A Systematic Approach to Constructing Families of Incremental Topology Control Algorithms Using Graph Transformation

Roland Kluge · **Michael Stein** · **Gergely Varró** ·
Andy Schürr · **Matthias Hollick** · **Max Mühlhäuser**

Abstract This document corresponds to the accepted manuscript of the article *Kluge, R., Stein, M., Varró, G., Schürr, A., Hollick, M., Mühlhäuser, M.: "A systematic approach to constructing families of incremental topology control algorithms using graph transformation," in: SoSyM 2017*. The URL of the formal version is `https://doi.org/10.1007/s10270-017-0587-8`. This document is made available under the CC-BY-NC-ND 4.0 license `http://creativecommons.org/licenses/by-nc-nd/4.0/`.

In the communication systems domain, constructing and maintaining network topologies via topology control algorithms is an important cross-cutting research area. Network topologies are usually modeled using attributed graphs whose nodes and edges represent the network nodes and their interconnecting links. A key requirement of topology control algorithms is to fulfill certain consistency and optimization properties to ensure a high quality of service. Still, few attempts have been made to constructively integrate these properties into the development process of topology control algorithms. Furthermore, even though many topology control algorithms share substantial parts (such as structural patterns or tie-breaking strategies), few works constructively leverage these commonalities and differences of topology control algorithms systematically. In previous work, we addressed the constructive integration of consistency properties into the development process. We outlined a constructive, model-driven methodology for designing individual topology control algorithms. Valid and high-quality topologies are characterized using declarative graph constraints; topology control algorithms are specified using programmed graph transformation. We applied a well-known static analysis technique to refine a given topology control algorithm in a way that the resulting algorithm preserves the specified graph constraints.

R. Kluge
Merckstraße 25, 64283 Darmstadt
Tel.: +49-6151-16-22354
Fax: +49-6151-16-22352
E-mail: roland.kluge@es.tu-darmstadt.de

In this paper, we extend our constructive methodology by generalizing it to support the specification of families of topology control algorithms. To show the feasibility of our approach, we reneging six existing topology control algorithms and develop e-kTC, a novel energy-efficient variant of the topology control algorithm kTC. Finally, we evaluate a subset of the specified topology control algorithms using a new tool integration of the graph transformation tool EMOFLON and the SIMONSTRATOR network simulation framework.

Keywords Graph transformation · Graph constraints · Static analysis · Model-driven engineering · Wireless networks · Network simulation

1 Introduction

In the communication systems domain, wireless sensor networks (WSNs) [66, 80] are a highly active research area. For instance, WSNs are applied to monitor physical or environmental conditions using distributed, autonomous, battery-powered sensor nodes that cooperatively transmit their collected data to a central location. To improve important properties (e.g., the battery lifetime of these devices), a topology control (TC) algorithm [66] inactivates redundant communication links of a WSN. Key requirements on a TC algorithm are to (i) handle continuously changing network topologies, (ii) operate in a highly distributed environment, in which each node can only observe and modify its local neighborhood, and (iii) guarantee important local and global formal properties (e.g., bounded node degree or connectivity of the topology) for their neighborhood and the whole network, respectively.

The design and implementation of a TC algorithm are, therefore, challenging and elaborate tasks, which are typically carried out by highly skilled experts. The development of a new TC algorithm is usually an iterative process. In each iteration, (i) a new variant must be individually designed and implemented for a distributed environment, (ii) the preservation of required formal properties must be proved, and (iii) performance measurements must be carried out in a corresponding runtime environment, which is either a network simulator or a hardware testbed. On the one hand, the *specification* of a TC algorithm often builds on a mathematically well-founded framework (e.g., graph or game theory), which allows the TC developer to prove formal properties. On the other hand, the *implementation* of a TC algorithm is typically written in a general-purpose programming language (e.g., Java for simulation [61] or C for testbed evaluation [21]). Additionally, the TC algorithm and the runtime environment (often continuously) interact: The TC algorithm activates and inactivates links in the topology, and the runtime environment causes context events (e.g., node removals due to battery depletion). A *dynamic TC algorithm* has to handle such context events. In many application scenarios of WSNs (e.g., environmental monitoring), these context events are small compared to the size of the entire topology. Therefore, it is crucial that a dynamic TC algorithm reacts to context events in an *incremental* manner, i.e., by retaining unaltered parts of the topology as much as possible.

Current shortcomings State-of-the-art TC literature reveals that the traditional development process of TC algorithms exposes two major shortcomings.

S1 A systematic mapping between the specification and the implementation is missing. This makes it extraordinarily difficult to verify that both representations are indeed equivalent. Incremental TC algorithms are considerably more difficult to develop compared to their batch version, which complicates the traceability between specification and implementation even more.

S2 Novel TC algorithms tend to build on former TC algorithms. Still, these inherent commonalities and differences of TC algorithms are not specified systematically. This reduces reusability among and comparability of TC algorithms. Moreover, such a systematic specification could also enable us to prove formal properties not only for individual TC algorithms but for whole families of TC algorithms at once.

Both shortcomings are well-known research challenges in the communication systems domain. For instance, S1 is addressed in [20, 42, 49, 51, 57, 85, 86] (see also Section 7.1), and S2 is addressed in [3, 4, 8, 18, 26, 53, 55, 64, 65] (see also Section 7.3).

Previous work on S1 In [42, 43], we showed that model-driven principles [10], as applied in many success stories [34, 79], are suitable to address S1. We describe topologies as graph-based models and possible operations of topology control algorithms as declarative graph transformation (GT) rules [63]. Although this approach provides a well-defined procedure for modeling the static and dynamic aspects of TC algorithms in general, it does not ensure that all required formal properties are fulfilled for the resulting topology. To this end, a well-known, constructive, static analysis technique [33] has been established in the GT community to formulate structural invariants in terms of graph constraints and to guarantee that these graph constraints hold. Graph constraints specify positive or negative graph patterns, which must be present in or missing from a valid graph, respectively. Based on these graph constraints and a set of GT rules, a refinement algorithm enriches the GT rules with additional application conditions. These application conditions are derived from the graph constraints and ensure that the applying the refined GT rules never produces invalid graphs w.r.t. the graph constraints. This technique could previously only be applied to scenarios where invariants must hold permanently (e.g., [44]). In the WSN domain, the situation is different because context events inevitably violate the specified graph constraints. To address this problem, we first relaxed the original constraints by introducing appropriate intermediate states of links. Then, we ensured that the specification of the TC algorithm always preserves the relaxed constraints, using the constructive approach described in [33]. Finally, we iteratively modified the state configuration of the topology to enforce the original, strong constraints. We illustrated the proposed constructive methodology by re-engineering a *single* existing TC algorithm, kTC [69]. Yet, S2, i.e., describing commonalities and differences of TC algorithms, still remained an open issue in [42, 43].

Contribution In this paper, we tackle S2 by generalizing the constructive, model-driven methodology for designing TC algorithms using graph transformation [42, 43] to support the development of families of TC algorithms. This paper has four major contributions:

(i) We model commonalities and differences of TC algorithms by extracting common structural constraints and specifying the individual part of each TC algorithm in terms of individual attribute constraints. We lift all steps described in our original approach [43] to operate on abstract representations of TC algorithm families.

(ii) To demonstrate the applicability of our approach, we re-engineer six existing TC algorithms (kTC [69], l-kTC [70, 71], XTC [81], RNG [39], GG [80], Yao Graph [84]) and propose e-kTC, a novel, energy-aware variant of kTC [69] that has been inspired by the CTCA algorithm [15].

(iii) We extend the constructive approach with a step that systematically derives context event handlers, which repair all constraint violations that may result from the context events.

(iv) We perform a comparative, simulation-based evaluation of kTC and e-kTC to showcase an integration of the GT tool EMOFLON [47] and the SIMONSTRATOR network simulation environment [61].
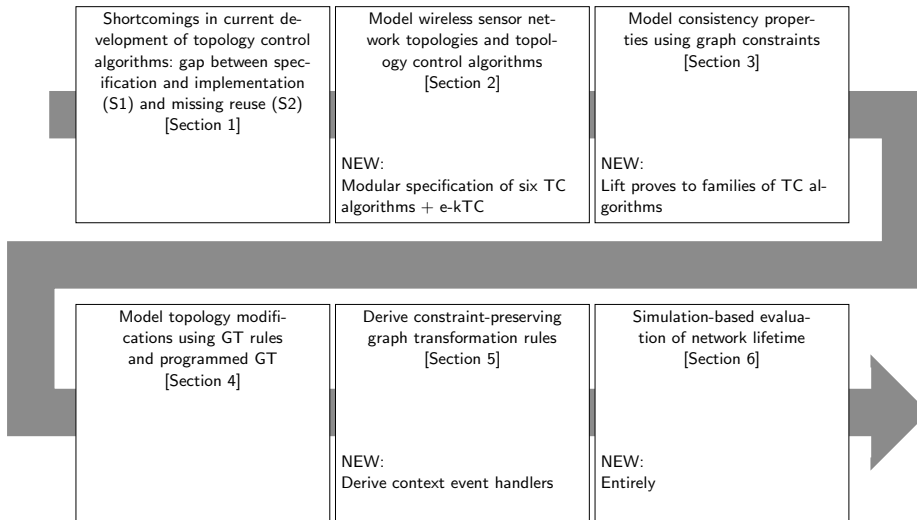
Shortcomings in current development of topology control algorithms: gap between specification and implementation (S1) and missing reuse (S2) [Section 1]	Model wireless sensor network topologies and topology control algorithms [Section 2] NEW: Modular specification of six TC algorithms + e-kTC	Model consistency properties using graph constraints [Section 3] NEW: Lift proves to families of TC algorithms
Model topology modifications using GT rules and programmed GT [Section 4]	Derive constraint-preserving graph transformation rules [Section 5] NEW: Derive context event handlers	Simulation-based evaluation of network lifetime [Section 6] NEW: Entirely

Fig. 1: Structure of this paper (NEW: extensions compared to [42, 43])

Structure Figure 1 maps the major contributions of this paper to the following sections. In Section 2, we specify valid topologies using metamodeling and introduce the six existing and one novel TC algorithm using first-order logic predicates. In Section 3 we specify the TC algorithms using graph constraints and conduct the prove of connectivity based an abstract specification of the family of TC algorithms. In Section 4, we specify topology modifications using GT rules and TC algorithms using programmed GT. In Section 5, we refine the GT rules based on the graph constraints to ensure that the refined GT rules preserve the graph constraints. Additionally, we derive handler operations for the context event rules and ensure that the refined TC algorithm terminates. In Section 6, we present the results of a simulation-based evaluation. In Section 7, we survey related work and conclude this paper in Section 8.

