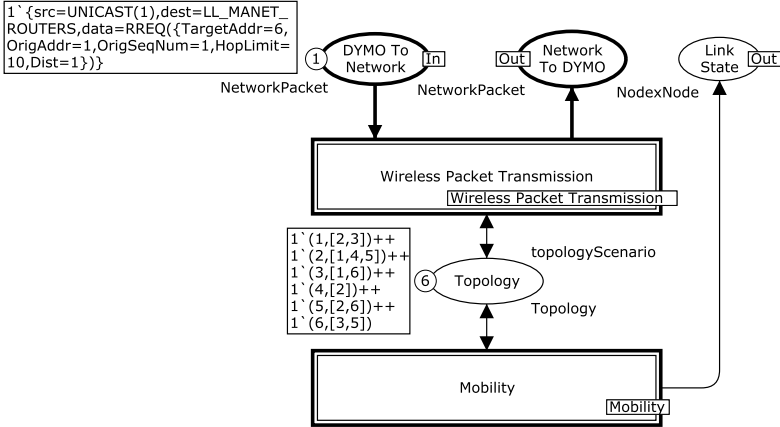


Fig. 12. The Mobile Wireless Network

```
colset NodeList = list Node;
colset Topology = product Node * NodeList;
```

Fig. 13. Colour set declarations for topology modelling

place `LinkState` models that a node can be informed about the reachability of its neighbouring nodes which is used in active link monitoring.

The `WirelessPacketTransmission` module models the actual transmission of packets and is shown in Fig. 14. The module captures how network packets are transmitted via the physical network from one node to the next. Packets are transmitted over the network according to the function `transmit` on the arc from the transition `Transmit` to the place `NetworkToDYMO`. When the `Transmit` transition occurs in a binding where the boolean variable `success` is set to `true`, then all nodes within reach of the sending node will receive the packet. Otherwise, no nodes will receive the packet. The transition `Transmit` is enabled in the marking shown in Fig. 14 (left) and the marking resulting from a successful transmission of the packet on `DYMOToNetwork` is shown in Fig. 14 (right). In this case two tokens are added to place `NetworkToDYMO` corresponding to nodes 2 and 3 receiving the packet being multi-casted from node 1.

In a real network, a transmission could be received by any subset of the neighbouring nodes (e.g., because of signal interference). Here it is only modelled that either all of the neighbouring nodes receive the packet or none of the nodes receive it. This is sufficient because the modelling of the dynamic topology means that a node can move out of reach of the transmitting node immediately




Fig. 14. Transmission of packets - before (left) and after (right) transmission

before the transmission occurs which has exactly the same effect as a signal interference in that the node does not receive the packet. Hence, signal interference and similar phenomena imply that a node does not receive a packet is in the model equivalent to the node moving out of reach of the transmitting node.

3.5 Lessons Learned and Perspectives

The development of the DYMO CPN model was based on the natural language specification provided in the Internet draft [15] specifying the DYMO protocol. The modelling work was done when version 10 [15] was the most recent DYMO specification. In the process of constructing the CPN model and simulating it, several issues and ambiguities in the specification were discovered. The most important ones are summarised in Table 1. These issues were submitted to the IETF MANET Working Group mailing list [82] and issue 1 and 3-7 were acknowledged by the DYMO developers and taken into account in the subsequent version DYMO specification [16] (version 11). Issue 2 was not considered critical as it causes route discovery to fail in scenarios which according to the experience of the DYMO developers would seldom occur in practise.

The modelling conducted with the DYMO protocol illustrates that the construction of a formal and executable model provides a very systematic and comprehensive way of reviewing a protocol design document (such as the DYMO

Table 1. DYMO CPN modelling [25]: issues identified in the modelling phase

Issue	Description
1	When processing a routing message, a DYMO router may respond with a RREQ flood, i.e., a RREQ addressed to the node itself, when it is target for a RREQ message (cf. [15], Sect. 5.3.4). It was not clear from the specification which information to put in the RREQ message, i.e., the originator address, hop limit, and sequence number of the RREQ.
2	When judging the usefulness of routing information, the target node is not considered. This means that a new request with a higher sequence number can make an older request for another node stale since the sequence number in the old message is smaller than the sequence number found in the routing table.
3	When creating a RREQ message the distance field in the message is set to zero. This means that for a given node n an entry in the routing table of a node n' connected directly to n may have a distance to n which is 0. Distance is a metric indicating the distance traversed before reaching n , and the distance between two directly connected nodes should be one.
4	In the description of the data structure route table entry (cf. [15], Sect. 4.1) it is suggested that the address field can contain more than one node. It was not clear why this was the case.
5	When processing RERR messages (cf. [15], Sect. 5.5.4) it is not specified whether the hop limit shall be decremented.
6	When retransmitting a RREQ message (cf. [15], Sect. 5.4), it was not explicitly stated whether the node sequence number should be increased.
7	Version 10 of DYMO introduced the concept of distance instead of hop count. Distance is a more general metric, but in the routing message processing (cf. [16], Sect. 5.3.4) it is incremented by one. We believe it should be up to the implementers how much distance is incremented depending on the metric used.

Internet draft) and how it can contribute to increasing the quality of a protocol design specification. Similar conclusions can also be drawn from other case studies where CPN modelling has been applied to protocols developed in the context of IETF. A CPN model of the DYMO protocol has also been developed in [12] where a considerably more compact CPN model of the DYMO protocol directly targeting state space exploration was developed. A number of other issues related to the functionality of the DYMO protocol were reported in [12]. In comparison to the CPN model in this section, the CPN model developed in [12] provides a more abstract modelling approach that does not use an explicit representation of MANET topology.

4 The GAN Protocol Architecture

This section focuses on how standard behavioural properties of CPNs in combination with explicit state space exploration can be used to verify basic properties of protocols. Furthermore, this section gives an example of how CPNs can be used to model a system spanning multiple protocols and protocol layers. The presentation is based on a project [30] in which CPN modelling and state space exploration was used at TietoEnator Denmark in early phases of developing an implementation corresponding to a particular instantiation [42] of the generic GAN architecture [2] aimed at integrating IP and telephone services.

4.1 GAN Secure Connection Establishment

The Generic Access Network (GAN) [2] architecture specified by the 3rd Generation Partnership Project (3GPP) [3] allows access to common telephone services such as SMS and voice-calls via IP networks. A central part of the GAN architecture is the establishment of a secure connection between a *mobile station* (e.g., a mobile phone) and a *GAN controller* through a *security gateway*. The GAN architecture relies on standardised protocols such as Dynamic Host Configuration Protocol (DHCP) for IP address configuration, IP Security (IPsec) [65] for encryption and authentication, and Internet Key Exchange v2 (IKEv2) protocol [64] for negotiation of IPsec parameters.

The purpose of the CPN model constructed in the project was two-fold. Firstly, to define the scope of the protocol software to be developed by TietoEnator. More specifically, the aim was to determine which parts of the generic GAN specification were to be included in the implementation to be developed by TietoEnator. Secondly, to specify the detailed design and usage of the involved protocol software components. The focus of the CPN model is on the establishment of a secure tunnel and the initial GAN message exchanges since this is where important details were not provided in the full GAN specification. In particular, the full GAN specification [2] contained no clear specification of the IKEv2 message exchange and the states that the protocol entities should be in when establishing a GAN connection (at the time of the project in 2007). Furthermore, the GAN specification only states that IKEv2 and IPsec are to be used, and in which operating modes.

4.2 CPN Model of the GAN Protocol Architecture

The CPN model of the secure connection establishment consists of 31 modules organised into four hierarchical levels. In the following, we present four selected modules from the CPN model. Our purpose is to illustrate how the phases that the protocol entities enter when establishing a GAN connection have been modelled, and provide sufficient detail on the CPN model in order for the reader to interpret the verification results presented later. A more in-depth presentation of the CPN model can be found in [30].




Fig. 15. Top-level module of the GAN model

Figure 15 shows the top-level module which is organised so that it mimics the GAN network architecture. The substitution transition **MobileStation** represents the mobile station which is connecting to the telephone network via an IP network. The place **Wireless Network** connected to **MobileStation** represents the wireless network which connects the mobile station to a wireless router represented by the substitution transition **WirelessRouter**. The wireless router is an arbitrary access point with routing functionality, and is connected to the **Provisioning Security Gateway**, through **NetworkB**. As part of establishing a GAN connection, an encrypted tunnel is established between the mobile station and the security gateway. The encrypted tunnel is provided by the Encapsulating Security Payload (ESP) mode of the IP security layer (IPSec) [65]. To provide such an encrypted tunnel, both ends have to authenticate each other, and agree on both an encryption algorithm and keys. This is achieved using the Internet Key Exchange v2 (IKEv2) protocol [64]. The provisioning security gateway is connected to the **Provisioning GAN Controller** via **NetworkC**. The GAN controllers are connected to the telephone network and perform the relay of traffic to/from the IP networks (**NetworkC** and the **WirelessNetwork**). This in turn allows mobile stations to access the services on the telephone network. The places with thin lines connected to the substitution transitions **Provisioning Security Gateway** and **Provisioning GAN Controller** are used to provide configuration information to the corresponding network nodes. The CPN model does not include modelling of the telephone network as the scope of the CPN model covers the components involved in establishing the connection with the GAN controller. Furthermore, as the purpose of the model was to represent the protocol entities present on each of the nodes in the network architecture, it sufficed that the model encompassed one mobile node, one wireless router, one provisioning security gateway, and one provisioning GAN controller.

The basic exchange of messages in establishing a GAN connection to the provisioning GAN controller involves three steps. The first step is for the mobile station to acquire an IP address on the wireless network using DHCP. The second phase is to create a secure tunnel to the provisioning security gateway. Having established the secure tunnel, the third phase is for the mobile station to open a secure connection to the GAN controller and register itself. Figure 16 (left) shows the **IKEInitiator** module of the mobile station and Fig. 16 (right) shows the **IKEResponder** module of the security gateway. These two peer modules

model the second step of the GAN connection establishment concerned with creating the secure tunnel. Incoming IP packets for the module arrive via the `ReceiveBuffer` input port places. Outgoing IP packets are put in the `SendBuffer` places. The states (phases) that the protocol entities goes through during the IKE message exchange when establishing the secure tunnel are represented by the places connecting the substitution transitions.

The state changes are represented by substitution transitions. The submodules of the substitution transitions specify the processing rules for messages during the individual phases. Figure 17 shows the `Send IKE_SA_INIT Packet` and `Handle_SA_INIT_Request` modules which are the submodules of the two top-most substitution transitions in Fig 16. The `Send IKE_SA_INIT Packet` transition in Fig. 17 (left) takes the IKE Initiator from the state `Ready` to `Await IKE_SA_INIT` and sends an IKE message to the security gateway initialising the communication. The IP address of the security gateway is retrieved from the `Ready` place. Figure 17 (right) shows how the `IKE_SA_INIT` packet is handled by the IKE Responder. The guard of the `HandleSA_INIT_Request` transition ensures that the transition is only enabled if the incoming packet (token) on `IncomingIKERequest` represents a `IKE_SA_INIT` packet. In that case, it sends an IKE packet back to the initiator as specified by the arc expression on the arc from `Handle_SA_INITRequest` and the responder enters the `Wait for EAP Auth` state. The submodules of the other substitution transitions in Fig. 16 are similar.

The establishment of a GAN connection involves multiple layers of the IP network stack. DHCP (used to configure the mobile station) and GAN are application layer protocols, IKE is a transport layer protocol, and IPSec belongs to the network layer. As a consequence, the CPN model of GAN connection spans multiple protocol layers. Furthermore, the protocol entities also access and manipulate the *routing table* and a *security policy database* (SPD) which is maintained at the IP network layer. The establishment of a GAN connection accesses the routing table of a node in order to ensures that packets are put into the secure tunnel, and extracted again at the other end. The SPD describes what packets are allowed to be sent and received by the IP protocol stack, and is also responsible for identifying which packets are to be tunnelled at the mobile station and the security gateway. Each entry in the SPD contains the source and destination addresses to use for matching packets, and an action to perform. Modelled actions are *bypass* (which means allow packet to pass without tunnelling) and *tunnel* (the matched packet is to be sent through an ESP tunnel). As we will see later, the content of the routing table and the SPD play an important role in validating the correctness of the GAN connection establishment. It was therefore required to explicitly represent them in the CPN model.

4.3 Verification of the GAN CPN Model

The goal of applying state space exploration was to verify the completeness of the design. This included verifying that all phases, steps, and messages involved in establishing a secure GAN connection were covered by the design, and the correctness of the connection establishment, i.e., that a GAN connection is

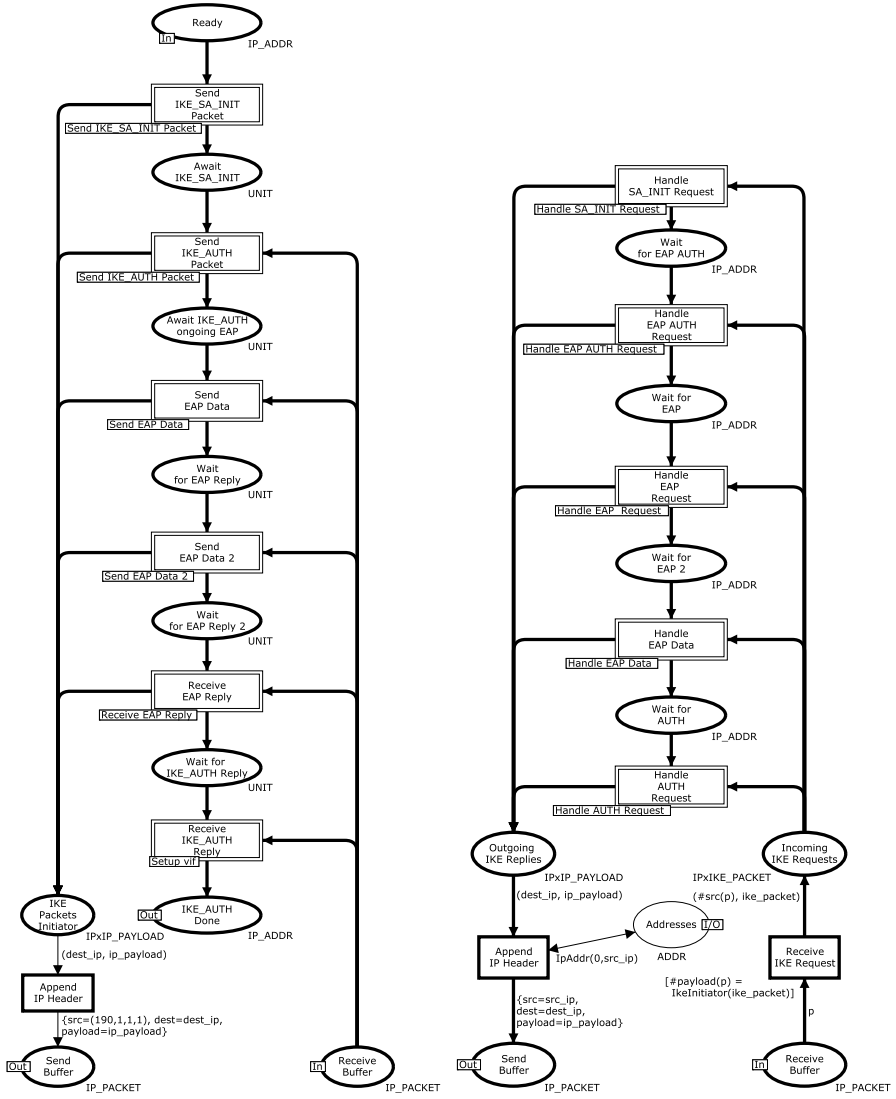


Fig. 16. IKE initiator (left) and IKE responder (right) modules




Fig. 17. Example of IKE initiator (left) and IKE responder (right) submodules

eventually established with the mobile station and the GAN controller being properly synchronised. Verification of the key properties of the design for secure connection establishment was done by *state space exploration*. The basic idea underlying state space exploration is to compute all reachable states and state changes of the CPN model and represent these as a directed graph, where nodes represent markings and arcs represent occurring binding elements. State spaces can be constructed fully automatically by the state space tool in CPN Tools. Verification of the GAN scenario modelling by means of state spaces relied on the use of the *state space report* that can be generated by CPN Tools. The generation of a state space report for the smallest possible configuration of a considered protocol is typically the first step performed when conducting verification of a CPN model.

The state space report is divided into several sections. In the following we present excerpts from the individual sections and explain how they can be used for the verification. Figure 18 shows the first part of the state space report for the CPN model. This part provides some *state space statistics* specifying how large the state space is. It can be seen that the state space consists of 3,854 nodes and 9,225 arcs. The construction of the state space took 4 seconds. Statistics for the strongly connected component graph (SCC-graph) are also specified. It has 3,514 nodes and 8,881 arcs, and was calculated in 2 seconds. The fact that there are fewer nodes in the SCC-graph than in the state space implies that there are non-trivial strongly connected components (SCCs), i.e., SCCs consisting of more than a single state space node. This means that infinite executions exist and that the GAN connection establishment may not terminate. We will investigate the reasons for this at the end of this subsection.

The *boundedness properties* section of the state space report specifies how many and which tokens a place may hold – when considering all reachable states (markings). Figure 19 lists the best upper and best lower integer bounds for selected places in the mobile station module. It can be seen that the first four places modelling the states of the mobile station contain at most one token and may contain zero tokens. Similarly, it can be seen that there is at most one token

State Space		Scc Graph	
Nodes:	3,854	Nodes:	3,514
Arcs:	9,225	Arcs:	8,881
Secs:	4	Secs:	2

Fig. 18. State space report – statistics

Best Integer Bounds	Upper	Lower
Down	1	0
Ready	1	0
VIF open to Prov. SG	1	0
VIF Closed	1	0
Send Buffer	1	0
Receive Buffer	1	0
Network Buffer	1	0
Routing Table	1	0
Security Policy Database	1	1

Fig. 19. State space report – integer bounds

in the send, received, and network buffers. The place `RoutingTable` has a lower integer bound of 0 and an upper integer bound of 1. The lower integer bound is 0 since in the initial marking there are no tokens on this place. During the start-up procedure of the mobile station, a token representing a list of routing table entries is put on this place. The place `SecurityPolicyDatabase` has a best upper and a best lower integer bound of 1. This means that there is always exactly one token present on this place. This is because the security policy database is modelled as a single token being a list containing the current entries in the security policy database.

The *best upper multi-set bound* of a place specifies for each colour in the colour set of the place the maximal number of tokens that is present on this place with the given colour in any reachable marking. This is specified as a multi-set, where the coefficient of each value is the maximal number of tokens with the given value. If the coefficient is zero, then the colour is omitted in the specification. Figure 20 shows part of the state space report providing the upper multi-set bounds for the security policy databases of the mobile station, wireless router, security gateway, and the GAN controller. The upper multi-set bounds specify the possible tokens that can reside on these places and by carefully inspecting these bounds it was possible to validate that the possible entries in the security

```

Mobile Station: Security Policy Database
  1' [{src=((0,0,0,0),0),dest=((0,0,0,0),0),
      nl_info=PayloadList([PAYLOAD_DHCP]),policy=SpdBypass}]++
  1' [{src=((80,1,1,1),32),dest=((12,0,0,0),8),
      nl_info=AnyNextLayer,policy=ESPTunnel(((190,1,1,1),(172,1,1,2))),
      {src=((190,1,1,1),0),dest=((0,0,0,0),0),
      nl_info=AnyNextLayer,policy=SpdBypass}]++
  1' [{src=((190,1,1,1),0),dest=((0,0,0,0),0),
      nl_info=AnyNextLayer,policy=SpdBypass}]

Wireless Router: Security Policy_Database
  1' [{src=((0,0,0,0),0),dest=((0,0,0,0),0),
      nl_info=AnyNextLayer,policy=SpdBypass}]

Security Gateway: Security Policy Database
  1' [{src=((13,0,0,0),8),dest=((80,1,1,1),32),
      nl_info=AnyNextLayer,policy=ESPTunnel(((172,1,1,2),(190,1,1,1))),
      {src=((0,0,0,0),0),dest=((0,0,0,0),0),
      nl_info=AnyNextLayer,policy=SpdBypass}]

GAN Controller: Security Policy Database
  1' [{src=((0,0,0,0),0),dest=((0,0,0,0),0),
      nl_info=AnyNextLayer,policy=SpdBypass}]

```

Fig. 20. State space report – best upper multi-set bounds

policy database were all as desired. Altogether, an inspection of the boundedness properties helped significantly in increasing confidence in the correctness of the design in terms of proper settings of the routing table and the security policy database.

Figure 21 shows the part of the state space report specifying the *home and liveness properties*. The home properties show that there exists a single *home marking*, which is state number 3854. A home marking is a state which can be reached from any reachable state. For the GAN scenario model this means that it is impossible to have an execution sequence starting from the initial state (initial marking) which cannot be extended to reach state 3854. The liveness properties tell us that there is a single *dead marking* which is also state number 3854. A dead marking is a state in which no transitions are enabled. This means that the marking corresponding to node 3854 is both a home and a dead marking.

To obtain information about the marking corresponding to node number 3854, the node number was transferred into the simulator of CPN Tools and displayed graphically on the CPN model. It was then checked (by inspecting the markings of the individual places) that the marking corresponded to the desired terminating state of the GAN connection establishment procedure, i.e., the state where

Home Properties	Liveness Properties
Home Markings: [3854]	Dead Markings: [3854]

Fig. 21. State space report – home and liveness properties

the mobile station has obtained an IP address, has successfully communicated with the provisioning GAN controller, all protocol modules are in a state corresponding to the GAN connection having been established, and the routing tables and security databases contain the correct entries. The fact that state 3854 is the only dead marking tells us that the protocol as specified by the CPN model is partially correct, i.e., if execution terminates we have the correct result. Furthermore, because node 3854 is also a home marking it is always possible to terminate the GAN connection establishment with the correct result.

The analysis above showed that it is always possible to terminate the GAN connection establishment procedure correctly, but there is no guarantee that it will eventually happen. The section of the state space report providing information about *fairness properties* showed that the two transitions `RejectDiscoveryRequest` and `HandleGARCReject` which are part of the GAN controller module were *impartial*. This means that these two transitions occur infinitely often in any infinite occurrence sequence. The two transition occurs if the GAN controller decides to reject an incoming connection from a mobile station. Hence, if the connection establishment procedure does not terminate in the single home and dead marking identified, then it is because the GAN controller keeps rejecting the connection.

4.4 Lessons Learned and Perspectives

The validation of secure connection establishment in the considered GAN scenario is representative for how validation of protocols is typically performed with CPN Tools – as it in practise involves a combination of both simulation and state space exploration. As part of the construction of the GAN model, the support for interactive simulation in CPN Tools was used to perform detailed checks to ensure that the model behaviour was as desired. Even though the use of interactive simulations (and simulation in general) cannot be used to prove correct behaviour, it proved to be very useful in identifying situations related to improper manipulations of the entries in the routing tables and security policy database - or when additional detail not present in the GAN specification had to be worked out and specified. Furthermore, interactive simulation was helpful in identifying issues that led the GAN connection establishment procedure to terminate prematurely, e.g., because a certain phase of the connection establishment was missing in the CPN model. These issues manifested themselves in markings where the GAN connection had not yet been established, but where

no transitions were enabled. This was in particular effective in making explicit where further specification of the message exchanges were required.

The interactive simulation was in later phases replaced with automatic simulation where a number of random executions of the CPN model were performed with the purpose of checking whether the execution of the CPN model resulted in a state in which the GAN connection was properly established. Eventually state space exploration of the CPN model was conducted which succeeded in establishing the key property that a GAN connection will eventually be established provided that the GAN controller does not keep rejecting the connection request. The verification conducted also illustrated the general observation that in many cases, the use of basic state space exploration and the state space report (i.e., investigating standard behavioural properties of Petri nets) are sufficient in establishing key properties of a protocol design. In this case the state space was small in size and could be generated in a few seconds without the use of advanced state space exploration techniques.

5 The Routing Interoperability Protocol

The section show how a CPN model can be augmented with application-specific behavioural visualisation reflecting the execution of the CPN model. This section is based on a project conducted at Ericsson Telebit A/S addressing the specification of the Routing Interoperability Protocol (RIP)¹ for routing packets between IP core networks and mobile ad-hoc networks. The CPN model of RIP augmented with behavioural visualisation was used as an early model-based prototype of RIP. It allowed the protocol design to be discussed among protocol engineers unfamiliar with CPNs, and it also enabled the protocol design to be presented to customers with the purpose of soliciting requirements of the services to be provided by the protocol.

5.1 CPN Model of the RIP Protocol

The main purpose of the routing interoperability protocol is to ensure that a packet flow between a host in the core network and a mobile node in an ad-hoc network is always relayed via one of the closest gateways that connect the core network and the mobile ad-hoc network. Figure 22 shows the top level module of the CPN model which reflects the network architecture that the RIP protocol is designed to operate in. The network architecture consists of three parts: an IPv6 core network represented by the `CoreNetwork` substitution transition (left) and its submodules, a mobile ad-hoc network represented by the `AdHocNetwork` substitution transition (right) and its submodules, and two gateways represented by the substitution transitions `Gateway1` and `Gateway2`. The basic idea in the interoperability protocol is that the mobile nodes register the IPv6 address in the Domain Name Server (DNS) server in the core network that corresponds to

¹ RIP as discussed in this section should not be confused with the Routing Information Protocol[85].