2 Metamodeling and Topology Control

In this section, we introduce basic concepts of metamodeling and TC. Afterwards, we introduce the considered TC algorithms and analyze them w.r.t. recurring substructures.

2.1 Basic Metamodeling Concepts

A *model* describes a set of related entities as a graph whose nodes are *objects* and whose edges are *references*. A *metamodel* specifies all well-formed models of the considered domain as a multi-graph whose nodes are the *classes*, which describe possible entities and serve as *object type*, and whose edges are *associations*, which describe possible relations between entities and serve as *reference type*. Object and reference types have to be *compatible*, i.e., the types of the source and target of a reference are the source and target class of its corresponding association. A class may have multiple typed *attribute*s, which represent properties of its instances. An association end is labeled with a *role name*, which further

describes the corresponding relation, and a *multiplicity*, which restricts the number of corresponding references in a model.

2.2 Topologies

A (network) *topology* represents the state of a communication system as an attributed graph consisting of *nodes* and (communication) *links* [66]. In this paper, we consider topologies of WSNs, i.e., nodes in the topology correspond to battery-powered wireless sensor nodes, and links correspond to the possible direct wireless communication connections between sensor nodes. This implies that a topology is a simple graph, which neither contains loops nor parallel links, i.e., the source and target node of a link are unequal, and each pair of nodes is connected by at most one link. We denote links with the letter e[1] in running text, e.g., e_{12}, e_{ab}, e_{AB}, and as arrow-headed lines in compact notation, e.g., ①→②. By convention, a link e_{ab} has source node n_a and target node n_b. A *path* $P(n_a, n_z) = (e_{ab}, e_{bc}, \ldots, e_{yz})$ from node n_a to node n_z is a list of links where the target node of one link in P is the source node of its successor link in P. In the following, we introduce node and link properties that are required to model the TC algorithms in this paper. A sensor node n_a exposes the following properties:

– An integer-valued unique *identifier* $\mathrm{id}(n_a) = a$ allows to distinguish n_a from other nodes. The identifier of a node is shown in subscript notation in running text, e.g., n_1, and as white label inside the corresponding solid black circle in compact notation, e.g., ①.
– A real-valued *energy* property $\mathrm{E}(n_a)$ stores the current energy level of node n_a, which is typically measured in Joule.
– The real-valued *latitude* $\mathrm{lat}(n_a)$ and *longitude* $\mathrm{long}(n_a)$ capture the position of node n_a (e.g., Euclidean or GPS coordinates).
– The integer-valued *hop count* $\mathrm{hops}(n_a)$ stores the shortest distance (w.r.t. the number of hops) between n_a and a dedicated second node n_0. A *hop* is the traversal of a single link. This property is required in application scenarios such as data collection, where each sensor node periodically sends collected data of its environment to a dedicated base station node n_0. A routing protocol (e.g., RPL [83], AODV [13]) operates on top of the output topology of the current TC algorithm and determines the path between n_a and the base station.

A link e_{ab} exposes the following properties:

– A real-valued generic *weight* $\mathrm{w}(e_{ab})$ stores the cost of using e_{ab} for message transfer. For example, the weight of e_{ab} may be derived from the distance of its incident nodes or the received signal strength indicator (RSSI) at n_b.
– The real-valued *angle* $\alpha(e_{ab})$ of a link e_{ab} can be derived from the positions of its incident nodes as follows:

$$\alpha(e_{ab}) = \mathrm{atan}(\mathrm{lat}(n_a) - \mathrm{lat}(n_b), \mathrm{long}(n_a) - \mathrm{long}(n_b)) + 180°$$

With atan, we denote the arcus tangens operator, which maps a pair of latitudinal and longitudinal differences to the corresponding angle.
– The *state* $\mathrm{s}(e_{ab})$ stores the processing state of e_{ab} during the execution of the TC algorithm. Details follow in Section 2.3.

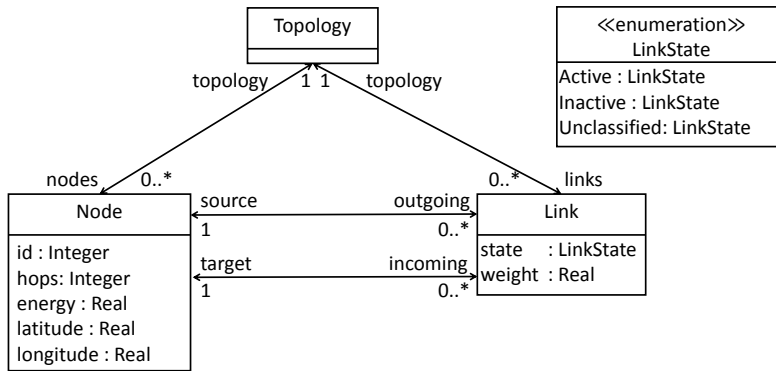[1] We use e instead of l or ℓ for better readability.

Fig. 2: Topology metamodel

Topology metamodel Figure 2 depicts the metamodel of topologies and contains three classes, Topology[2], Node, and Link (depicted as rectangular boxes). Its eight associations (depicted as four bidirectional arrows) specify that (i) each Node and Link is contained in a single Topology, (ii) each Topology contains an unlimited number of Nodes and Links, and (iii) a Node serves as the unique source (target) of any of its zero or more outgoing (incoming) Links. The class Node has two integer-valued attributes (id and hopCount), and three real-valued attributes (energy, latitude, longitude). The class Link has a real-valued weight and a state that can take values Active, Inactive, and Unclassified, specified in the enumeration type LinkState. All attributes and types correspond to the aforementioned node and link properties of the same name.

Figure 3 shows a *(directed) triangle* of links in object and compact notation. In this example, latitude and longitude are Euclidean coordinates, link weights represent the Euclidean distance between the incident nodes, and the hop count is relative to node n_1. Throughout this paper, we assume that every node and link is part of a single topology G with node set V and edge set E. For brevity, we use the compact notation and depict only the relevant attribute values (e.g., the link weight in this case) in the following.

2.3 Topology Control

Topology control (TC) is the discipline of adapting wireless sensor network topologies to optimize network metrics. As described earlier, wireless sensor nodes are typically battery-powered, and often the energy source is not (easily) exchangeable or rechargeable. This makes prolonging the network lifetime a key optimization goal for WSNs [66].

Figure 4 sketches the three phases of the TC process: topology monitoring, planning, and execution. The *topology monitoring* detects *context events*, which are external modifications of the physical topology. In this paper, we consider the following six types of *context event*s: node addition, node removal, link addition, link removal, node property modification, and link property modification. The execution of the TC process is *triggered* either periodically or on-demand, e.g., when a batch of context events has finished. In the *planning step*, the TC algorithm analyzes the input topology and produces a corresponding output topology.

---

[2] We use sans-serif font when referring to metamodel elements.
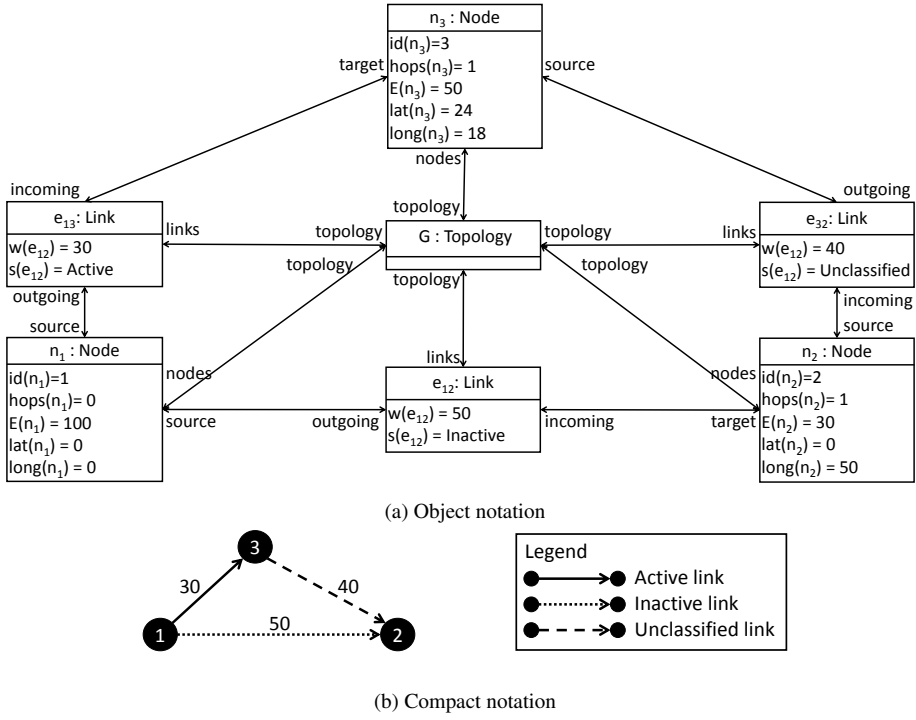
(a) Object notation

(b) Compact notation

Fig. 3: Triangle topology in object and compact notation. Hop count is relative to n_1.

Fig. 4: Topology control process

The output topology contains all links of the input topology that are necessary to fulfill the specified consistency properties (e.g., reduced node degree or connectivity). We distinguish between batch and incremental TC algorithms: A *batch TC algorithm* analyzes the entire input topology and outputs an entire output topology. An *incremental TC algorithm* selectively analyzes the modified parts of the topology and updates the output topology accordingly. This behavior requires that all links that are modified, added, and removed between two iterations of the planning step are marked. In the *execution step*, the sensor node ensures that all links in the output topology are available for message transfer. The planning step is a link classification problem, whereas the execution step and the topology monitoring are

highly platform-specific tasks. Therefore, we focus on the planning step in this paper and leave the necessary refinement of the metamodel for the other two steps as future work.

In Figure 4, we use a joint representation for the input and output topology. The physically possible links in the topology are shown as directed lines. The decision of the TC algorithm is stored as *state* $s(e_{ab})$ attribute for each link e_{ab}. This representation allows us to store previous decisions of the TC algorithm for each link, which is essential for incremental TC algorithms. More precisely, we say that a link e_{ab} in the output topology is

- *active* if $s(e_{ab})=$ Active if it is part of the output topology (denoted as solid line in compact notation, e.g., ①——▷②),
- *inactive* if $s(e_{ab})=$ Inactive if is not part of the output topology (denoted as dotted line in compact notation e.g., ①··▷②),
- *classified* if $s(e_{ab}) \in \{$Active, Inactive$\}$ if the TC algorithm has made a decision for e_{ab} (denoted as mixed dotted-solid line in compact notation, e.g., ①·-▷②),
- *unclassified* $s(e_{ab}) =$ Unclassified if either the TC algorithm has not considered link e_{ab} yet or a context event has invalidated the decision of the TC algorithm (denoted as dashed line in compact notation, e.g., ①- ▷②), and
- *undefined* if we either do not know or do not care about the state of e_{ab} (denoted as gray line in compact notation, e.g., ①——▷②).