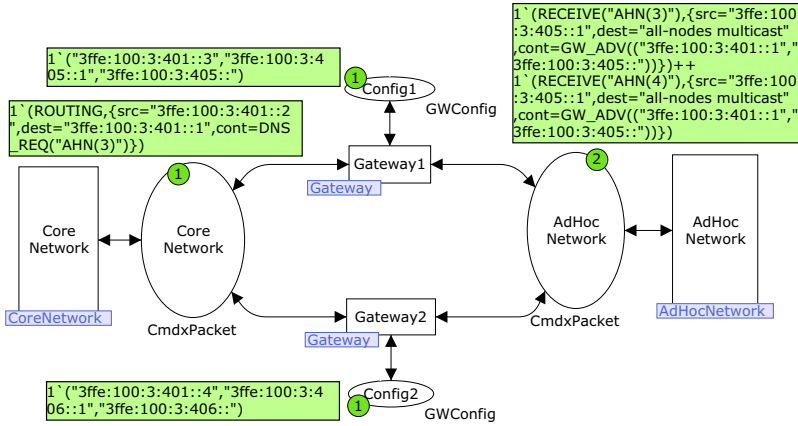


Fig. 22. The System module – top-level module of the CPN model

an IPv6 address prefix announced by the closest (preferred) gateway. Updates to the DNS database managed by the DNS server rely on the Dynamic Domain Name System Protocol [108].

The places `CoreNetwork` and `AdHocNetwork` are used for modelling the packets in transit on the core network and ad-hoc network, respectively. Figure 22 depicts a state in which there is one token on place `CoreNetwork` and two tokens on place `AdHocNetwork`. As an example, place `CoreNetwork` contains one token with the colour:

```
(RECEIVE("3ffe:100:3:401::1"), {src="3ffe:100:3:401::2",
  dest="3ffe:100:3:401::1", cont=DNSREQ("AHN(3)")})
```

representing a DNS request (`DNSREQ`) in transit on the core network from a host with source IPv6 address `3ffe:100:3:401::2` to a DNS server with destination IPv6 address `3ffe:100:3:401::1`. IPv6 addresses are 128-bit and by convention written in hexadecimal notation in groups of 16-bits separated by a colon (:). Leading zeros are skipped within each group and a double colon (::) is a shorthand for a sequence of zeros. Addresses consist of an *address prefix* and an *interface identifier*.

The place `AdHocNetwork` contains two tokens representing gateway advertisements in transit to nodes in the ad-hoc network. The gateways periodically announce their presence to nodes in the mobile ad-hoc network by sending *gateway advertisements* containing an IPv6 *address prefix*. The two `Config` places contain a token representing the configuration of the corresponding gateway. It consists of the IPv6 address of the gateway interface connected to the core network, the IPv6 address of the gateway interface connected to the ad-hoc network, and the address prefix announced by the gateway. Address prefixes are written in the form x/y where x is an IPv6 address and y is the length of the prefix. The mobile nodes in the ad-hoc network configure IPv6 addresses based

on the received gateway advertisements. In the marking depicted in Fig. 22, Gateway1 is announcing the 64-bit address prefix `3ffe:100:3:405::/64` and Gateway2 is announcing the prefix `3ffe:100:4:406::/64`. Each of the gateways has configured an address on the interface to the ad-hoc network based on the prefix they are announcing to the ad-hoc network. Gateway1 has configured the address `3ffe:100:3:405::1` and Gateway2 has configured the address `3ffe:100:3:406::1`. The gateways have also configured addresses on the interface to the core network based on the `3ffe:100:3:401::/64` prefix of the core network.

Figure 23 lists the definitions of the colour sets used in the System module. IP addresses, prefixes, and symbolic IP addresses are represented by colour sets `IPAdr`, `Prefix`, and `Symname` all defined as the set of strings. The colour set `PacketCont` and `Packet` are used for modelling the IP packets. The five different kinds of packets used in RIP are modelled by the `PacketCont` colour set:

`DNS_REQ` modelling a DNS request packet. It contains the symbolic IP address to be resolved to a (numerical) IP address by a DNS server.

`DNS_REP` modelling a DNS reply. It contains the symbolic IP address and the resolved IP address.

`DNS_UPD` modelling a DNS update. It contains the symbolic IP address to be updated and the new IP address to be bound to the symbolic address.

`GW_ADV` modelling the advertisements disseminated from the gateways. An advertisement contains the IP address of the gateway and the announced prefix.

```

colset Prefix = string; (* address prefixes *)
colset IPAdr = string; (* IP addresses *)
colset SymName = string; (* symbolic names *)

colset SymNamexIPAdr = product SymName * IPAdr;
colset IPAdrxPrefix = product IPAdr * Prefix;

colset PacketCont = union DNS_REQ : SymName + (* DNS Request *)
                          DNS_REP : SymNamexIPAdr + (* DNS Reply *)
                          DNS_UPD : SymNamexIPAdr + (* DNS Update *)
                          GW_ADV : IPAdrxPrefix + (* Advertisements *)
                          PACKET; (* Generic payload *)

colset Packet = record src : IPAdr * dest : IPAdr * cont : PacketCont;
colset Cmd = union ROUTING + RECEIVE : IPAdr +
                FLOODING : IPAdr + GWAHRROUTING : IPAdr +
                AHNGWROUTING : IPAdr;

colset CmdxPacket = product Cmd * Packet;
colset GWConfig = product IPAdr * IPAdr * Prefix;

```

Fig. 23. Colour set definitions used in the System module

PACKET modelling generic payload packets belonging to packet flows between hosts and the mobile nodes.

The colour set `Packet` models the packets as a record containing the source, destination, and content. The actual payload (content) and layout of packets are not essential for modelling the interoperability protocol and has therefore been abstracted away. The colour set `Cmd` is used to control the operation of the various modules in the CPN model. The colour set `GWConfig` models the configuration information of the gateway.

The Core Network. Figure 24 shows the `CoreNetwork` module modelling the core network. This module is the immediate submodule of the substitution transition `CoreNetwork` of the `System` module shown in Fig. 22. The port place `CoreNetwork` is assigned to the `CoreNetwork` socket place in the `System` module (see Fig. 22). The substitution transition `Routing` represents the routing mechanism in the core network. The substitution transition `Host` represents the host on the core network, and the substitution transition `DNS Server` represents the DNS server that maintains the DNS database.

The Mobile Ad-hoc Network. Figure 25 depicts the `AdHocNetwork` module modelling the mobile ad-hoc network. The place `Nodes` is used to represent the nodes in the mobile ad-hoc network. The place `RoutingInformation` is used to represent the routing information in the ad-hoc network which is assumed to be available via some routing protocol executed in the ad-hoc network. This routing information enables among other things the nodes to determine the distance to the reachable gateways. Detailed information about the colour of the token on place `RoutingInformation` has been omitted.

Figure 26 lists the definition of the colour sets used in the `AdHocNetwork` module. The colour set `AHNConfig` is used to model the configuration information for the mobile ad-hoc nodes. Each ad-hoc node is represented by a token on place `Nodes` and the colour of the tokens specifies the name of the node and a list of configured IP addresses. Each configuration specifies the IP address configured, and the IP address and prefix of the corresponding gateway. It is possible for




Fig. 24. Core Network module – modelling the core network




Fig. 25. AdHocNetwork module – modelling the ad-hoc network

```
(* --- ad-hoc nodes --- *)
color AHId = int with 1..5;
color AHNNode = union AHN : AHId;

(* --- configuration information for ad-hoc nodes --- *)
color AHNIPConfig = product IPAdr * IPAdr * Prefix;
color AHNIPConfigs = list AHNIPConfig;
color AHNConfig = product AHNNode * AHNIPConfigs;
```

Fig. 26. Colour definitions used in the AdHocNetwork module

a mobile ad-hoc node to configure an IP address for multiple gateways. The mobile node must ensure that the DNS database always contains the IP address corresponding to the *preferred gateway*. In the marking shown in Fig. 25, it can be seen from the labels below the mobile nodes that Ad-hoc Node 3 and Ad-hoc Node 4 have configured IP addresses based on the prefix announced by Gateway1, whereas Ad-hoc Node 5 has configured an IP address based on the prefix announced by Gateway2. For an example, Ad-hoc Node 3 has configured the address 3ffe:100:3:405::3.

There are four substitution transitions in the AdHocNetwork module corresponding to the components of the ad-hoc network. The substitution transition AHNnodes represents the behaviour of the nodes in the mobile ad-hoc network. The substitution transition Mobility models the mobility of nodes in the ad-hoc network, i.e., that the nodes may move closer or further away from the gateways. The substitution transition Routing represents the routing protocol executed in the ad-hoc network. The purpose of the routing protocol in the context of the RIP protocol is to provide the nodes with information about distances to the

gateways. The routing is abstractly modelled in a similar way as the routing mechanism in the core network and will not be discussed further in this paper. The substitution transition **Flooding** models the dissemination of advertisements from the gateways. A detailed presentation of this part of the model has been omitted here. The complete CPN model of the RIP protocol is hierarchically structured into 18 modules. A detailed presentation of the CPN model can be found in [74].

5.2 Behavioural Visualisation of the RIP Protocol

In the routing interoperability project, the BRITNeY Suite animation framework [111] was used to create an *animation GUI* on top of the CPN model. The animation GUI allows a user to observe the execution of the constructed CPN model using a graphical representation of the network architecture. The graphics is updated by the underlying CPN model according to the execution of the formally specified protocol, and the CPN model is also able to react to stimuli provided by the user via the animation GUI.

Figure 27 shows a representative snapshot of the application-specific graphics during the execution of the CPN model. The IP addresses configured by the individual nodes are shown as labels below the nodes. For an example, **Ad-hoc Node 3** has configured two IP addresses: `3ffe:100:3:405:3` and `3ffe:100:3:406:3`. The convention is that the preferred IP address is the topmost address in the list below the node. The entries in the DNS database are shown in the upper left corner. It shows the entries for each of the three ad-hoc nodes. The two numbers written at the top of each node are counters that provide information about the number of packets on the incoming (left) and outgoing (right) interfaces of the nodes. Transmissions of advertisements from the gateways are visualised by green dots. Fig. 27 shows an example where **Gateway2** is transmitting an advertisement. Transmission of payload packets is visualised using red dots, and DNS packets are visualised using blue dots.

In addition to observing feedback on the execution of the CPN model in the animation GUI, it is also possible to provide input to the CPN model directly via the animation GUI. The user can move the nodes in the ad-hoc network thereby changing the distances to the two gateways. It is also possible to define a packet flow from the host in the core network to one of the nodes in the mobile ad-hoc network by clicking on the red square positioned next to each of the ad-hoc nodes. The square will change its colour to green once the CPN model has registered the flow. The flow can be stopped again by clicking on the (now green) square next to the mobile ad-hoc node. Finally, it is possible to force the transmission of an advertisement from a gateway by clicking on the gateway.

A more generic form of high-level graphical feedback in the form of MSCs was also used in this project. Figure 28 shows an example of an MSC diagram based on a simulation of the CPN model. The MSC shows a scenario where **Ad-hoc Node 3** makes a **Move** and discovers that **Gateway 2** is now the closest gateway. This causes it to send a **DNS update** to the DNS server. The last part of the MSC shows the host initiating a packet flow to **Ad-hoc Node 3**. One benefit of

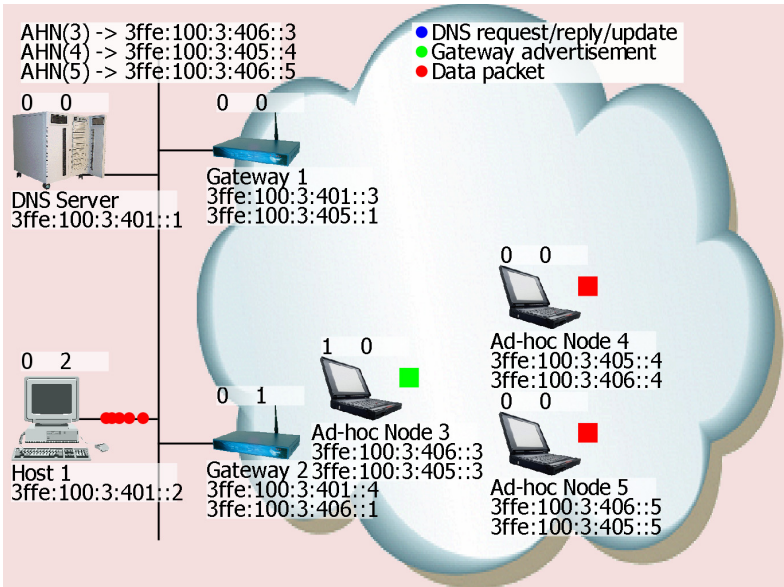


Fig. 27. Snapshot of the interaction graphics

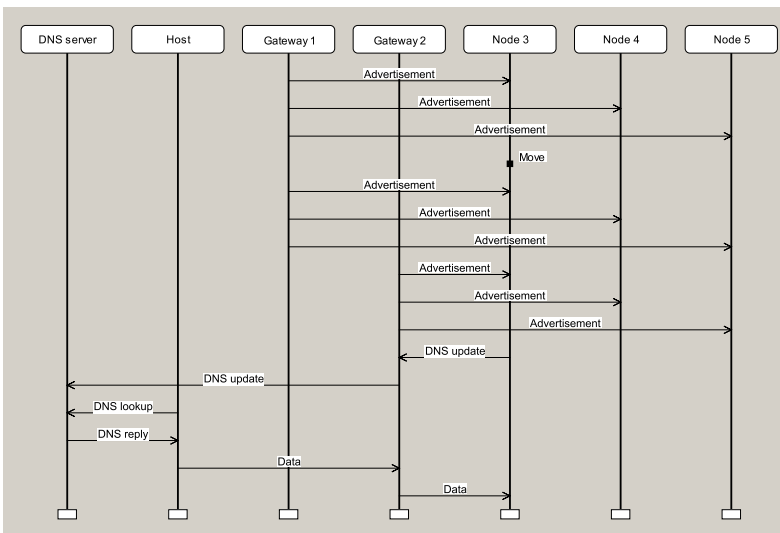


Fig. 28. Message sequence chart generated by the animation GUI

using MSCs is that they provide an event-based view that records the execution history. This is in contrast to the state-based view on the CPN model that one obtains during an interactive simulation. The two forms of feedback therefore complement each other and MSCs have been widely used in projects where CPNs were applied to protocol design.

Graphical feedback from the execution of the CPN model is achieved by attaching *code segments* to the transitions in the CPN model. These code segments are sequential pieces of code that are executed whenever the corresponding transition occurs in the simulation/execution of the CPN model. As an example consider the CNRouting module in Fig. 29. The transition *Route* models the routing of the packet on the core network. It uses the routing information on place *RoutingInformation* to direct the packet to the proper gateway. The SML function *FindNextHop* in the guard expression of the transition computes the IP address of the next hop gateway using the routing information and destination IP address of the packet. The *Route* transition has an attached code segment which is executed whenever the transition occurs. The code segment invokes the primitives in the animation package for animating the transmission of packets in the core network.

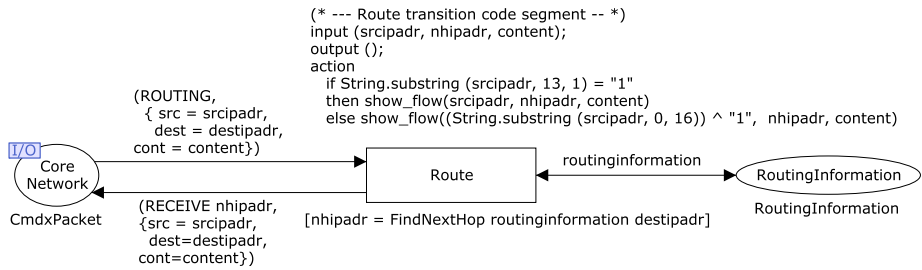


Fig. 29. The CNRouting module – Routing in the core network

The CPN model receives input from the animation GUI by polling the animation GUI for events. An event queue has been implemented between the animation GUI and the CPN model. The code segment of transition *Produce* in the *Poll* module shown in Fig. 30 polls the animation GUI for events at regular intervals during the execution of the CPN model. Events are put into a list-token representing an event queue on the place *Events*. The parts of the CPN model that are to react on events from the animation GUI are linked via place fusion to the *Event* place and are able to consume events from the event queue. The occurrence of the transition *Produce* corresponds to a poll to the animation GUI for events.

5.3 Lessons Learned and Perspectives

The CPN model combined with the animation GUI that was developed in the RIP project served as an early model-based executable prototype. The domain




Fig. 30. The Poll module – Polling the animation GUI for events

specific graphical user interface (the animation GUI) made it possible to explore and demonstrate the design of the interoperability protocol with the underlying formal model being transparent for the observer and the demonstrator. In particular, it made it possible for persons without knowledge of the CPN modelling language to experiment with the proposed design. The use of an animation GUI on top of the CPN model has the advantage that the behaviour observed by the user is as defined by the underlying model that formally specifies the design. The alternative would have been to implement a separate visualisation application totally detached from the CPN model. This would have led to double representation of the dynamics of the interoperability protocol which could in turn lead to inconsistencies between the two representation of the design.

Another advantage offered by the development of a model-based prototype is ease of control compared to a physical prototype, in particular in the case of mobile nodes and wireless communication where scenarios can be very difficult to control and reproduce. The use of a model means that there is no need to invest in physical equipment and there is no need to setup the actual physical equipment early in the project. The use of a model also makes it possible to investigate larger scenarios, e.g., scenarios that may not be feasible to investigate with the available physical equipment. An additional general advantage of the approach taken in the RIP project is that at an early stage of development, the implementation details can be abstracted away and only the key part of the design have to be specified in detail. As an example, the CPN model of the interoperability protocol abstracted away the routing mechanisms in the core and ad-hoc networks, and the mechanism used for distribution of advertisements. Instead, the service assumed from these components for the interoperability protocol to work was modelled. The possibility of making abstraction means that it is possible to obtain an executable prototype without implementing all the components.

6 The Edge Router Discovery Protocol

The previous sections have demonstrated how modelling, simulation, state space exploration, and behavioural visualisation can be applied for validating the functional design of protocols. This section summarises a project [67] conducted with Ericsson Telebit A/S where a combination of the techniques introduced in the previous sections were applied for the design of an Edge Router Discovery Protocol (ERDP). The CPN model of ERDP was developed in close cooperation

with the protocol engineers at Ericsson Telebit A/S based on a natural-language specification that would normally have served as a basis for the implementation of the protocol. Simulation and MSCs were used in initial investigations of the ERDP protocol behaviour. Then state space exploration was used to conduct a formal verification of the key behavioural properties of ERDP. The aim of this section is to show how modelling, simulation, visualisation, and state space exploration all can help to identify omissions and behavioural errors in a design, and how they are typically used in conjunction in a protocol design process.

6.1 CPN Model of the ERDP Protocol

ERDP is based on the IPv6 Neighbour Discovery Protocol (NDP) [88] and supports *edge routers* residing on the boundary of an *IP core network* in configuring *gateways* with an IPv6 address prefix. This address prefix can in turn be used by mobile nodes in ad-hoc networks to configure global IPv6 unicast addresses and obtain Internet access via the core network. Figure 31 shows the ERDP module which is the top-level module of the CPN model. The substitution transition *Gateway* represents the gateway, and the substitution transition *EdgeRouter* represents the edge router. The wireless communication link between the edge router and the gateway is represented by the substitution transition *GW_ER_Link*. The four socket places *GWIn*, *GWOut*, *ERIn*, and *EROut* model packet buffers between the link layer and the gateway and edge router. Both the gateway (GW) and the edge router (ER) have an incoming and an outgoing packet buffer.

All four places in Fig. 31 have the colour set *IPv6Packet*, used to model the IPv6 packets exchanged between the edge routers and gateways. Since ERDP is based on the IPv6 Neighbour Discovery Protocol, the packets are carried as Internet Control Message Protocol (ICMP) packets. The definitions of the colour sets for NDP, ICMP, and IPv6 packets were derived directly from RFC 2460 [22] by using record type constructors for representing fields within packets

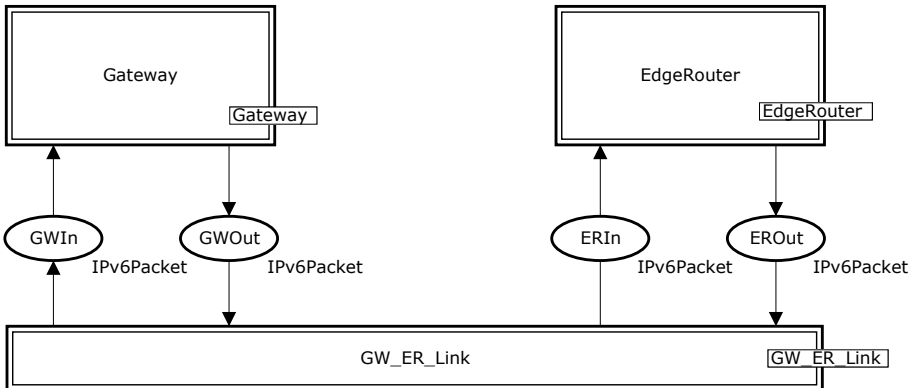


Fig. 31. Top-level module of the ERDP CPN model

and union type constructors for representing the different kinds of packets (see [67] for detail). It was considered important by the protocol engineers for later implementation that the definition of the packets followed closely the structure of IPv6 packets instead of a more abstract representation.

Figure 32 shows the `EdgeRouter` module. The port places `ERIn` and `EROut` are related to the accordingly named socket places in the `ERDP` module (see Fig. 31). The place `Config` models the configuration information associated with the edge router, and the place `PrefixCount` models the number of prefixes still available in the edge router for distribution to gateways. The place `PrefixAssigned` is used to keep track of which prefixes are assigned to which gateways.

Figure 33 shows the declarations of the colour sets for the three places in Fig. 32. The configuration information for the edge router (modelled by the colour set `ERConfig`) is a record consisting of the IPv6 link-local address and the link-layer address of the edge router. A list of pairs (colour set `ERPrefixAssigned`) consisting of a link-local address and a prefix is used to keep track of which prefixes are assigned to which gateways. A counter modelled by the place `PrefixCount` with the colour set `PrefixCount` is used to keep track of the number of prefixes still available. When this counter reaches 0, the edge router has no further prefixes available for distribution. The number of available prefixes can be modified by changing the initial marking of the place `PrefixCount`, which is set to 1 by default.

The substitution transition `SendUnsolicitedRA` (in Fig. 32) corresponds to the multicasting of periodic *unsolicited router advertisements* (RAs) by the edge router such that gateways can discover the presence of the edge router. When a gateway receives an unsolicited RA, it responds with a unicast *router solicitation* (RS). The substitution transition `ProcessRS` models the reception at the edge

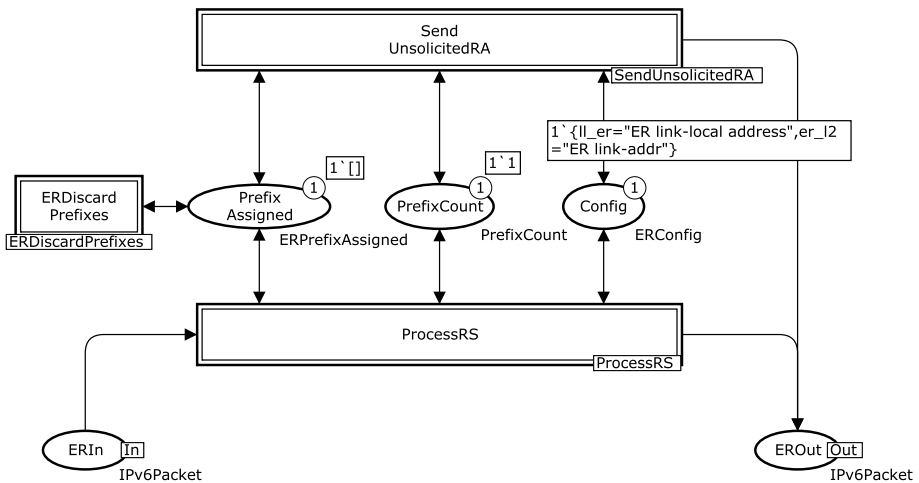


Fig. 32. The `EdgeRouter` module


```

colset LinkAddr      = string;

colset ERConfig = record ll_er : IPv6Addr * (* link-local address *)
                        er_l2 : LinkAddr; (* link-addr (layer 2) *)

colset ERPrefixEntry = product IPv6Addr * IPv6Prefix;
colset ERPrefixAssigned = list ERPrefixEntry;

colset PrefixCount = int;
    
```

Fig. 33. Colour set definitions for edge routers

router of unicasted RSs from gateways, and the sending of a unicast RA to the gateway in response. The substitution transition `ERDiscardPrefixes` models the expiration of prefixes on the edge router side.

The marking shown in Fig. 32 has a single token on each of the three places used to model the internal state of the edge router protocol entity. In the marking shown, the token on the place `PrefixAssigned` with the colour `[]` corresponds to the edge router not having assigned any prefixes to the gateways. The token on the place `PrefixCount` with colour `1` indicates that the edge router has a single prefix available for distribution. Finally, the colour of the token on the place `Config` specifies the link-local and link addresses of the edge router. In this case the edge router has the symbolic link-local address `ER link-local address`, and the symbolic link-address `ER link-addr`.