A *link state modification* is the modification of the state of a single link, i.e., activation, inactivation, or unclassification of the link.

2.4 Specifying Valid Output Topologies with First-Order Logic Predicates

In the following, we specify required properties of output topologies in terms of first-order logic predicates. We begin with two general properties that must hold for any TC algorithm. Afterwards, we introduce additional algorithm-specific conditions for seven TC algorithms (six existing TC algorithms and one new TC algorithm variant).

2.4.1 General Required Properties of Output Topologies

Upon termination of every TC algorithm, each link in the topology should be classified and the output topology should be connected. These requirements are described by the following two predicates.

Complete Classification Constraint ϕ_{CC}: A TC algorithm should make a definite decision for each link in the topology, i.e., the output topology of every TC algorithm should only contain classified links. This postcondition ensures that a TC algorithm may only terminate after completely classifying the input topology; more formally:

$$\phi_{CC}(G(V,E)) \Leftrightarrow \forall e_{ab} \in E : s(e_{ab}) \in \{\text{Active}, \text{Inactive}\} \tag{1}$$

A-connectivity Predicate $\phi_{A\text{-}conn}$: The output topology must be connected, i.e., each pair of nodes $n_a, n_b \in V$ must be connected by a path $P_{out}(n_a, n_b)$ of active links if a path $P_{in}(n_a, n_b)$

of edges exists in the input topology. This requirement can be described by the following A-connectivity predicate $\phi_{\text{A-conn}}$:

$$\phi_{\text{A-conn}}(G(V,E)) \Leftrightarrow \forall (n_{\text{a}}, n_{\text{b}}) \in V \times V :$$
$$\left(\exists P_{\text{in}}(n_{\text{a}}, n_{\text{b}}) \Rightarrow \exists P_{\text{out}}(n_{\text{a}}, n_{\text{b}}) : \forall e \in P_{\text{out}} : \text{s}(e) = \texttt{Active} \right) \quad (2)$$

Note that A-connectivity can only be evaluated based on a global view of the topology, whereas complete classification can be checked based on local knowledge of each node's outgoing links.

2.4.2 Algorithm-Specific Properties

Each TC algorithm has specific optimization goals, which jointly describe when a link may be inactive in a valid output topology; all links that do not fulfill these conditions have to be active. As an example, we consider the TC algorithm kTC [69]. In a valid output topology of kTC, a link e_{ab} is inactive if and only if (i) it is the weight-maximal link in a triangle, together with classified links e_{ac} and e_{cb}, and (ii) its weight is additionally k times larger than the weight of the weight-minimal link in the same triangle; more formally:

$$\forall e_{\text{ab}} \in E : \ \text{s}(e_{\text{ab}}) = \texttt{Inactive}$$
$$\Leftrightarrow e_{\text{ab}} \text{ is in a triangle with classified links } e_{\text{ac}}, e_{\text{cb}}$$
$$\wedge \ \text{w}(e_{\text{ab}}) \geq \max\left(\text{w}(e_{\text{ac}}), \text{w}(e_{\text{cb}})\right) \quad (3)$$
$$\wedge \ \text{w}(e_{\text{ab}}) \geq k \cdot \min\left(\text{w}(e_{\text{ac}}), \text{w}(e_{\text{cb}})\right).$$

The core idea of kTC is that it is often beneficial to use multiple shorter (i.e., more energy-efficient) links (here: e_{ac} and e_{cb}) instead of one long link (here: e_{ab}) for transferring a message because the required transmission power grows at least quadratically with the length of a link [25].

A closer look at Equation (3) reveals that the algorithm-specific condition contains a structural predicate (the first line, here: a triangle) and an additional attribute predicate (the remaining two lines), which refers to the links identified by the structural predicate. In fact, this is a recurring property in specifications of TC algorithms. We express this separation in the following reformulation of Equation (3).

$$\forall e_{\text{ab}} \in E : \ \text{s}(e_{\text{ab}}) = \texttt{Inactive}$$
$$\Leftrightarrow \exists e_{\text{ac}}, e_{\text{cb}} : \phi_{\Delta}(e_{\text{ab}}, e_{\text{ac}}, e_{\text{cb}}) \wedge \phi_{\text{kTC}}(e_{\text{ab}}, e_{\text{ac}}, e_{\text{cb}})$$
$$\text{with}$$
$$\phi_{\Delta}(e_{\text{ab}}, e_{\text{ac}}, e_{\text{cb}}) = e_{\text{ab}} \text{ is in a triangle with classified links } e_{\text{ac}}, e_{\text{cb}} \quad (4)$$
$$\phi_{\text{kTC}}(e_{\text{ab}}, e_{\text{ac}}, e_{\text{cb}}) = \text{w}(e_{\text{ab}}) \geq \max\left(\text{w}(e_{\text{ac}}), \text{w}(e_{\text{cb}})\right)$$
$$\wedge \ \text{w}(e_{\text{ab}}) \geq k \cdot \min\left(\text{w}(e_{\text{ac}}), \text{w}(e_{\text{cb}})\right).$$

The *directed-triangle predicate* ϕ_{Δ} reflects the structural condition that an inactive link e_{ab} must be part of a triangle. The *kTC predicate* ϕ_{kTC} specifies the condition that e_{ab} must be the weight-maximal link and at least k times larger than the weight-maximal link among e_{ac} and e_{cb}.

Example: Incremental TC using kTC Figure 5 shows the evolution of a sample topology (Figure 5a). For conciseness, we show a link and its reverse link as a single double-headed line. The topology is first optimized by invoking kTC ($k = 2$) (Figure 5b). Then, two links e_{79} and e_{97} are added (e.g., because an obstacle between node n_7 and node n_9 has moved out of the way), and node n_{10} is removed (e.g., because its battery is empty). The resulting topology is shown in Figure 5c. The context event handling has unclassified the new links e_{79} and e_{97} and all links around the removed node n_{10}. Finally, the topology is processed by kTC again (Figure 5d). Now, the added links e_{79} and e_{97} as well as the formerly inactive links $e_{3,11}$, $e_{11,3}$, $e_{9,11}$ and $e_{11,9}$ are active, and the links $e_{3,9}$ and $e_{9,3}$ are inactive.

2.4.3 The Tie-Breaking Predicate $\phi_{tie\text{-}break}$

A recurring issue while developing TC algorithms is that more than one link, e.g., in a triangle, may fulfill the algorithm-specific predicate, which may cause multiple links to be inactivated. In case of kTC, this may even lead to a disconnected output topology. As a resort, tie breaker are applied in such situations. For instance, a link e_{ab} is only inactivated if it has the largest identifier compared to all other links in the triangle that fulfill the algorithm-specific predicate. For triangles of weighted links, the *tie-breaking predicate $\phi_{tie\text{-}break}$* is defined as follows:

$$\phi_{\text{tie-break}}(e_{ab}, e_{ac}, e_{cb}) = (\text{w}(e_{ab}) = \text{w}(e_{ac}) \Rightarrow \text{id}(e_{ab}) > \text{id}(e_{ac}))$$
$$\wedge (\text{w}(e_{ab}) = \text{w}(e_{cb}) \Rightarrow \text{id}(e_{ab}) > \text{id}(e_{cb})) \tag{5}$$

The following Equation (6) shows how the tie-breaking predicate $\phi_{\text{tie-break}}$ can be used to compose a variant of kTC that is guaranteed to inactivate only one weight-maximal link in each triangle.

$$\text{s}(e_{ab}) = \texttt{Inactive} \Leftrightarrow \exists e_{ac}, e_{cb} : \phi_{\Delta}(e_{ab}, e_{ac}, e_{cb}) \wedge \phi_{\text{ kTC}}(e_{ab}, e_{ac}, e_{cb}) \wedge \phi_{\text{tie-break}}(e_{ab}, e_{ac}, e_{cb})$$

$$\text{with}$$

$$\phi_{\Delta} e_{ab}, e_{ac}, e_{cb} = e_{ab} \text{ is in a triangle with classified links } e_{ac}, e_{cb}$$

$$\phi_{\text{ kTC}}(e_{ab}, e_{ac}, e_{cb}) = \text{w}(e_{ab}) \geq \max(\text{w}(e_{ac}), \text{w}(e_{cb}))$$
$$\wedge \text{ w}(e_{ab}) \geq k \cdot \min(\text{w}(e_{ac}), \text{w}(e_{cb}))$$

$$\phi_{\text{tie-break}}(e_{ab}, e_{ac}, e_{cb}) = (\text{w}(e_{ab}) = \text{w}(e_{ac}) \Rightarrow \text{id}(e_{ab}) > \text{id}(e_{ac}))$$
$$\wedge (\text{w}(e_{ab}) = \text{w}(e_{cb}) \Rightarrow \text{id}(e_{ab}) > \text{id}(e_{cb}))$$

$$\tag{6}$$

2.4.4 Maxpower Topology Control

The Maxpower TC algorithm activates all links in a topology. Its name derives from the fact that its output topology contains all links that are available if the node transmits with maximum power. Maxpower TC is a generally accepted baseline for performing network evaluations. The predicate ϕ_{Maxpower} of this algorithm is *false* because it never inactivates a link: Note that Maxpower TC does not even require the additional triangle-identifying predicate ϕ_{Δ}, i.e., the full specification of Maxpower TC looks as follows:

$$\forall e_{ab} \in E : \text{s}(e_{ab}) = \texttt{Inactive} \Leftrightarrow \phi_{\text{Maxpower}}(e_{ab})$$