Figure 34 depicts the `SendUnsolicitedRA` module which is the submodule of the substitution transition `SendUnsolicitedRA` in Fig. 32. The transition `SendUnsolicitedRA` models the sending of the periodic unsolicited router advertisements. The variable `erconfig` is of type `ERConfig`, and the variable `prefixleft` is

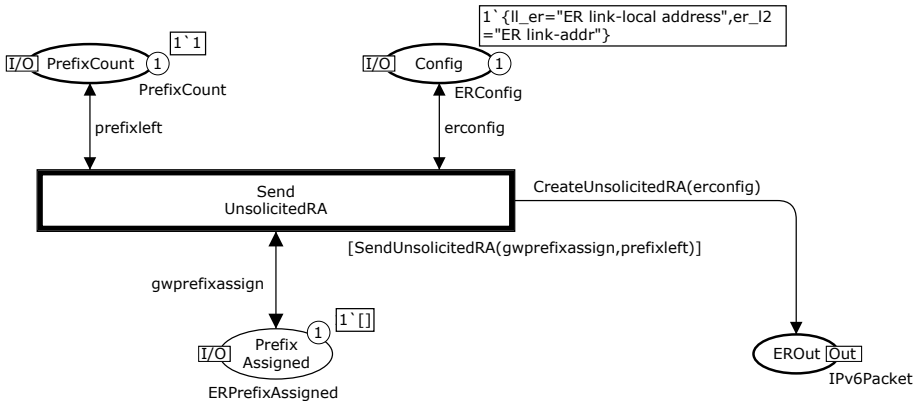


Fig. 34. Initial marking of the `SendUnsolicitedRA` module

of type `PrefixCount`. The transition `SendUnsolicitedRA` is enabled only if the edge router has prefixes available for distribution, i.e., `prefixleft` is greater than 0. This is ensured by the function `SendUnsolicitedRA` in the guard of the transition.

Figure 35 depicts the marking of the `SendUnsolicitedRA` module after the occurrence of the transition `SendUnsolicitedRA` in the marking shown in Fig. 34. An unsolicited router advertisement has been put in the outgoing buffer of the edge router. It can be seen that the `DestinationAddress` is the address `all-nodes-multicast`, the `SourceAddress` is ER link-local address, and the `LinkLayerAddress` (in the options part) is ER link-addr.

Figure 36 shows the part of the `GW_ER_Link` module that models transmission of packets from the edge router to the gateway across the wireless link. Transmission of packets from the gateway to the edge router is modelled similarly. The port places `GWIn` and `EROut` are linked to the similarly named socket places in Fig. 31. The transition `ERtoGW` models the successful transmission of packets, whereas the transition `LossERtoGW` models the loss of packets. The variable `ipv6packet` is of type `IPv6Packet`. A successful transmission of a packet from the edge router to the gateway corresponds to moving the token modelling the packet from the place `EROut` to `GWIn`. If the packet is lost, the token will only be removed from the place `EROut`.

Wireless links, in general, have a lower bandwidth and higher error rate than wired links. These characteristics have been abstracted away in the CPN model since the purpose is to reason not about the performance of ERDP but rather its logical correctness. Duplication and reordering of messages are not possible on typical one-hop wireless links, since the detection of duplicates and the preservation of order are handled by the data-link layer. The modelling of the wireless links does allow overtaking of packets, but this overtaking was elimi-

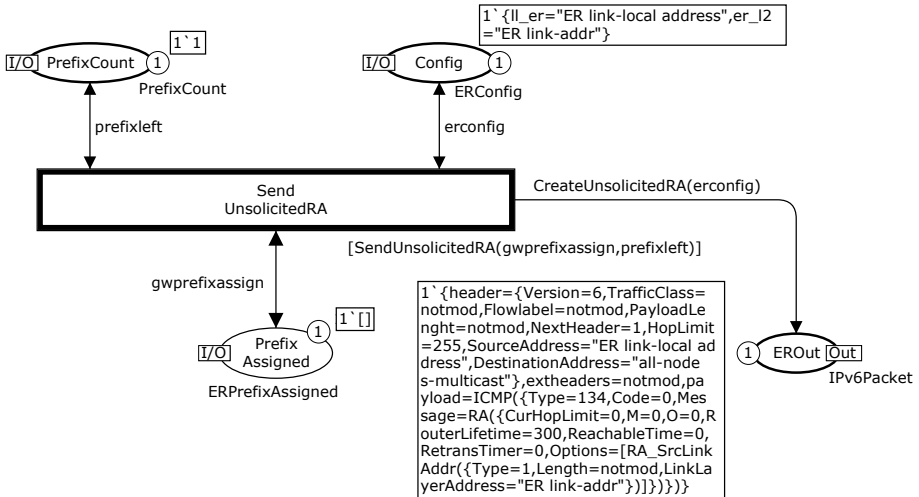


Fig. 35. Module `SendUnsolicitedRA`, after occurrence of `SendUnsolicitedRA`




Fig. 36. Part of the GW_ER.Link module

Table 2. ERDP project [67] – design issues identified in the modelling phase

Category	Review 1	Review 2	Total
Errors in protocol specification/operation	2	7	9 issues
Incompleteness and ambiguity in specification	3	6	9 issues
Simplifications of protocol operation	2	0	2 issues
Additions to the protocol operation	4	0	4 issues
Total	11	13	24 issues

nated in the state space exploration phase where bounds were imposed on the capacity of the input and output packet buffers.

The CPN model was developed as an integrated part of the development of ERDP. The creation of the CPN model was done in cooperation with the protocol engineers at Ericsson in parallel with the development of the ERDP specification. Altogether 70 person-hours were spent on CPN modelling. Prior to the development of the CPN model, the protocol engineers at Ericsson were given a 6 hour course on CPNs that made them capable of reading CPN models. This means that CPN models could be used actively in discussion related to the design of the ERDP protocol. MSCs (to be illustrated shortly), integrated with simulation were used in both review steps to investigate the behaviour of ERDP in detail. The use of MSCs in the project was of particular relevance since it presented the operation of the protocol in a form well known to the protocol engineers. Altogether 24 design issues were identified during three iterations on the CPN model. Table 2 categorises and enumerates the issues encountered during two review phases (Review 1 and Review 2) of the protocol design. The issues were identified in the process of constructing the CPN model, performing single-step executions of the CPN model, and engaging in discussions of the CPN model with the protocol engineers at Ericsson.

6.2 Verification of the ERDP CPN Model

State space exploration was conducted after the three iterations of modelling as discussed in the previous section. The purpose of the state space exploration was to conduct a more thorough investigation of the operation of ERDP, including verification of its key properties. The key behavioural property of ERDP is

proper configuration of the gateway with prefixes. This means that for a given prefix and state where the gateway has not yet been configured with that prefix, the protocol must be able to configure the gateway with that prefix. Furthermore, when the gateway has been configured with the prefix, the edge router and the gateway should be *consistently configured*, i.e., the assignment of the prefix must be recorded both in the gateway and in the edge router protocol entity. Whether a marking represents a consistently configured state for a given prefix can be checked by inspecting the marking of the place `PrefixAssigned` in the edge router and the marking of the place `Prefixes` in the gateway.

Obtaining a finite state space. The first step towards state space exploration of the CPN model was to obtain a finite state space. The CPN model as presented above has an infinite state space, since an arbitrary number of tokens (packets) can be put on the places modelling the packet buffers. As an example, the edge router may initially send an arbitrary number of unsolicited router advertisements. To obtain a finite state space, an upper integer bound of 1 was imposed on each of the places `GWIn`, `GWOut`, `ERIn`, and `EROut` (see Fig. 31) which model the packet buffers. This also prevents overtaking among the packets transmitted across the wireless link. Furthermore, the number of packets simultaneously present in the four input/output buffers was limited to 2. Technically, this was done by using the *branching options* available in the CPN state space tool to prevent the processing of enabled transitions whose occurrence in a given marking would violate the imposed bounds on the buffer places.

No packet loss and prefix expire. The second step was to consider the simplest possible configurations of ERDP, starting with a single prefix and assuming that there is no packet loss on the wireless link and that prefixes do not expire. The full state space for this configuration had 46 nodes and 65 arcs. Inspection of the state space report showed that there was a single dead marking represented by node 36. Inspection of this node showed that it represented a state where all of the packet buffers were empty. However, the edge router and gateway were inconsistently configured in this state in that the edge router had assigned the prefix `P1` (the single prefix), while the gateway was not configured with that prefix. This was an error in the protocol. To locate the source of the problem, query functions in the state space tool were used to obtain a counter example leading from the node representing the initial marking to node 36. Figure 37 shows the resulting error trace, visualised by means of an MSC. This MSC was generated automatically from the extracted counter example. The column labelled `GW-Buffer` represents the packet buffer between the gateway protocol entity and the underlying protocol layers. Similarly, the `ERBuffer` column represents the packet buffer in the edge router. The problem is that the edge router sends two unsolicited RAs. The first one gets through and the gateway is configured with the prefix, which can be seen from the event marked with `*A*` in the lower part of the MSC. However, when the second RS, without any prefixes, is received by the edge router (the event marked with `*B*`), the corresponding solicited RA will not contain any prefixes. Because of the way the protocol was specified, the

gateway will therefore update its list of prefixes to the empty list (the event marked with *C*), and the gateway is no longer configured with a prefix.

To fix the error, the protocol was modified so that the edge router always replies with the list of all prefixes that it has currently assigned to the gateway. The state space for the modified protocol consisted of 34 nodes and 49 arcs, and there were no dead markings in the state space. The state space report specified that there were 11 home markings. Inspection of these 11 markings showed that they all represented consistently configured states for the prefix P1. The markings were contained in the single terminal SCC of the state space. A terminal SCC is an SCC of the state space where all successors of states in the SCC belong to the SCC itself. This shows that, from the initial marking it is always possible to reach a consistently configured state for the prefix, and that when such a marking has been reached, the protocol entities will remain in a consistently configured state. To verify that a consistently configured state would eventually be reached, it was checked that the single terminal SCC was the only non-trivial SCC. A trivial SCC is a SCC consisting of just a single state. This showed that all cycles in the state space (which correspond to non-terminating executions of the protocol) were contained in the single terminal SCC, which (from above) contained only consistently configured states. The reason why the protocol is not supposed to terminate in a consistently configured state represented by a dead marking is that the gateway may, at any time, when it is configured, send a router solicitation back to the edge router to have its prefixes refreshed.

Increasing the number of prefixes. When the correctness of the protocol had been established for a single prefix, the number of prefixes was increased. When there is more than one prefix available it no longer holds that a marking will *eventually* be reached where *all* prefixes are consistently configured. The reason is that with more than one prefix, the edge router may at any time decide not to configure the gateway with additional prefixes. Hence, a state where all prefixes have been consistently configured might not eventually be reached. Instead, firstly, it was verified that there was a single terminal SCC, all markings of which represent states where all prefixes have been consistently configured. This shows that it is always possible to reach such a marking, and when the protocol has consistently configured all prefixes, the protocol entities will remain consistently configured. Secondly, it was checked that all markings in each non-trivial SCC represented markings where the protocol entities were consistently configured with a subset of the prefixes available in the edge router. The properties above was checked using a number of user-defined queries in the state space tool of CPN Tools.

Adding packet loss. The third step was to allow packet loss on the wireless link between the edge router and the gateway. First, the case was considered in which there is only a single prefix for distribution. The state space for this configuration had 40 nodes and 81 arcs. Inspection of the state space report showed that there was a single dead marking. This marking represented an undesired terminal state where the prefix had been assigned by the edge router, but the gateway




Fig. 37. MSC showing an execution leading to an undesired terminal state

was not configured with the prefix. The source of the problem was located by extracting a counter example and visualising it in a similar manner as shown in Fig. 37. The problem was fixed by ensuring that the edge router would resend an unsolicited RA to the gateway as long as it had prefixes assigned to the gateway. The state space of the revised CPN model had 68 nodes and 160 arcs. Inspection of the state space report showed that there were no dead markings and no home markings. Investigation of the terminal SCCs showed that there were two terminal SCCs, each containing 20 markings. The nodes in one of them all represented states where the edge router and gateway were consistently configured with the single prefix P1, whereas the nodes in the other terminal SCC all represented states where the protocol entities were not consistently configured. The markings in the undesired terminal SCC represent a livelock in the protocol, i.e., if one of the markings in the undesired terminal SCC is reached, it is no longer possible to reach a state where the protocol entities are consistently configured with the prefix. The source of the livelock was related to the control fields used in the router advertisements for refreshing prefixes and their interpretation on the gateway. This was identified by obtaining the MSC for a path leading from the initial marking to one of the markings in the undesired terminal SCC. As a result, the processing of router advertisements in the gateway was modified. The state space for the protocol with the modified processing of router advertisements also had 68 nodes and 160 arcs. The state space had a single terminal SCC containing 20 nodes, which all represented states where the protocol entities were consistently configured with the single prefix.

When packet loss is present, it is not immediately possible to verify that the two protocol entities will eventually be consistently configured. The reason is that any number of packets can be lost on the wireless link. Each of the non-trivial SCCs was inspected using a user-defined query to investigate the circumstances under which the protocol entities would not eventually be consistently configured. This query checked that either all nodes in the non-trivial SCC represented consistently configured states or none of the nodes in the SCC represented a consistently configured state. For those non-trivial SCCs where no node represented a consistently configured state, it was checked that all cycles contained the occurrence of a transition corresponding to loss of a packet. Since this was the case, it can be concluded that any failure to reach a consistently configured states will be due to packet loss only. Hence, if finitely many packets are lost, a consistently configured state for some prefix will *eventually* be reached.

Adding prefix expire. The fourth and final step in the analysis was to allow prefixes to expire. The analysis was conducted first for a configuration where the edge router had only a single prefix to distribute. The state space for this configuration had 173 nodes and 513 arcs. The state space had a single dead marking, and inspection of this dead marking showed that it represented a state where the edge router has no further prefixes to distribute, it has no prefixes recorded for the gateway, and the gateway is not configured with any prefix.

This marking is a desired terminating state of the protocol, as we expect prefixes to eventually expire. Since the edge router has only finitely many prefixes to distribute, the protocol should eventually terminate in such a state. The single dead marking was also a home marking, meaning that the protocol can always enter the expected terminal state. When prefixes can expire, it is possible that the two protocol entities may never enter a consistently configured state. The reason is that a prefix may expire in the edge router (although this is unlikely) before the gateway has been successfully configured with that prefix. Hence, we are only able to prove that for any marking where a prefix is still available in the edge router, it is possible to reach a marking where the gateway and the edge router are consistently configured with that prefix.

Table 3 lists statistics for the size of the state space in the three verification steps for different numbers of prefixes. The column ‘|P|’ specifies the number of prefixes. The columns ‘Nodes’ and ‘Arcs’ give the numbers of nodes and arcs, respectively, in the state space. For the state spaces obtained in the first verification step, it can be seen that 38 markings and 72 arcs are added for each additional prefix. The reason for this is that ERDP proceeds in phases where the edge router assigns prefixes to the gateway one at a time. Configuring the gateway with an additional prefix follows exactly the same procedure as that for the assignment of the first prefix. Once the state space had been generated, the verification of properties could be done in a few seconds. It is also worth observing that as the assumptions are relaxed, i.e., we move from one verification step to the next, the sizes of the state spaces grow. This, combined with the identification of errors in the protocol even in the simplest configuration, without packet loss and without expiration of prefixes, shows the benefit of starting state space exploration from the simplest configuration and gradually lifting assumptions. Furthermore, the state explosion problem was not encountered during the verification of ERDP, and the key properties of ERDP were verified for the number of prefixes that were envisioned to appear in practise.

6.3 Lessons Learned and Perspectives

The project at Ericsson highlights the benefits of a formal modelling and validation approach. Furthermore, the project emphasised the benefits of the model construction phase which is often underestimated (or not reported) in literature on protocol validation. As illustrated by the ERDP project, the modelling phase itself lead to significant insight into the protocol design, and contributed to a simpler and more complete protocol design. The construction of a CPN model and subsequent state space exploration can be seen as a very thorough and systematic way of reviewing the ERDP design specification. The project showed that the process of constructing a CPN model based on the ERDP specification provided valuable input to the ERDP design, and the use of simulation added further insight into the operation of the protocol. State space exploration, starting with the simplest possible configuration of the protocol, identified additional errors in the protocol. The results from state space exploration also

Table 3. State space statistics for the three verification steps

P	No loss/No expire		Loss/No Expire		Loss/Expire	
	Nodes	Arcs	Nodes	Arcs	Nodes	Arcs
1	34	49	68	160	173	531
2	72	121	172	425	714	2 404
3	110	193	337	851	2 147	7 562
4	148	265	582	1 489	5 390	19 516
5	186	337	926	2 390	11 907	43 976
6	224	409	1 388	3 605	23 905	89 654
7	262	481	1 987	5 185	44 550	169 169
8	300	553	2 742	7 181	78 211	300 072
9	338	625	3 672	9 644	130 732	505 992
10	376	697	4 796	12 625	209 732	817 903

demonstrate that errors are often present in the smallest configurations of a protocol system.

Using an iterative process where both a conventional natural-language specification and a CPN model were developed (as in this project) turned out to be an effective way of integrating CPN modelling and validation into the development of a protocol. In general, the combination of an executable formal model (such as a CPN model) and a natural-language specification seems to provide a useful way to develop a protocol. One reason why both are required is that the software engineers that are eventually going to implement the protocol (which may be different from those that design the protocol) in many cases will not be familiar with the CPN modelling language. Secondly, in many cases there are important implementation elements of the protocol specification that are not reflected in the CPN model, such as the layout of packets.

It can be argued whether or not the issues and errors discovered in the process of modelling and conducting state space exploration would have been identified if additional conventional reviews of the ERDP specification had been conducted. Some of them probably would have been, but more subtle problems such as the inconsistent configurations discovered during state space exploration would probably not have been discovered until the first implementation of ERDP was operational. The reason for this is that discovering these problems requires one to consider subtle execution sequences of the protocol.

Overall, the application of CPNs in the development of ERDP was considered a success for three main reasons. Firstly, it was demonstrated that the CPN modelling language and supporting computer tools were powerful enough to specify and verify a real-world protocol being developed in an industrial project, and that integration into the conventional protocol development process is not difficult. Secondly, the act of constructing the CPN model, executing it, and discussing it led to the identification of several non-trivial design errors and

issues that, under normal circumstances, would not have been discovered until, at best, the implementation phase. Finally, the effort of constructing the CPN model and conducting state space exploration was represented by approximately 100 person-hours. This is a relatively small investment compared with the many issues that were identified and resolved early as a consequence of constructing and analysing the CPN model.

7 Related Work on CPN Protocol Validation

The four protocol examples presented in this paper constitute only a small subset of the examples that have been published in the literature on the use of CPNs for specification and validation of protocols - in particular in relation to protocols developed in the context of IETF and other protocol standardisation bodies.

The Datagram Congestion Control Protocol (DCCP) developed by the IETF has been investigated in [11]. DCCP is intended to provide an unreliable transport service with congestion control mechanisms. The work in [11] was done in parallel with the development of the emerging DCCP standard, and concentrated on modelling and verification of the connection establishment and synchronisation procedures of DCCP. It resulted in the identification of several functional errors in the protocol design, including discovery of deadlocks, non-progress behaviour (chatter), and problems with connection establishment in relation to sequence number wraps. The formal validation resulted in the IETF working group making small (but important) changes to the connection establishment and synchronisation procedures of DCCP. The work also included the development of a formal service specification for DCCP [33] and application of the sweep-line method [105] for on-the-fly checking of the protocol conformance to the developed service specification.

The classical Transmission Control Protocol (TCP) has also been modelled and verified using CPNs [10]. Similar to the work on DCCP, this work concentrated on the connection establishment procedures. It resulted in verifying the absence of deadlocks and livelocks in connection establishment, and a detailed specification of the circumstances under which TCP connection establishment may not be successful. Another example of transport layer protocol modelling and validation can be found in [104] which considers the Stream Transmission Control Protocol (SCTP).

The Internet Open Trading Protocol (IOTP) designed to provide an interoperability framework for Internet commerce was formally modelled and validated using CPNs in [90–92]. IOTP is designed to handle common trading procedures and encompass trading roles such as consumer, merchant, payment handler, and delivery handler. IOTP is organised around a collection of eight baseline transactions consisting of Purchase, Refund, Value exchange, Authentication, Withdrawal, Deposit, Inquiry, and Ping. These transactions comprise a minimal set of transactions for an Internet commerce protocol. A formal specification of the service provided by IOTP was developed using CPN in [92]. The service

was specified in the form of a finite-state automaton labelled with service primitives. The automaton was extracted from the state space of the CPN model by identifying the binding elements corresponding to service primitives of the protocol. A CPN model of the IOTP protocol itself was presented in [90, 91]. State space exploration focused on the termination properties and absence of livelocks in the IOTP transactions. The use of state space exploration revealed deficiencies related to termination of transactions. A verification of the IOTP protocol CPN model [90, 91] against the formal service specification from [92] was presented in [89]. Finite-state automata language comparison was used as the criterion for conformance following the methodology of [9]. Application of the sweep-line state space method on IOTP was investigated in [34] exploiting an inherent progression from the start of an IOTP transaction to termination of the transaction.

The Wireless Application Protocol (WAP) has been considered in [40, 41]. WAP is designed to provide Internet services to a wide range of hand-held devices. The work of [40, 41] concentrates on the Wireless Transaction Protocol (WTP) which is an important element of the WAP architecture and protocol suite. The work in [40] presents a formal modelling of the WTP service and a formal modelling of the WTP protocol. Checking the conformance of the WTP protocol against the WTP service was done using finite-state automata language comparison. This approach succeeded in detecting several inconsistencies between the protocol and the service which was provided as input to the WAP forum responsible for the development of WAP. The sweep-line method was used in [41] to alleviate the state explosion problem and allow for the verification of larger configurations of WTP. The application of the sweep-line method allowed configurations with parameter settings of retransmission counters corresponding to the recommended setting for GSM and IP network to be verified.

The Session Initiation Protocol (SIP) is a widely used protocol for the establishment of Internet multimedia session, and has been subject to formal modelling and validation in [23, 77]. The INVITE transactions have been formally analysed using state space exploration in [23, 77] leading to identification of undesired terminating states of the protocol when operating over an unreliable communication medium. Security aspects of SIP have been investigated in [78]. The work of [37] focuses on the formal modelling of a SIP-based protocol for multi-channel service oriented architectures. A formalisation of SIP with the purpose of providing a framework model for present architectures in mobile computing is presented in [36]. Another multimedia control protocol, the Capability Exchange Signalling (CES) protocol, has been formally modelled using CPNs and verified using state space exploration in [79]. The work on the CES protocol led to the identification of protocol errors in presence of sequence number wrap. Suggested changes were incorporated in a revised CPN model, and it was formally verified showing that the discovered errors have been eliminated.

The NEO protocol which is part of the distributed transactional object database management system NEOPPOD was investigated using high-level Petri

Nets in [17]. The Coloane environment was used for the construction of the models, and verification was performed using the CPN-AMI and Helena tools. The NEO protocol is used to coordinate data storage and retrieval in a decentralised and distributed system where data can be stored on a number of data nodes and data is accessed through the primary master node. The focus of [17] was on the protocol used for the election of the primary master node. The model of the election part of the NEO protocol consisted of eighteen modules. Since there existed no specification document for the protocol, the Petri net model was reverse-engineered from a prototype implementation. The validation process which relied on the use of state spaces discovered two flaws in the implementation of the protocol. These were subsequently provided to the software engineers responsible for the implementation of the protocol component.

The Resource Reservation Protocol (RSVP) was formally modelled and verified in [106, 107]. The modelling and verification concentrates on verifying the absence of deadlocks and livelocks in relation to the setup, maintenance and path release procedures of RSVP. In addition, a number of RSVP specific behavioural properties were investigated which considered in detail the internal state of the sender, router, and receiver protocol entities of the protocol. The main contribution of the work was the development of a formal specification of the RSVP path procedures. Another example on the modelling of routing protocols can be found in [76] which uses Mobile Petri Nets to construct a formal model of the Mobile IP protocol. Mobile IP allows transport layer connections to be preserved when mobile nodes change their point of attachment to the Internet. CPNs have also recently been used for the verification of security protocols. Privacy enhancing protocols were considered in [99], and [39] addresses the modelling and validation of PANA Authentication and Authorisation Protocol. Examples of protocols for which parametric verification has been pursued in the context of CPNs can be found in [31, 32].

8 Conclusions and Outlook

Functional validation of protocol designs is one of the main application areas of CPNs and supporting computer tools [28]. In this paper, we have surveyed a selection of recent projects on modelling and functional validation of industry relevant protocols. The examples demonstrate how the elements of protocols can be modelled using CPNs, and they illustrate how a combination of simulation, application-specific behavioural visualisation, and state space exploration is typically applied in protocol validation with CPNs. From a modelling perspective, the protocol examples have ranged from models representing two (or few) peer protocol entities (e.g., GAN, EDRP, and RIP) having an explicit representation in the net structure, to parameterised models capable of modelling an arbitrary number of peer protocol entities by setting a model parameter (e.g., DYMO). The latter was based on constructing a folded model where the identity of the protocol entities is encoded explicitly as part of the token colours. The CPN

models have also illustrated modelling at different protocols layer ranging from models operating at a single protocol layer (e.g., DYMO and ERDP) to protocol system design involving multiple protocol layers and protocols (e.g., GAN and RIP). An important aspect of the examples is that the process of modelling and conducting single step simulation is an important (but often underestimated) activity in the validation of a protocol design.

The main technique available for functional verification of CPN models is that of explicit state space exploration. The examples presented in this paper show how basic state space exploration combined with the generation of a state space report relying on a number of standard behavioural properties of Petri Nets, provides a light-weight approach which in many cases is an important step in verifying key properties of a protocol design. The main reason for the wide spread application of state space exploration has been the presence of mature computer tool support combined with the main advantages of state space exploration in terms of being a highly systematic approach, being able to provide counter examples, and allowing for a high degree of automation. The compact modelling of protocols enabled by CPNs has, in many cases, had the effect that the full state space can be explored for at least the smallest configuration of the considered protocol. The GAN and ERDP examples presented in this paper are concrete examples illustrating this. Practise have shown that the primary capability offered by the advanced state space methods is the possibility of verifying larger configurations of the protocol - and in some cases [71] the configurations of the system that are expected to occur in practise. The ERDP example considered in this paper is another example of this. Hence, despite the fact that explicit state space exploration methods requires one to conduct verification relative to a particular configuration of the protocol, the current suite of available state space methods combined with the power of modern computing platforms in many situations allows for the practical validation of industrial-sized protocols.