$$\text{with } \phi_{\text{Maxpower}}(e_{ab}) = \text{false}.$$

(a) Initial topology

(b) After initial TC

(c) After addition of e_{79}, e_{97} and rem. of n_{10}

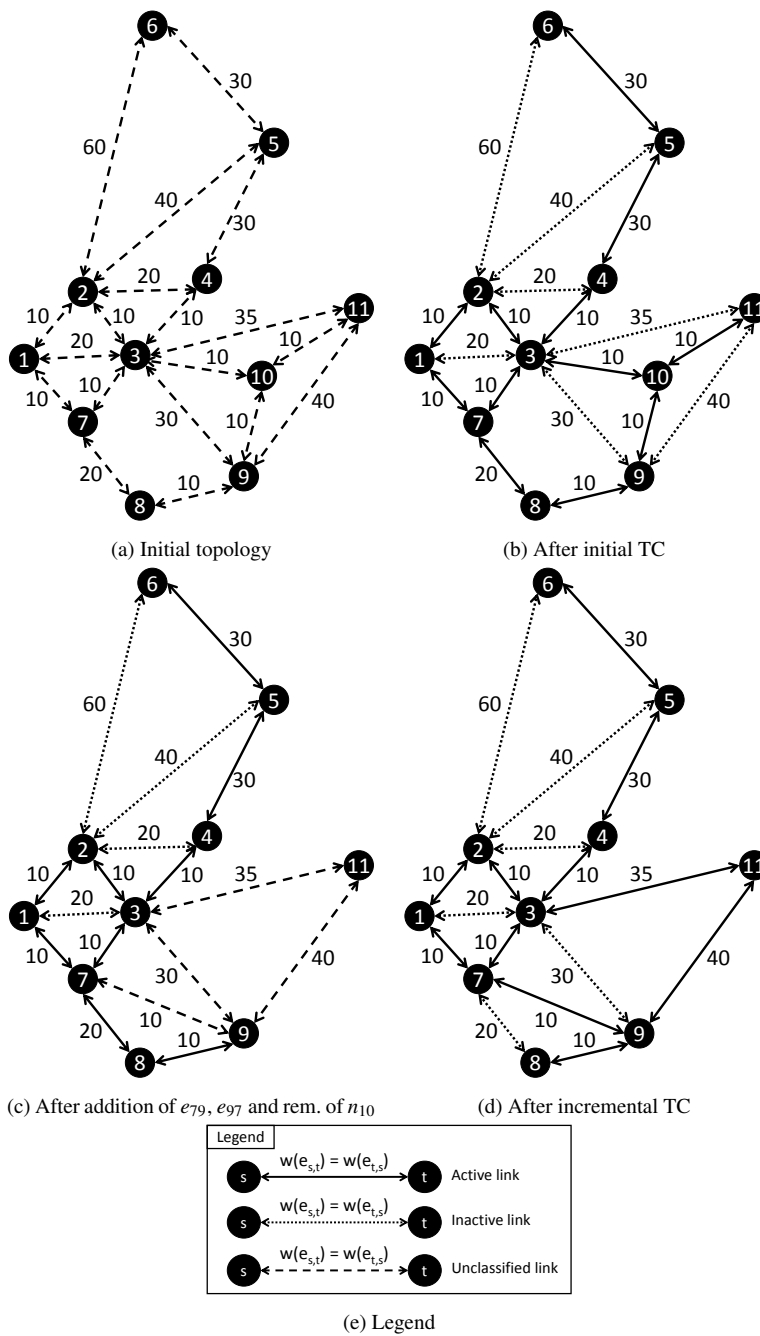(d) After incremental TC

(e) Legend

Fig. 5: Example of incremental TC with kTC ($k = 2$)

2.4.5 XTC Algorithm

The idea behind the XTC algorithm [81, Sec. 3] is that a large link weight indicates a low link quality. A link e_{ab} is inactive in the output topology of XTC if there exist links of higher quality, i.e., smaller weight, that connect the source with the target of e_{ab}, possibly via multiple intermediate links. This is equivalent to the following property: A link in the output topology of the XTC algorithm is inactive if it is the weight-maximal link in some triangle; more formally: $\mathrm{w}(e_{ab}) > \max(\mathrm{w}(e_{ac}), \mathrm{w}(e_{cb}))$. In [81, Sec. 4], the authors refine the XTC algorithm to using the same tie breaking predicate as in Section 2.4.3: Whenever the triangle contains multiple links of the same minimum quality (i.e., of the same maximum weight), only the link with the unique maximum ID is considered. Therefore, we define the XTC predicate ϕ_{XTC} as

$$\phi_{\mathrm{XTC}}(e_{ab}, e_{ac}, e_{cb}) = \mathrm{w}(e_{ab}) \geq \max(\mathrm{w}(e_{ac}), \mathrm{w}(e_{bc}))$$
$$\wedge\ \phi_{\text{tie-break}}(e_{ab}, e_{ac}, e_{cb}).$$

2.4.6 Gabriel Graph Algorithm

The output graph of the Gabriel Graph (GG) algorithm [62, 80] conforms to the following geometric definition. A graph is a GG if for each link e_{ab}, the circle with diameter e_{ab} and center between n_a and n_b contains no nodes apart from n_a and n_b [28]. The original formulation is position-based, i.e., each node requires knowledge about its latitude and longitude. By Thales' theorem [1, p. 50], the following equivalent formulation can be obtained. In each triangle, a link is inactive if its squared weight is smaller than the sum of the squared weights of the other links e_{ac} and e_{cb}:

$$\phi_{\mathrm{GG}}(e_{ab}, e_{ac}, e_{cb}) = \mathrm{w}^2(e_{ab}) > \mathrm{w}^2(e_{ac}) + \mathrm{w}^2(e_{cb}).$$

2.4.7 Relative Neighborhood Algorithm

The output topology of the Relative Neighborhood Graph (RNG) algorithm [39, 80] is an RNG, which is defined as follows. In each triangle, the weight of the weight-maximal link e_{ab} in the triangle must be less than or equal to the weight of the other links e_{ac} and e_{cb}. This means that a link e_{ab} is inactive if it is part of a triangle with shorter links e_{ac} and e_{cb}:

$$\phi_{\mathrm{RNG}}(e_{ab}, e_{ac}, e_{cb}) = \mathrm{w}(e_{ab}) > \mathrm{w}(e_{ac}) \wedge \mathrm{w}(e_{ab}) > \mathrm{w}(e_{cb})$$

Note that the RNG predicate is similar to XTC, which may lead to the impression that the output topology of XTC is (almost) identical to RNG. However, this is true only if (i) the link weight (as used by XTC) correlates strictly negatively with the Euclidean distance (used by RNG), and (ii) link-distances are unique because RNG applies the $>$-operator while XTC applies the $\geq$-operator with ID-based tie breaking.

2.4.8 Local kTC Algorithm

l*-kTC [71] is a variant of kTC that is tailored to many-to-one communication scenarios (e.g., data collection). Here, the hop count attribute $\mathrm{hops}(n_a)$ stores the length (in hops) of the shortest path from n_a to a dedicated base station node. In the output topology of l*-kTC,

a link e_{ab} is inactive if (i) e_{ab} fulfills ϕ_{kTC} and (ii) if its inactivation does not extend the length of the path to the base station by more than a factor a:

$$
\begin{aligned}
\phi_{1*\text{-}kTC}(e_{ab}, e_{ac}, e_{cb}) = {}& \phi_{kTC}(e_{ab}, e_{ac}, e_{cb}) \\
& \wedge \min(\text{hops}(n_a), \text{hops}(n_b), \text{hops}(n_c)) \geq 0 \\
& \Big(\text{hops}(n_a) = \text{hops}(n_b) \Rightarrow true \\
& \wedge \text{hops}(n_a) > \text{hops}(n_b) \Rightarrow \frac{\text{hops}(n_c) + 1}{\max(1, \text{hops}(n_a))} < a \\
& \wedge \text{hops}(n_a) < \text{hops}(n_b) \Rightarrow \frac{\text{hops}(n_c) + 1}{\max(1, \text{hops}(n_b))} < a \Big)
\end{aligned}
\tag{7}
$$

The second line of Equation (7) ensures that the hop count is defined for each participating node. The third line covers the case that n_a and n_b have the same distance to the base station. The fourth line considers the case when n_a is farther away from the base station than n_b. In this case, we may estimate the path length after inactivating e_{ab} as $\text{hops}(n_c) + 1$. As n_a may be the base station itself (i.e., $\text{hops}(n_a) = 0$), we ensure that the denominator is always at least 1. The fifth line is symmetric to the fourth line. More details of the algorithm can be found in [71, p.5].

2.4.9 Yao Graph Algorithm

The Yao graph algorithm [84] is the only location-dependent TC algorithm considered in this paper. This means that it requires information about the latitude and longitude of each node. The Yao graph algorithm separates the environment of a node into *cone*s of uniform angle. If we denote the cone count with n_{cone}, each cone covers an angle of $\frac{360°}{n_{\text{cone}}}$. A link e_{ac} is inactive if a link e_{ac} in the same cone exists that has a smaller weight:

$$
\begin{aligned}
\phi_{\text{Yao}}(e_{ab}, e_{ac}, e_{cb}) = {}& \text{w}(e_{ab}) > \text{w}(e_{ac}) \\
& \wedge \exists x \in \{1, 2, \ldots, n_{\text{cone}}\} : \\
& \Big(\alpha_{\text{cone}} \cdot (x - 1) \leq \alpha(e_{ab}) < \alpha_{\text{cone}} \cdot x \\
& \wedge \alpha_{\text{cone}} \cdot (x - 1) \leq \alpha(e_{ac}) < \alpha_{\text{cone}} \cdot x \Big) \\
& \text{with } \alpha_{\text{cone}} = \frac{360°}{n_{\text{cone}}}
\end{aligned}
$$

2.4.10 e-kTC Algorithm

As the last TC algorithm considered in this paper, we derive a novel, energy-aware variant of kTC, called e-kTC. Its distinctive feature is that it considers the remaining energy of nodes. We begin with an illustrative example that highlights a situation in which the output topology of kTC is suboptimal w.r.t. network lifetime. Afterwards, we present e-kTC and show that it improves the network lifetime of the sample topology.

Network Lifetime In the WSN community, extending the lifetime of a network is a key optimization goal. There are many alternative definitions of network lifetime [14]. We apply a definition that is tailored to the per-node lifetime. A node n_a is *alive* if its remaining energy is positive, i.e., $E(n_a) > 0$. Likewise, a node n_a is *dead* if its battery is empty, i.e., $E(n_a) = 0$. The *d-lifetime* L_d of a network is defined as the first point in time at which at least d nodes are dead.[3] The following values of d are of special interest:

 – $d = 1$ because the network is fully intact before L_1 and
 – $d = |V|$ because the network is no longer operational after $L_{|V|}$.