While CPNs have been successfully applied to modelling and validating protocol designs, there has been relatively few attempts at using the constructed CPN models in an automated or semi-automated manner as a basis for the actual implementation of protocols. Some simulation-based approaches were used in [87] and [70] for generating server-side implementations. Here, the simulation code for the CPN model generated by CPN Tools was extracted, and after undergoing automatic modifications (e.g., linking the code to external libraries), the generated simulation code is used as the system implementation. A limitation of this approach is that the execution speed is affected because each step in the execution of the program involves the computation and execution of enabled transitions (as done by a CPN simulator) in order to determine the next state. Secondly, the approach ties the target platform to that of the CPN Tools simulator which may make the approach impractical for many application domains due to resource consumption of the CPN simulator. The SML/NJ compiler used for the simulator in CPN Tools has a large memory footprint making it ill-suited, e.g., for the domain of embedded systems. Some initial work on a

translation-based approach can be found in [73]. Here a restricted form of CPNs was used for obtaining an Erlang implementation of the DYMO routing protocol. The approach in [73] relies on the use of Process-Partitioned CPNs which enforces a detailed modelling of the protocol design which is very close to an implementation level model. An area that will be important as part of efforts in developing capabilities for automated code generation is the development of CPN protocol modelling methodology on which only limited research has been undertaken [18].

Acknowledgements. The authors acknowledge the contribution of Kristian L. Espensen and Mads. K. Kjeldsen in the project on the DYMO protocol, Paul Fleischer for his contribution in the GAN project, Michael Westergaard and Peder Christian Nørgaard for their contribution to the project on the RIP protocol, and Kurt Jensen for his contributions in the ERDP project.

References

1. ISO/IEC 15437. Information technology. Enhancements to LOTOS (E-LOTOS) (September 2001)
2. 3GPP. Digital Cellular Telecommunications System (Phase 2+); Generic Access to the A/Gb Interface; Stage 2. 3GPP TS 43.318 version 6.9.0 Release 6 (March 2007)
3. 3GPP. Website of 3GPP (May 2007), <http://www.3gpp.org>
4. Alur, R., Holzmann, G., Peled, D.: An analyzer for message sequence charts. *Software - Concepts and Tools* 17(2), 70–77 (1996)
5. Ardis, M.A.: Formal Methods for Telecommunication System Requirements: A Survey of Standardised Languages. *Annals of Software Engineering* 3 (1997)
6. Baier, C., Katoen, J.-P.: Principles of Model Checking. The MIT Press (2008)
7. Bengtsson, J., Yi, W.: Timed automata: Semantics, algorithms and tools. In: Desel, J., Reisig, W., Rozenberg, G. (eds.) *Lectures on Concurrency and Petri Nets*. LNCS, vol. 3098, pp. 87–124. Springer, Heidelberg (2004)
8. Billington, J., Gallasch, G., Kristensen, L.M., Mailund, T.: Exploiting Equivalence Reduction and the Sweep-Line Method for Detecting Terminal States. *IEEE Transactions on Systems, Man, and Cybernetics. Part A: Systems and Humans* 34(1), 23–38 (2004)
9. Billington, J., Gallasch, G.E., Han, B.: A Coloured Petri Net Approach to Protocol Verification. In: Desel, J., Reisig, W., Rozenberg, G. (eds.) *Lectures on Concurrency and Petri Nets*. LNCS, vol. 3098, pp. 210–290. Springer, Heidelberg (2004)
10. Billington, J., Han, B.: Modelling and Analysing the Functional Behaviour of TCPs Connection Management Procedures. *International Journal on Software Tools for Technology Transfer* 9(3-4), 269–304 (2007)
11. Billington, J., Vanit-Anunchai, S.: Coloured Petri Net Modelling of an Evolving Internet Protocol Standard: The Datagram Congestion Control Protocol. *Fundamenta Informaticae* 88(3), 357–385 (2008)
12. Billington, J., Yuan, C.: On Modelling and Analysing the Dynamic MANET On-Demand (DYMO) Routing Protocol. In: Jensen, K., Billington, J., Koutny, M. (eds.) *Transactions on Petri Nets and Other Models of Concurrency III*. LNCS, vol. 5800, pp. 98–126. Springer, Heidelberg (2009)

13. Bochmann, G.: Finite state description of protocols. *Computer Networks*, 361–372 (1978)
14. Bolognesi, T., Brinksma, E.: Introduction to the ISO Specification Language LOTOS. *Computer Networks* 14, 25–59 (1987)
15. Chakeres, I.D., Perkins, C.E.: Dynamic MANET On-demand (DYMO) Routing. Internet-Draft. Work in Progress (July 2007), <http://www.ietf.org/internet-drafts/draft-ietf-manet-dymo-10.txt>
16. Chakeres, I.D., Perkins, C.E.: Dynamic MANET On-demand (DYMO) Routing. Internet-Draft. Work in Progress (November 2007), <http://www.ietf.org/internet-drafts/draft-ietf-manet-dymo-11.txt>
17. Choppy, C., Dedova, A., Evangelista, S., Hong, S., Klai, K., Petrucci, L.: The NEO Protocol for Large-Scale Distributed Database Systems: Modelling and Initial Verification. In: Lilius, J., Penczek, W. (eds.) *PETRI NETS 2010*. LNCS, vol. 6128, pp. 145–164. Springer, Heidelberg (2010)
18. Choppy, C., Petrucci, L., Reggio, G.: A Modelling Approach with Coloured Petri Nets. In: Kordon, F., Vardanega, T. (eds.) *Ada-Europe 2008*. LNCS, vol. 5026, pp. 73–86. Springer, Heidelberg (2008)
19. Christensen, S., Kristensen, L.M., Mailund, T.: A Sweep-Line Method for State Space Exploration. In: Margaria, T., Yi, W. (eds.) *TACAS 2001*. LNCS, vol. 2031, pp. 450–464. Springer, Heidelberg (2001)
20. Clarke, E.M., Enders, R., Filkorn, T., Jha, S.: Exploiting Symmetries in Temporal Logic Model Checking. *Formal Methods in System Design* 9, 77–104 (1996)
21. Comer, D.E.: *Internetworking with TCP/IP vol. 1: Principles, Protocols, and Architecture*, 5th edn. Prentice-Hall (2005)
22. Deering, S., Hinden, R.: Internet Protocol, Version 6 (IPv6) Specification. RFC 2460 (December 1998)
23. Ding, L.G., Liu, L.: Modelling and Analysis of the INVITE Transaction of the Session Initiation Protocol Using Coloured Petri Nets. In: van Hee, K.M., Valk, R. (eds.) *PETRI NETS 2008*. LNCS, vol. 5062, pp. 132–151. Springer, Heidelberg (2008)
24. Emerson, E.A., Sistla, A.P.: Symmetry and Model Checking. *Formal Methods in System Design* 9, 105–131 (1996)
25. Espensen, K.L., Kjeldsen, M.K., Kristensen, L.M.: Modelling and Initial Validation of the DYMO Routing Protocol for Mobile Ad-Hoc Networks. In: van Hee, K.M., Valk, R. (eds.) *PETRI NETS 2008*. LNCS, vol. 5062, pp. 152–170. Springer, Heidelberg (2008)
26. ETSI. ETSI ES 201 873-1: Methods for Testing and Specification; The Testing and Test Control Notation version 3; Part 1: TTCN-3 Core Language
27. Evangelista, S., Westergaard, M., Kristensen, L.M.: The ComBack Method Revisited: Caching Strategies and Extension with Delayed Duplicate Detection. *Transactions on Petri Nets and Other Models of Concurrency* 3, 189–215 (2009)
28. Examples of Industrial Use of CPNs, <http://cs.au.dk/cpnets/industrial-use/>
29. Fehnker, A., van Glabbeek, R., Hofner, P., McIver, A., Portmann, M., Tan, W.: Modelling and Analysis of AODV in UPPAAL. In: *Proc. of 1st Workshop on Rigorous Protocol Engineering* (2011)
30. Fleischer, P., Kristensen, L.M.: Modelling and Validation of Secure Connection Establishment in a Generic Access Network Scenario. *Fundamenta Informaticae* 94(3-4), 361–386 (2009)
31. Gallasch, G.E., Billington, J.: Using Parametric Automata for the Verification of the Stop-and-Wait Class of Protocols. In: Peled, D.A., Tsay, Y.-K. (eds.) *ATVA 2005*. LNCS, vol. 3707, pp. 457–473. Springer, Heidelberg (2005)

32. Gallasch, G.E., Billington, J.: A Parametric State Space for the Analysis of the Infinite Class of Stop-and-Wait Protocols. In: Valmari, A. (ed.) SPIN 2006. LNCS, vol. 3925, pp. 201–218. Springer, Heidelberg (2006)
33. Gallasch, G.E., Billington, J., Vanit-Anunchai, S., Kristensen, L.M.: Checking Safety Properties On-the-fly with the Sweep-Line Method. *International Journal on Software Tools for Technology Transfer (STTT)* 9(3-4), 371–392 (2007)
34. Gallasch, G.E., Ouyang, C., Billington, J., Kristensen, L.M.: Experimenting with Progress Mappings for the Sweep-Line Analysis of the Internet Open Trading Protocol. In: Proc. of 5th Workshop and Tutorial on Practical Use of Coloured Petri Nets and CPN Tools (CPN 2004), pp. 19–38 (2004)
35. Gallasch, G.E., Han, B., Billington, J.: Sweep-Line Analysis of TCP Connection Management. In: Lau, K.-K., Banach, R. (eds.) ICFEM 2005. LNCS, vol. 3785, pp. 156–172. Springer, Heidelberg (2005)
36. Gehlot, V., Hayrapetyan, A.: A Formalized and Validated Executable Model of the SIP-based Presence Protocol for Mobile Applications. In: Proceedings of the 45th Annual ACM Southeast Regional Conference, pp. 185–190. ACM (2007)
37. Gehlot, V., Hayrapetyan, A.: A CPN Model of a SIP-Based Dynamic Discovery Protocol for Webservices in a Mobile Environment. In: Proc. of 7th Workshop and Tutorial on Practical Use of Coloured Petri Nets and the CPN Tools, CPN 2006 (2006)
38. Genest, B., Muscholl, A., Peled, D.: Message sequence charts. In: Desel, J., Reisig, W., Rozenberg, G. (eds.) Lectures on Concurrency and Petri Nets. LNCS, vol. 3098, pp. 537–558. Springer, Heidelberg (2004)
39. Gordon, S.: Formal Analysis of PANA Authentication and Authorisation Protocol. In: Proc. of 9th International Conference on Parallel and Distributed Computing, Applications and Technologies, pp. 277–284. IEEE Computer Society (2008)
40. Gordon, S., Billington, J.: Analysing the WAP Class 2 Wireless Transaction Protocol Using Coloured Petri Nets. In: Nielsen, M., Simpson, D. (eds.) ICATPN 2000. LNCS, vol. 1825, pp. 207–226. Springer, Heidelberg (2000)
41. Gordon, S., Kristensen, L.M., Billington, J.: Verification of a Revised WAP Wireless Transaction Protocol. In: Esparza, J., Lakos, C.A. (eds.) ICATPN 2002. LNCS, vol. 2360, pp. 182–202. Springer, Heidelberg (2002)
42. Grimstrup, P.: Interworking Description for IKEv2 Library. In: Ericsson Internal. Document No. 155 10-FCP 101 4328 Uen (September 2006)
43. Harel, D.: Statecharts: A Visual Formalism for Complex Systems. *Science of Computer Programming* 8(3), 231–274 (1987)
44. Hoare, C.A.R.: *Communicating Sequential Processes*. Prentice-Hall International (1985)
45. Holzmann, G.J.: *The SPIN Model Checker*. Addison-Wesley (2004)
46. Holzmann, G.J.: *Design and Validation of Computer Protocols*. Prentice-Hall (1991)
47. The Internet Engineering Task Force (IETF), <http://www.ietf.org>
48. Ip, C.N., Dill, D.L.: Better Verification Through Symmetry. *Formal Methods in System Design* 9, 41–75 (1996)
49. ISO9074. Information Processing Systems - Open Systems Interconnection: ESTELLE (FOrmal Description Technique Based on an Extended State Transition Model)
50. ISO89 ISO/IEC. Information Processing Systems - Open Systems Interconnection: LOTOS, a Formal Description Technique Based on the Temporal Ordering of Observational Behaviour. IS 8807 (February 1989)

51. ITU-T. Z.120: Message Sequence Charts (MSC) (1996)
52. ITU-T. Z.109: SDL-2000 Combined with UML (2000)
53. ITU-T. X.680 to X.683: Abstract Syntax Notation One (2002)
54. ITU-T. X.692 - Encoding Control Notation (2002)
55. ITU-T. Z.100-Z.106: Specification and Description Language (SDL) (2010)
56. Jensen, K.: Coloured Petri Nets. Basic Concepts, Analysis Methods and Practical Use. vol. 1: Basic Concepts. Monographs in Theoretical Computer Science. Springer (1992)
57. Jensen, K.: Coloured Petri Nets. Basic Concepts, Analysis Methods and Practical Use. Analysis Methods. Monographs in Theoretical Computer Science, vol. 2. Springer (1994)
58. Jensen, K.: Condensed State Spaces for Symmetrical Coloured Petri Nets. Formal Methods in System Design 9, 7–40 (1996)
59. Jensen, K., Kristensen, L.M.: Coloured Petri Nets – Modelling and Validation of Concurrent Systems. Springer (2009)
60. Jensen, K., Kristensen, L.M., Mailund, T.: The sweep-line state space exploration method. Theoretical Computer Science 429, 169–179 (2012)
61. Jensen, K., Kristensen, L.M., Wells, L.: Coloured Petri Nets and CPN Tools for Modelling and Validation of Concurrent Systems. International Journal on Software Tools for Technology Transfer (STTT) 9(3-4), 213–254 (2007)
62. Jørgensen, J.B., Kristensen, L.M.: Computer Aided Verification of Lamport’s Fast Mutual Exclusion Algorithm Using Coloured Petri Nets and Occurrence Graphs with Symmetries. IEEE Transactions on Parallel and Distributed Systems 10(7), 714–732 (1999)
63. Jørgensen, J.B., Kristensen, L.M.: Verification of Coloured Petri Nets Using State Spaces with Equivalence Classes. In: Petri Net Approaches for Modelling and Validation, Lincoln Europa. LINCOS Studies in Computer Science, ch. 2, vol. 1, pp. 17–34 (2003)
64. Kaufman, C.: Internet Key Exchange Protocol. RFC 4306 (December 2005)
65. Kent, S., Seo, K.: Security Architecture for the Internet Protocol. RFC 4301 (December 2005)
66. Kristensen, L.M.: A Perspective on Explicit State Space Exploration of Coloured Petri Nets: Past, Present, and Future. In: Lilius, J., Penczek, W. (eds.) PETRI NETS 2010. LNCS, vol. 6128, pp. 39–42. Springer, Heidelberg (2010)
67. Kristensen, L.M., Jensen, K.: Specification and Validation of an Edge Router Discovery Protocol for Mobile Ad Hoc Networks. In: Ehrig, H., Damm, W., Desel, J., Große-Rhode, M., Reif, W., Schnieder, E., Westkämper, E. (eds.) INT 2004. LNCS, vol. 3147, pp. 248–269. Springer, Heidelberg (2004)
68. Kristensen, L.M., Mailund, T.: A Generalised Sweep-Line Method for Safety Properties. In: Eriksson, L.-H., Lindsay, P.A. (eds.) FME 2002. LNCS, vol. 2391, pp. 549–567. Springer, Heidelberg (2002)
69. Kristensen, L.M., Mailund, T.: Efficient Path Finding with the Sweep-Line Method Using External Storage. In: Dong, J.S., Woodcock, J. (eds.) ICFEM 2003. LNCS, vol. 2885, pp. 319–337. Springer, Heidelberg (2003)
70. Kristensen, L.M., Mechlenborg, P., Zhang, L., Mitchell, B., Gallasch, G.E.: Model-based Development of COAST. STTT 10(1), 5–14 (2007)
71. Kristensen, L.M., Jørgensen, J.B., Jensen, K.: Application of Coloured Petri Nets in System Development. In: Desel, J., Reisig, W., Rozenberg, G. (eds.) ACPN. LNCS, vol. 3098, pp. 626–685. Springer, Heidelberg (2004)

72. Kristensen, L.M., Valmari, A.: Finding Stubborn Sets of Coloured Petri Nets Without Unfolding. In: Desel, J., Silva, M. (eds.) ICATPN 1998. LNCS, vol. 1420, pp. 104–123. Springer, Heidelberg (1998)
73. Kristensen, L.M., Westergaard, M.: Automatic Structure-Based Code Generation from Coloured Petri Nets: A Proof of Concept. In: Kowalewski, S., Roveri, M. (eds.) FMICS 2010. LNCS, vol. 6371, pp. 215–230. Springer, Heidelberg (2010)
74. Kristensen, L.M., Westergaard, M., Nørgaard, P.C.: Model-Based Prototyping of an Interoperability Protocol for Mobile Ad-Hoc Networks. In: Romijn, J.M.T., Smith, G.P., van de Pol, J. (eds.) IFM 2005. LNCS, vol. 3771, pp. 266–286. Springer, Heidelberg (2005)
75. Lai, R., Jirachiefpattana, A.: Communication Protocol Specification and Verification. Kluwer Academic Publishers (1998)
76. Lakos, C.: Modelling Mobile IP with Mobile Petri Nets. *Transactions on Petri Nets and Other Models of Concurrency* 5800, 127–158 (2009)
77. Liu, L.: Verification of the SIP Transaction Using Coloured Petri Nets. In: Mans, B. (ed.) Thirty-Second Australasian Computer Science Conference (ACSC 2009), Wellington, New Zealand. CRPIT, vol. 91, pp. 63–72. ACS (2009)
78. Liu, L.: Security Analysis of Session Initiation Protocol - A Methodology Based on Coloured Petri Nets. In: Proc. of the 2010 International Cyber Resilience Conference (2010)
79. Liu, L., Billington, J.: Verification of the Capability Exchange Signalling protocol. *STTT* 9(3-4), 305–326 (2007)
80. Liu, M.T.: Protocol Engineering. *Advances in Computers* 29, 79–195 (1989)
81. Lorentsen, L., Kristensen, L.M.: Modelling and Analysis of a Danfoss Flowmeter System Using Coloured Petri Nets. In: Nielsen, M., Simpson, D. (eds.) ICATPN 2000. LNCS, vol. 1825, pp. 346–366. Springer, Heidelberg (2000)
82. IETF Mobile Ad-hoc Networks Discussion Archive,
<http://www1.ietf.org/mail-archive/web/manet/current/index.html>
83. Mailund, T.: Analysing infinite-state systems by combining equivalence reduction and the sweep-line method. In: Esparza, J., Lakos, C.A. (eds.) ICATPN 2002. LNCS, vol. 2360, pp. 314–333. Springer, Heidelberg (2002)
84. Malik, R., Mühlfeld, R.: A case study in verification of uml statecharts: the profisafe protocol. *Universal Computer Science* 9(2), 138–151 (2003)
85. Malkin, G.: RIP Version 2. RFC 4822 (February 2007)
86. Milner, R.: *Communication and Concurrency*. Prentice-Hall International (1989)
87. Mortensen, K.H.: Automatic Code Generation Method Based on Coloured Petri Net Models Applied on an Access Control System. In: Nielsen, M., Simpson, D. (eds.) ICATPN 2000. LNCS, vol. 1825, pp. 367–386. Springer, Heidelberg (2000)
88. Narten, T., Nordmark, E., Simpson, W.: Neighbor Discovery for IP Version 6 (IPv6), RFC 2461 (December 1998)
89. Ouyang, C., Billington, J.: On Verifying the Internet Open Trading Protocol. In: Bauknecht, K., Tjoa, A.M., Quirchmayr, G. (eds.) EC-Web 2003. LNCS, vol. 2738, pp. 292–302. Springer, Heidelberg (2003)
90. Ouyang, C., Billington, J.: Formal Analysis of the Internet Open Trading Protocol. In: Núñez, M., Maamar, Z., Pelayo, F.L., Pousttchi, K., Rubio, F. (eds.) FORTE 2004. LNCS, vol. 3236, pp. 1–15. Springer, Heidelberg (2004)

91. Ouyang, C., Kristensen, L.M., Billington, J.: A Formal and Executable Specification of the Internet Open Trading Protocol. In: Bauknecht, K., Tjoa, A.M., Quirchmayr, G. (eds.) EC-Web 2002. LNCS, vol. 2455, pp. 377–387. Springer, Heidelberg (2002)
92. Ouyang, C., Kristensen, L.M., Billington, J.: A Formal Service Specification for the Internet Open Trading Protocol. In: Esparza, J., Lakos, C.A. (eds.) ICATPN 2002. LNCS, vol. 2360, pp. 352–373. Springer, Heidelberg (2002)
93. Petri, C.A.: Kommunikation mit Automaten. Bonn: Institut für Instrumentelle Mathematik, Schriften des IIM Nr. 2 (1962)
94. Popovic, M.: Communication Protocol Engineering. CRC Press (2006)
95. Ratzner, A.V., Wells, L., Lassen, H.M., Laursen, M., Qvortrup, J.F., Stissing, M.S., Westergaard, M., Christensen, S., Jensen, K.: CPN Tools for Editing, Simulating, and Analysing Coloured Petri Nets. In: van der Aalst, W.M.P., Best, E. (eds.) ICATPN 2003. LNCS, vol. 2679, pp. 450–462. Springer, Heidelberg (2003), <http://www.cpntools.org>
96. Ravn, A.P., Srba, J., Viglio, S.: Modelling and verification of web services business activity protocol. In: Abdulla, P.A., Leino, K.R.M. (eds.) TACAS 2011. LNCS, vol. 6605, pp. 357–371. Springer, Heidelberg (2011)
97. Reisig, W.: Petri Nets - An Introduction. EATCS Monographs on Theoretical Computer Science, vol. 4. Springer (1985)
98. Stern, U., Dill, D.L.: Improved Probabilistic Verification by Hash Compaction. In: Camurati, P.E., Eveking, H. (eds.) CHARME 1995. LNCS, vol. 987, pp. 206–224. Springer, Heidelberg (1995)
99. Suriadi, S., Ouyang, C., Smith, J., Foo, E.: Modeling and Verification of Privacy Enhancing Protocols. In: Breitman, K., Cavalcanti, A. (eds.) ICFEM 2009. LNCS, vol. 5885, pp. 127–146. Springer, Heidelberg (2009)
100. Ullman, J.D.: Elements of ML Programming. Prentice-Hall (1998)
101. Valmari, A.: A Stubborn Attack on State Explosion. In: Clarke, E., Kurshan, R.P. (eds.) CAV 1990. LNCS, vol. 531, pp. 156–165. Springer, Heidelberg (1991)
102. Valmari, A.: Stubborn Sets of Coloured Petri Nets. In: Proc. of ICATPN 1991, pp. 102–121 (1991)
103. Valmari, A.: The State Explosion Problem. In: Reisig, W., Rozenberg, G. (eds.) APN 1998. LNCS, vol. 1491, pp. 429–528. Springer, Heidelberg (1998)
104. Vanit-Anunchai, S.: Towards Formal Modelling and Analysis of SCTP Connection Management. In: Proc. of CPN 2009, pp. 163–182 (2008)
105. Vanit-Anunchai, S., Billington, J., Gallasch, G.E.: Analysis of the Datagram Congestion Control Protocols Connection Management Procedures Using the Sweep-line Method. International Journal on Software Tools for Technology Transfer 10(1), 29–56 (2008)
106. Villapol, M.E., Billington, J.: A Coloured Petri Net Approach to Formalising and Analysing the Resource Reservation Protocol. CLEI Electron. J. 6(1) (2003)
107. Villapol, M.E., Billington, J.: Analysing Properties of the Resource Reservation Protocol. In: van der Aalst, W.M.P., Best, E. (eds.) ICATPN 2003. LNCS, vol. 2679, pp. 377–396. Springer, Heidelberg (2003)
108. Vixie, P.: Dynamic Updates in the Domain Name System. RFC 2136 (April 1997)
109. Westergaard, M., Evangelista, S., Kristensen, L.M.: ASAP: An Extensible Platform for State Space Analysis. In: Franceschinis, G., Wolf, K. (eds.) PETRI NETS 2009. LNCS, vol. 5606, pp. 303–312. Springer, Heidelberg (2009), <http://www.daimi.au.dk/~ascoveco/download.html>

110. Westergaard, M., Kristensen, L.M., Brodal, G.S., Arge, L.A.: The ComBack Method – Extending Hash Compaction with Backtracking. In: Kleijn, J., Yakovlev, A. (eds.) ICATPN 2007. LNCS, vol. 4546, pp. 445–464. Springer, Heidelberg (2007)
111. Westergaard, M., Lassen, K.B.: The BRITNeY Suite Animation Tool. In: Donatelli, S., Thiagarajan, P.S. (eds.) ICATPN 2006. LNCS, vol. 4024, pp. 431–440. Springer, Heidelberg (2006)
112. Westergaard, M.: A Game-theoretic Approach to Behavioural Visualisation. *Electr. Notes Theor. Comput. Sci.* 208, 113–129 (2008)
113. Wolper, P., Leroy, D.: Reliable Hashing without Collision Detection. In: Courcoubetis, C. (ed.) CAV 1993. LNCS, vol. 697, pp. 59–70. Springer, Heidelberg (1993)