As a shorthand, $L_{x\%}$ denotes the point in time when $x\% \cdot |V|$ nodes are dead, i.e., $L_{x\%} = L_{x\% \cdot |V|}$. For an energy-aware TC algorithm, it is important to estimate the remaining lifetime of the topology. This allows the TC algorithm to proactively relieve nodes that would otherwise fail soon. The *expected d-lifetime* $\widehat{L}_d(G)$ *of a topology G* estimates the d-lifetime of the topology. For simplicity, we focus on $\widehat{L}_1(G)$ in this paper. The *expected transmission power* $\widehat{P}(e_{ab})$ for each link e_{ab} represents the power that is required to reach n_b from n_a. According to Friis' free space propagation model, $\widehat{P}(e_{ab})$ grows at least proportionally to the squared distance (here: weight) of e_{ab} [25]:

$$\widehat{P}(e_{ab}) \propto w^2(e_{ab})$$

We may estimate the *expected lifetime* $\widehat{L}(e_{ab})$ *of a node n_a w.r.t. a link e_{ab}* as follows:

$$\widehat{L}(e_{ab}) = \frac{E(n_a)}{\widehat{P}(e_{ab})} \qquad (8)$$

Here, we estimate the number of messages that can be transmitted with the remaining energy $E(n_a)$ of n_a. Equation (8) presumes that transmitting a message consumes energy only at the sending node. On real hardware, transmitting a message will also consume energy at the receiving node. In this paper, we neglect this additional cost for simplicity, which is common in the network community (e.g., [15]). We define the *expected lifetime* $\widehat{L}(n_a)$ *of a node n_a* as the minimum expected lifetime $\widehat{L}(e_{ab})$ of its outgoing links:

$$\widehat{L}(n_a) = \min_{e_{ab} \in E} \widehat{L}(e_{ab})$$

Finally, we lift this definition to topologies: The *expected 1-lifetime* $\widehat{L}_1(G)$ *of a topology G* is the minimum expected lifetime $\widehat{L}(n_a)$ of its nodes:

$$\widehat{L}_1(G) = \min_{n_a \in V} \widehat{L}_1(n_a)$$

Motivating example with kTC In the following, we analyze the remaining 1-lifetime L_1 of the sample topology shown in Figure 6a. Each node n is annotated with its remaining energy $E(n)$, and each link e_{ab} is annotated with its weight $w(e_{ab})$ and its expected 1-lifetime $\widehat{L}_1(e_{ab})$. We simulate the behavior of the network over a number of discrete time steps until the first node runs out of energy, using the following workload. In each time step, the remaining energy of a node decreases by the greatest required transmission power among all of its active outgoing links. This is a simplified simulation, e.g., of a Gossiping protocol [37] that broadcasts a message in each time step from a node to all its neighbors. For instance, in each time step, the remaining energy of node n_4 decreases by 9. Only one

---

[3] The variable d alludes to the metric *dead node count*.

(a) Initial topology   (b) After executing kTC   (c) After 12 time steps

(d) Legend
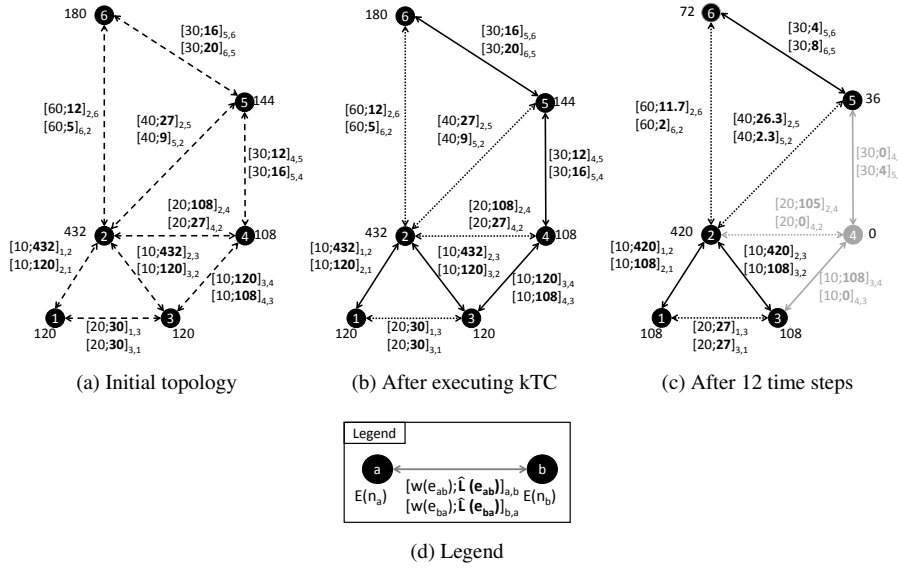
Fig. 6: Applying kTC to the sample topology ($k = 2$, $L_1 = 12$)

execution of kTC ($k = 2$) is required because the link weights are constant (Figure 6b). The example shows that n_4 is the first node to run out of energy after 12 time steps, i.e., $L_1(G) = 12$ (Figure 6c).

Specification of e-kTC The problem of unbalanced energy consumption is well-known in the WSN literature. A number of TC algorithms have been proposed that explicitly take the remaining energy of nodes into account (e.g., [15, 41]). Inspired by [15], we propose to modify the predicate of kTC to take the expected remaining lifetime of nodes into account. We call this energy-aware variant *e-kTC*. After executing e-kTC, a link is inactive if and only if this link is part of a triangle in which it has the minimum expected remaining lifetime among the links in the triangle and if its expected remaining lifetime is at least k times shorter than the maximum expected remaining lifetime of the other links in the triangle:

$$\phi_{\text{e-kTC}}(e_{\text{ab}}, e_{\text{ac}}, e_{\text{cb}}) = \widehat{L}_1(e_{\text{ab}}) \leq \min(\widehat{L}_1(e_{\text{ac}}), \widehat{L}_1(e_{\text{cb}}))$$
$$\wedge\ \widehat{L}_1(e_{\text{ab}}) \leq k \cdot \max(\widehat{L}_1(e_{\text{ac}}), \widehat{L}_1(e_{\text{cb}})). \tag{9}$$

Motivating example with e-kTC Figure 7 illustrates the processing of the same topology as in Figure 6 with e-kTC ($k = 2$). The expected remaining lifetime of a link changes in each time step and opposite links may have different states. To establish comparability with kTC, we invoke e-kTC only once in the beginning (Figure 7b). After 12 time steps, the remaining energy of node n_4 is now 96 (Figure 7c), and the minimal remaining energy among all nodes is 36. Therefore, executing e-kTC increases the remaining 1-lifetime L_1 of the sample topology.

(a) Initial topology
(b) After executing e-kTC
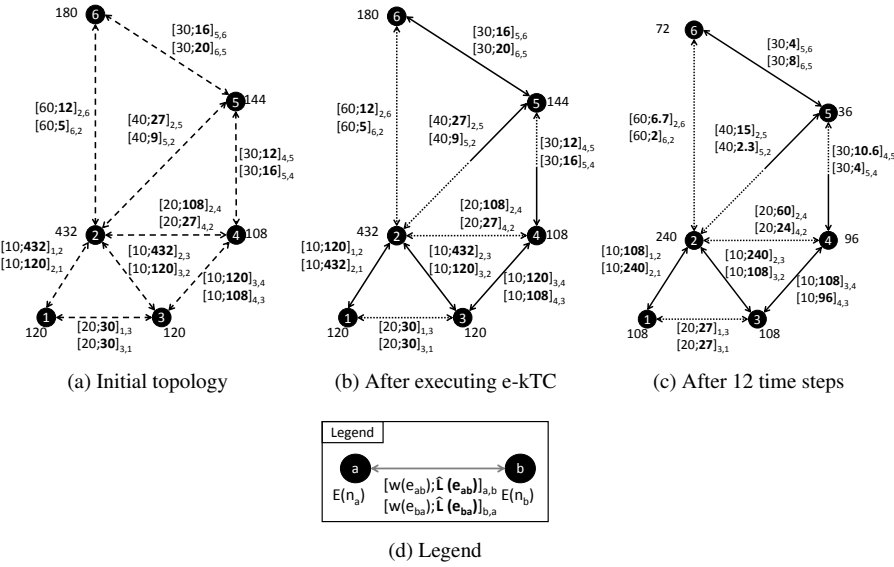(c) After 12 time steps

(d) Legend

Fig. 7: Applying e-kTC to the sample topology ($k = 2$, $L_1 = 16$)

Lifetime preservation with e-kTC In fact, the benefit of applying e-kTC in the previous example can be generalized: For each triangle of links, consisting of e_{ab}, e_{ac}, and e_{cb}, e-kTC preserves the expected lifetime of n_a without decreasing the lifetime of n_b and n_c, compared to applying Maxpower TC. The expected lifetime of n_a even increases if the inactivated link e_{ab} has the minimum expected lifetime among all outgoing links of n_a.

2.4.11 The Minimum-Weight Predicate $\phi_{min\text{-}weight}$

In the following, we introduce the novel minimum-weight predicate $\phi_{min\text{-}weight}$, which can be combined freely with all specified TC algorithms. It serves to reduce the memory footprint and the runtime of the TC algorithm. Working memory is a highly limited resource on wireless sensor nodes. For this reason, keeping the entire neighborhood of a sensor node in working memory may be infeasible if the topology is dense. Fortunately, it is often unnecessary to store links to close neighbors because the energy consumption is typically predominated by links to distant neighbors. Additionally, reducing the size of the processed neighborhood may speed up the planning step.

The following *minimum-weight predicate* $\phi_{min\text{-}weight}$ formalizes this reduction step. The parameter w_{thres} represents the configurable minimal weight of a link to be included in the considered neighborhood.

$$\phi_{min\text{-}weight}(e_{ab}, e_{ac}, e_{cb}) = \min\left(w(e_{ab}), w(e_{ac}), w(e_{cb})\right) \geq w_{thres} \tag{10}$$

This predicate may now be used to compose new variants of the previously specified TC algorithms. For instance, a modified version of kTC with minimum-weight predicate is

$$\phi_{kTC+min\text{-}weight}(e_{ab}, e_{ac}, e_{cb}) = \phi_{kTC}(e_{ab}, e_{ac}, e_{cb}) \wedge \phi_{min\text{-}weight}(e_{ab}, e_{ac}, e_{cb}).$$
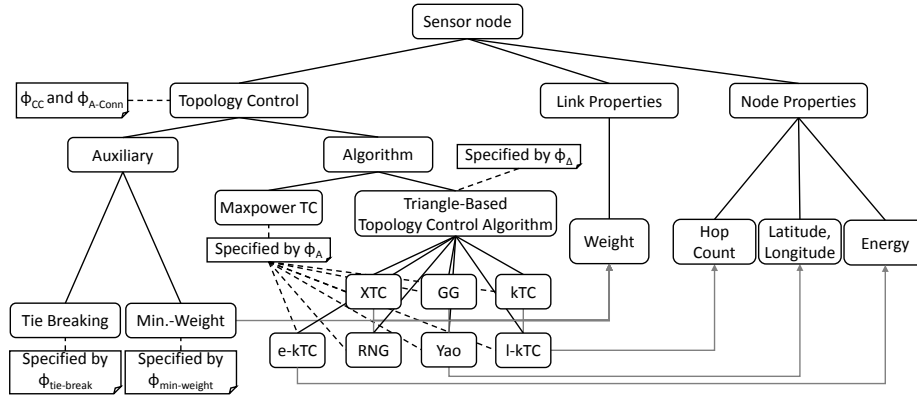
Fig. 8: Overview of configuration options of a sensor node

2.5 Summary of TC Algorithms

Figure 8 and Table 1 summarize the results of this section. Figure 8 illustrates the configuration options of a single sensor node: the possible TC algorithms, the minimum-weight optimization, and the relevant node and link properties. Black solid lines indicate the hierarchical decomposition relation between configuration options. Gray lines indicate dependencies from TC algorithms to properties. Dashed black lines indicate the relationship between the components of the considered TC algorithms and their corresponding predicate.

The output topology of any TC algorithm must fulfill the complete-classification predicate ϕ_{CC} and the A-connectivity predicate $\phi_{A\text{-conn}}$ (Section 2.4.1). The triangle-identifying predicate ϕ_Δ is common to all considered algorithms except for Maxpower TC. In contrast, the predicate $\phi_A \in \{\phi_{kTC}, \phi_{Maxpower}, \phi_{XTC}, \phi_{GG}, \phi_{RNG}, \phi_{l*\text{-}kTC}, \phi_{Yao}, \phi_{e\text{-}kTC}\}$ describes the algorithm-specific attribute constraints. The auxiliary predicates $\phi_{tie\text{-}break}$ and $\phi_{min\text{-}weight}$ describe the ID-based tie-breaking (Section 2.4.3) and the minimum-weight optimization (Section 2.4.11). Table 1 summarizes the algorithm-specific predicates ϕ_A. For simplicity, we write ϕ instead of ϕ_A in the following.

3 Characterizing Valid Topologies with Graph Constraints

In this section, we introduce graph patterns and graph constraints to specify the desired formal properties of topologies. Compared to first-order logic predicates, the main advantage of graph constraints is that they can be constructively combined with GT rules to produce refined constraint-preserving GT rules. The specification using graph constraints allows us to prove that all considered TC algorithms produce connected output topologies. As one main contributions of this paper, we lift the proof of connectivity to entire families of TC algorithms.

Table 1: Attribute predicates of the considered TC algorithms

Algorithm A	Algorithm-Specific Predicate $\phi_A(e_{ab}, e_{ac}, e_{cb})$
Maxpower TC	*false*
XTC [81]	$\mathrm{w}(e_{ab}) \geq \max(\mathrm{w}(e_{ac}), \mathrm{w}(e_{cb}))$ $\wedge \ (\mathrm{w}(e_{ab}) = \mathrm{w}(e_{ac}) \Rightarrow \mathrm{id}(e_{ab}) > \mathrm{id}(e_{ac}))$ $\wedge \ (\mathrm{w}(e_{ab}) = \mathrm{w}(e_{cb}) \Rightarrow \mathrm{id}(e_{ab}) > \mathrm{id}(e_{cb}))$
GG [80]	$\mathrm{w}^2(e_{ab}) > \mathrm{w}^2(e_{ac}) + \mathrm{w}^2(e_{cb})$
RNG [39]	$\mathrm{w}(e_{ab}) > \max(\mathrm{w}(e_{ac}), \mathrm{w}(e_{cb}))$
kTC [69]	$\mathrm{w}(e_{ab}) \geq \max(\mathrm{w}(e_{ac}), \mathrm{w}(e_{cb}))$ $\wedge \ \mathrm{w}(e_{ab}) \geq k \cdot \min(\mathrm{w}(e_{ac}), \mathrm{w}(e_{cb}))$ $\wedge \ (\mathrm{w}(e_{ab}) = \mathrm{w}(e_{ac}) \Rightarrow \mathrm{id}(e_{ab}) > \mathrm{id}(e_{ac}))$ $\wedge \ (\mathrm{w}(e_{ab}) = \mathrm{w}(e_{cb}) \Rightarrow \mathrm{id}(e_{ab}) > \mathrm{id}(e_{cb}))$
l*-kTC [71]	$\phi_{\mathrm{kTC}}(e_{ab}, e_{ac}, e_{cb})$ $\wedge \min(\mathrm{hops}(n_a), \mathrm{hops}(n_b), \mathrm{hops}(n_c)) \geq 0$ $\Big(\mathrm{hops}(n_a) = \mathrm{hops}(n_b) \Rightarrow true$ $\wedge \mathrm{hops}(n_a) > \mathrm{hops}(n_b) \Rightarrow \dfrac{\mathrm{hops}(n_c) + 1}{\max(1, \mathrm{hops}(n_a))} < a$ $\wedge \mathrm{hops}(n_a) < \mathrm{hops}(n_b) \Rightarrow \dfrac{\mathrm{hops}(n_c) + 1}{\max(1, \mathrm{hops}(n_b))} < a \Big)$
Yao graph [84]	$\mathrm{w}(e_{ab}) > \mathrm{w}(e_{ac}) \ \wedge$ $\exists x \in \{1, 2, \dots, n_{\mathrm{cone}}\} :$ $\big(\dfrac{360°}{n_{\mathrm{cone}}} \cdot (x-1) \leq \alpha(e_{ab}) < \dfrac{360°}{n_{\mathrm{cone}}} \cdot x$ $\wedge \ \dfrac{360°}{n_{\mathrm{cone}}} \cdot (x-1) \leq \alpha(e_{ac}) < \dfrac{360°}{n_{\mathrm{cone}}} \cdot x \ \big)$
e-kTC [this paper]	$\widehat{\mathrm{L}}_1(e_{ab}) < \min(\widehat{\mathrm{L}}_1(e_{ac}), \widehat{\mathrm{L}}_1(e_{cb}))$ $\widehat{\mathrm{L}}_1(e_{ab}) \leq k \cdot \max(\widehat{\mathrm{L}}_1(e_{ac}), \widehat{\mathrm{L}}_1(e_{cb}))$

3.1 Graph Constraint Concepts

A *pattern* is a graph consisting of node and link variables together with a set of attribute constraints. A *node (link) variable* serves as a placeholder for a node (link) in a topology. An *attribute constraint* is a predicate over attributes of node and link variables. A *match m of a pattern p in a topology G* injectively maps the node and link variables of p to the nodes and links of G, respectively, such that all attribute constraints are fulfilled. Additionally, a match must preserve the end nodes of link variables, i.e., map the incident node variables of each link variable to the incident nodes of the corresponding link.

A *graph constraint C_x* consists of a *premise* pattern p_x and a *conclusion c_x* that consists of zero or more *conclusion pattern*s $c_{x,y}$. The premise is a subgraph of each conclusion pattern, and the attribute constraints of each conclusion pattern jointly imply the attribute constraints of the premise. A graph constraint with no (at least one) conclusion patterns is called *negative (positive) graph constraint*. A *graph constraint C_x is fulfilled on a topology G* (alternatively: a topology G fulfills a graph constraint C_x) if every match of its premise p_x can be extended to a match of at least one conclusion pattern $c_{x,y}$. This implies that a
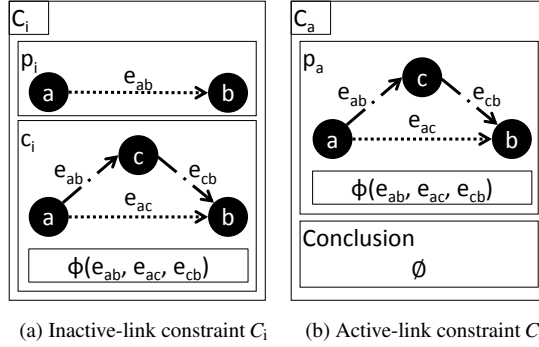
Fig. 9: No-unclassified-links constraint C_u

(a) Inactive-link constraint C_i (b) Active-link constraint C_a

Fig. 10: Abstract algorithm-specific graph constraints for triangle-based TC algorithms

topology fulfills a negative graph constraint C_x if it does not contain *any* matches of p_x. Note that this requirement is sufficient but not necessary for positive graph constraints. A *graph constraint C_x is fulfilled at a match m of a pattern p in a topology G if every match* of p_x that maps all node (link) variables that appear also in p to the same nodes (links) in G can be extended to a match of at least one conclusion pattern $c_{x,y}$. A *match m′ of a pattern p′ in G extends a match m of a pattern p if every node (link) variable that appears in p′ and in p is bound to the same node (link), respectively.*

No-unclassified-links constraint C_u   The no-unclassified-links constraint C_u is equivalent to the complete-classification property ϕ_{CC}, which must be fulfilled by every TC algorithm (Equation (1)). Its premise p_{no-u} matches any unclassified link e_{ab}, and its conclusion is empty because it is a negative constraint. This means that C_u can only be fulfilled on a topology if this topology does not contain any unclassified links.

Algorithm-specific graph constraints   Figure 10 shows two graph constraints that describe the algorithm-specific properties of a triangle-based TC algorithm. The *inactive-link constraint C_i* (Figure 10a) states that each inactive link e_{ab} must be part of a triangle together with classified links e_{ac} and e_{cb} for which $\phi(e_{ab}, e_{ac}, e_{cb})$ is true. Symmetrically, the *active-link constraint C_a* (Figure 10b) states that each active link e_{ab} may not be part of a triangle together with classified links e_{ac} and e_{cb} for which $\phi(e_{ab}, e_{ac}, e_{cb})$ is true.

3.2 Consistency of Topologies

In the following, we categorize topologies according to which of the specified graph constraints they fulfill. We distinguish between the following three levels of *consistency*:

- A topology is *strongly consistent* if it fulfills the no-unclassified-links constraint C_u, the inactive-link constraint C_i, and the active-link constraint C_a. A valid output topology of a TC algorithm must be strongly consistent.
- A topology is *weakly consistent* if it fulfills the inactive-link constraint C_i and the active-link constraint C_a. For instance, an entirely unclassified topology is weakly consistent because it contains matches of the premises of neither C_i nor C_a.
- A topology is *inconsistent* if it is neither weakly nor strongly consistent, i.e., if it fails to fulfill at least one of the algorithm-specific constraints C_i and C_a.

In the following, we require that the topology is weakly consistent before and strongly consistent after invoking a TC algorithm, i.e., weak consistency is the precondition and strong consistency is the postcondition of every TC algorithm.

3.3 Connectivity of Topologies

Next, we show that the specified graph constraints already allow us to prove the important property that the output topology of a TC algorithm must be A-connected (see also Section 2.4.1). A-connectivity is a hard constraint that may be violated by context events because context events may unclassify links. For such situations, we need a second, softer constraint for defining connectivity in the presence of context events. A topology is *A-U-connected* if the subgraph consisting of its active and unclassified links is connected. According to this definition, A-connectivity implies A-U-connectivity. The idea behind A-U-connectivity is that unclassified links should be treated as if they were active. The context event handlers of a TC algorithm must ensure that the topology remains A-U-connected. The following Theorems 1 and 2 show that the output topology of the considered TC algorithms is A-connected if the input topology is A-U-connected and weakly consistent. This property must be proved w.r.t. the TC algorithm A because weak or strong consistency is always evaluated in the context of the current TC algorithm A.

Theorem 1 *For each considered TC algorithm A, a strongly consistent and A-U-connected topology is also A-connected.*

Sketch of Proof. Let A be the TC algorithm whose active-link constraint C_a and inactive-link constraint C_i are fulfilled. From strong consistency follows that the topology fulfills the no-unclassified-links constraint C_u. Therefore, it suffices to show the following claim: The end nodes of each link are connected by a path of active links in the output topology. This trivially holds for the end nodes of active links. By induction, we show that the claim also holds for all inactive links.

Let $e_{i_1} \prec_A e_{i_2} \prec_A \ldots \prec_A e_{i_k}$ be a strict ordering of the inactive links in the topology such that in each match m of the conclusion c_i of the inactive-link constraint C_i, the link corresponding to e_{ab} is larger (w.r.t. $\prec_A$) than the links corresponding to e_{ac} and e_{cb}, i.e., $m(e_{ac}) \prec_A m(e_{ab}) \land m(e_{cb}) \prec_A m(e_{ab})$. Corresponding definitions of $\prec_A$ for all considered algorithms in this paper are shown in Table 2.

Induction start: The minimal inactive link w.r.t. ϕ_A, e_{i_1}, is part of a triangle with two active links that connect the end nodes of link e_{i_1} because the inactive-link constraint C_i is fulfilled. Thus, the claim holds for link e_{i_1}.

Table 2: Algorithm-specific link order $\prec_A$ for the considered algorithms

Algorithm A	Link Order $\prec_A$
Maxpower TC, XTC [81], GG [62, 80], RNG [39], kTC [69], l-kTC [71], Yao graph [84]	$e_{ab} \prec_A e_{cd} \Leftrightarrow w(e_{ab}) < w(e_{cd})$ $\vee \left(w(e_{ab}) = w(e_{cd}) \wedge id(e_{ab}) < id(e_{cd}) \right)$
e-kTC (Section 2.4.10)	$e_{ab} \prec_A e_{cd} \Leftrightarrow \widehat{L}_1(e_{ab}) < \widehat{L}_1(e_{cd})$ $\vee \left(\widehat{L}_1(e_{ab}) = \widehat{L}_1(e_{cd}) \wedge id(e_{ab}) < id(e_{cd}) \right)$

Induction step: We now consider an inactive link $e_{i_{\ell+1}}$ with $1 \leq \ell \leq k-1$, which is part of a triangle with two classified links, e_1 and e_2. We assume that only e_1 is inactive.[4] Link e_1 appears before $e_{i_{\ell+1}}$ in $\prec_A$, i.e., there is some $s \leq \ell$ such that $e_1 := e_{i_s}$. Since the claim has been proved for all inactive links that appear before $e_{i_{\ell+1}}$ in $\prec_A$, there is a path of active links between the end nodes of e_{i_s}. A path of active links between the end nodes of link $e_{\ell+1}$ can be constructed by joining the two paths between the end nodes of e_1 and e_2. □

Theorem 2 *The output topology of each considered TC algorithm A is A-connected if the input topology is A-U-connected and weakly consistent.*

Sketch of Proof. By definition, a TC algorithm turns a weakly consistent input topology into a strongly consistent output topology. According to Theorem 1, a strongly consistent and A-U-connected topology is A-connected. Therefore, the claim follows. □

4 Specifying Topology Control with Programmed Graph Transformation

In this section, we specify TC operations and context events using graph transformation (GT) rules and TC algorithms using programmed graph transformation [23], which carries out basic topology modifications by applying graph transformation rules whose execution order is defined by an explicit control flow.

Programmed Graph Transformation Concepts A *graph transformation (GT) rule* [22, 63] consists of a *left-hand side (LHS)* pattern, a *right-hand side (RHS)* pattern, and *application conditions (ACs)*. A *positive (negative) application condition (PAC (NAC))* is a positive (negative) graph constraint. To enable meaningful attribute assignments, a predicate in the attribute constraints of the RHS pattern can only be an equation with an attribute of a single node or link variable on its left side. A *GT rule is applicable to a topology G* if a match of the LHS pattern exists in the topology that fulfills all application conditions. The *application of a GT rule at a match of its LHS pattern to a topology G produces a topology G'* as follows.

(i) All nodes (links) of the topology that map to a node (link) variable of the LHS pattern and have a corresponding node (link) variable in the RHS pattern are *preserved*.
(ii) All nodes (links) of the topology that are assigned to a node (link) variable of the LHS pattern and lack a corresponding node (link) variable in the RHS pattern are *removed*,
(iii) For each node (link) variable in the RHS that is missing in the LHS, a corresponding node (link) is *added* to the topology.

---

[4] If both links are active, the claim follows trivially. If both links are inactive, the argument applies for each link individually.
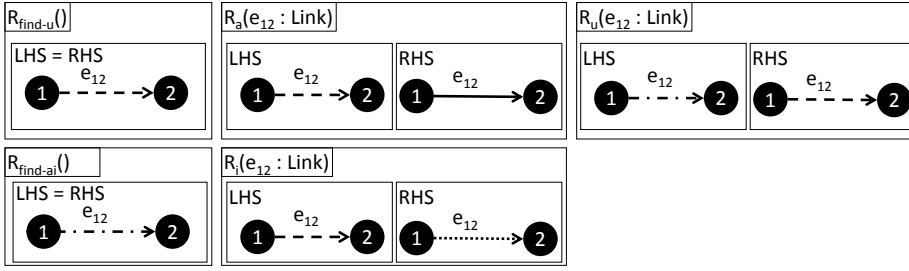
Fig. 11: TC rules

(iv) Node (link) attributes are *assigned* in such a manner that the attribute constraints of the RHS pattern are fulfilled (operator =).

Specifying control flow using programmed GT Control flow is specified in our approach with Story-Driven Modeling [23], an activity-diagram-based notation in which each activity node contains a graph transformation rule. A *regular activity node* (denoted by a single framed, rounded rectangle) with one unlabeled outgoing edge applies the contained GT rule once to one arbitrary match. A regular activity node with an outgoing [Success] and [Failure] edge applies the contained GT rule and follows the [Success] edge if the rule is applicable at an arbitrary match, and it follows the [Failure] edge if the rule is inapplicable. A *foreach activity node* (denoted by a double-framed rounded rectangle) applies the contained GT rule to all matches and traverses along the optional outgoing edge labeled with [Success] for each match. When all the matches have been completely processed, the control flow continues along the [Failure] outgoing edge. RHS node and link variables are *bound* by a successful rule application and can be reused in subsequent activity nodes.

Topology control rules Three of the five *TC rules*, shown in Figure 11, represent link state modifications inside the GT specification of a TC algorithm: the activation rule R_a, the inactivation rule R_i, and the unclassification rule R_u. The remaining two TC rules are the unclassified-link-identification rule R_{find-u} and the classified-link-identification rule $R_{find-ai}$, which serve to identify (but not to modify) unclassified and classified links and whose LHS and RHS are identical.

Context event rules Four of the nine *context event rules*, as shown in Figure 12, specify structural modifications of a topology: the node addition rule R_{+n}, the node removal rule R_{-n}, the link addition rule R_{+e}, and the link removal rule R_{-e}. The remaining five context event rules specify attribute value modifications: the weight modification rule R_{mod-w}, the hop count modification rule R_{mod-h}, the energy modification rule R_{mod-E}, the latitude modification rule $R_{mod-lat}$, and the longitude modification rule $R_{mod-long}$.

Specification of the Maxpower TC algorithm Figure 13 shows a specification of the Maxpower TC algorithm (see Section 2.4.4), which serves as starting point for the subsequent development steps. Link variable e_{12} is bound by the unclassified-link-identification rule R_{find-u} and reused in the activity nodes containing the activation rule R_a and the inactivation rule R_i. Note that R_i is only shown here for completeness, even though it is unreachable because R_a is always applicable.
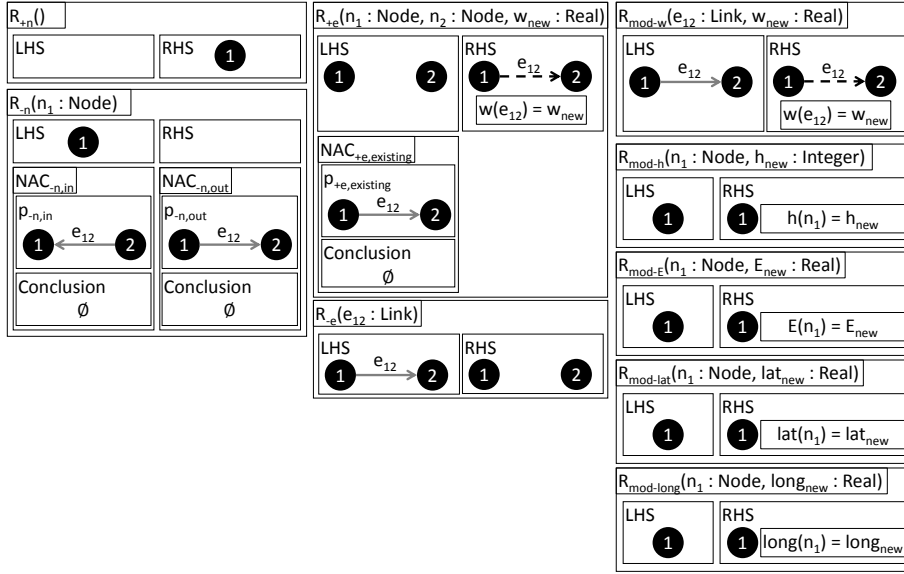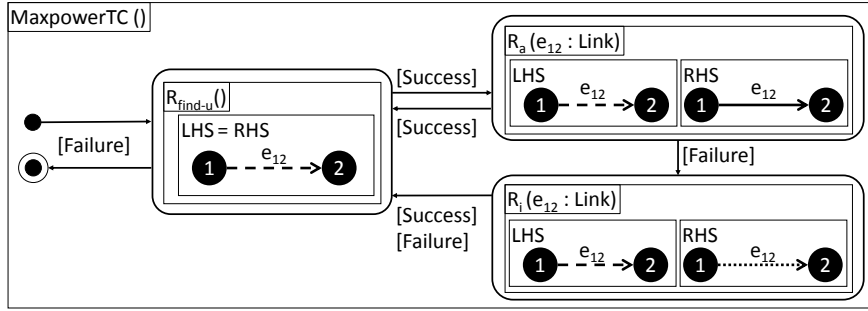
Fig. 12: Context event rules

Fig. 13: Specification of the Maxpower TC algorithm

5 Refining the Graph Transformation Rules to Preserve the Graph Constraints

The fourth step of our design methodology combines the TC and context event rules with the graph constraints to produce a refined version of the TC algorithm, which preserves all graph constraints. This step is split into three substeps, as shown in Figure 14.

– In Section 5.1, we enrich the TC and context event rules step-by-step with additional application conditions that are derived from the graph constraints to ensure that the refined rules preserve all graph constraints. After this step, the TC algorithm preserves weak consistency.
– In Section 5.2, we systematically transform the additional application conditions of the context event rules into equivalent GT-based handler operations. These handler operations ensure that the topology is constantly weakly consistent. This step is necessary be-
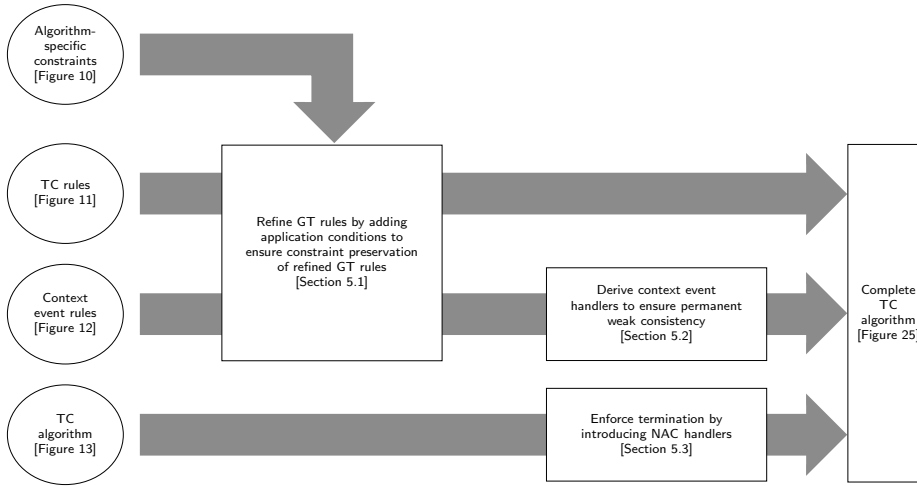
Fig. 14: Overview of refinement step (Section 5)

cause the applicability of context event rules may not be restricted because they present uncontrollable effects of the environment. The derivation of context event handlers is one of the main contributions of this paper.
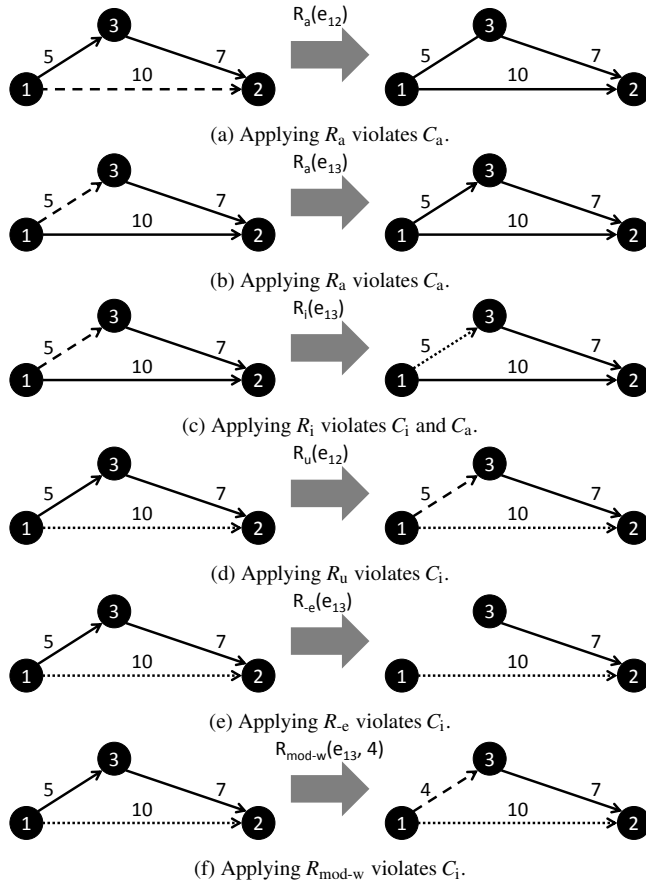
– In Section 5.3, we refine the structure of the TC algorithm to enforce its termination. Shared application conditions of the refined activation rule R_a and inactivation rule R_i may lead to a non-termination of the TC algorithm. We systematically transform the shared application conditions into handler operations that ensure the termination of the TC algorithm. Finally, we prove that the resulting algorithm terminates.

5.1 Refinement of TC and Context Event Rules

During this first refinement step, we (i) analyze which of the specified TC and context event rules preserve or violate the inactive-link constraint C_i and the active-link constraint C_a, which constitute weak consistency, and (ii) apply a well-known static analysis technique [33] that produces additional application conditions that prevent the GT rules from violating weak consistency.

5.1.1 Analysis of Consistency Preservation

The examples shown in Figure 15 illustrate situations in which applying one of the GT rules violates weak consistency. Every example is shown in the form of a weakly consistent initial topology, a rule application, and a constraint-violating final topology. These findings can be generalized to the results in Table 3. Figure 15 provides examples for each entry in the table that corresponds to a constraint-violating GT rule (denoted by a cross mark, ×). In the following, we discuss why the remaining GT rules are guaranteed to preserve the graph constraints, as indicated by the nine check marks (✓) in the table.

(a) Applying R_a violates C_a.

(b) Applying R_a violates C_a.

(c) Applying R_i violates C_i and C_a.

(d) Applying R_u violates C_i.

(e) Applying R_{-e} violates C_i.

(f) Applying R_{mod-w} violates C_i.

Fig. 15: Sample constraint violations caused by applying GT rules for kTC with $k = 2$

Table 3: Overview of constraint preservation ($\checkmark$: constraint-preserving, $\times$: constraint-violating, *: depends on TC algorithm, †: assuming unclassification of incident unclassified links)

Rule	C_i	C_a	Remarks
R_a	$\checkmark$	$\times$	Figures 15a and 15b
R_i	$\times$	$\times$	Figure 15c
R_u	$\times$	$\checkmark$	Figure 15d
R_{find-u}, $R_{find-ai}$	$\checkmark$	$\checkmark$	No modification
R_{+n}, R_{-n}	$\checkmark$	$\checkmark$	Added/removed node is isolated
R_{+e}	$\checkmark$	$\checkmark$	Added link is unclassified
R_{-e}	$\times$	$\checkmark$	Figure 15e
R_{mod-w}	$\times^*$	$\checkmark$	Figure 15f, nalogous to R_u
$R_{mod-lat}$, $R_{mod-long}$	$\times^*$	$\checkmark^{*\dagger}$	Analogous to R_u
R_{mod-E}, R_{mod-h}	$\times^*$	$\checkmark^{*\dagger}$	Analogous to R_u

- The activation rule R_a preserves the inactive-link constraint C_i because applying R_a neither creates a new match of the premise of C_i nor destroys a match of the conclusion c_i of C_i.
- The unclassification rule R_u preserves the active-link constraint C_a because unclassifying a link may never result in a new match of the premise p_a of C_a.
- The unclassified-link-identification rule $R_{\text{find-u}}$ and the classified-link-identification rule $R_{\text{find-ai}}$ do not modify the topology and, therefore, preserve both constraints.
- The node addition rule R_{+n} (node removal rule R_{-n}) only adds (removes) an isolated node to (from) the topology, which may neither produce a new match of p_a or p_i nor destroy a match of c_i.
- The link addition rule R_{+e} preserves C_i and C_a because the added link is unclassified.
- The link removal rule R_{-e} preserves the active-link constraint C_a because removing a link cannot establish a new match of its premise p_a.

The situation is more difficult for the attribute modification rules. We may assume that the sensor node is configured to suppress context events that modify attributes irrelevant for the current TC algorithm. Therefore, the following explanations only hold if the considered TC algorithm depends on the particular attribute. Independent of the considered TC algorithm, the weight modification rule $R_{\text{mod-w}}$ preserves the active-link constraint C_a but does not preserve the inactive-link constraint C_i because this rule unclassifies the modified link and, thereby, behaves similar to the unclassification rule R_u. For instance, whenever the hop count of a node n_a is modified (by applying the hop count modification rule $R_{\text{mod-h}}$), we may need to evaluate for each incoming and outgoing link of n_a whether its state needs to be updated. A conservative resolution strategy is to unclassify all incident links of a node whenever one of its attribute values changes. Under this assumption, the hop count modification rule $R_{\text{mod-h}}$, energy modification rule $R_{\text{mod-E}}$, the latitude modification rule $R_{\text{mod-lat}}$, and the longitude modification rule $R_{\text{mod-long}}$ behave equivalently to a sequence of applications of the unclassification rule R_u.

5.1.2 Refinement of Graph Transformation Rules

In this section, we show how the identified constraint-violating GT rules can be refined to preserve the graph constraints. In total, we have identified twelve problematic pairs of GT rules and graph constraints, where a GT rule may violate a particular graph constraint. These pairs serve as input for the constructive refinement algorithm that has first been presented by Heckel and Wagner [33] for purely structural graph constraints and later extended by Deckwerth and Varró [16] to support attribute constraints. The fundamental idea of translating global constraints into so-called *weakest preconditions* of an algorithm dates back to Dijkstra [19]. Figure 16 shows an overview of the refinement of a GT rule R_x and a graph constraint C_x.

(1) The (global) graph constraint C_y is combined with RHS_x, which results in a set of graph constraints that act as additional *postconditions* of R_x. If the set of postconditions is empty, then R_x already preserves C_y. This is the case for all pairs of GT rules and graph constraints that are labeled with a check mark (✓) in Table 3. Postconditions are similar to application conditions of a GT rule R_x, but they are checked *after* applying R_x. If a postcondition is violated, the application of R_x has to be rolled back by the GT engine. Avoiding this rollback is the purpose of the following transformation (2).

(2) The postconditions are transformed into an equivalent set of application conditions by applying R_x in reverse order to the postconditions. These application conditions block
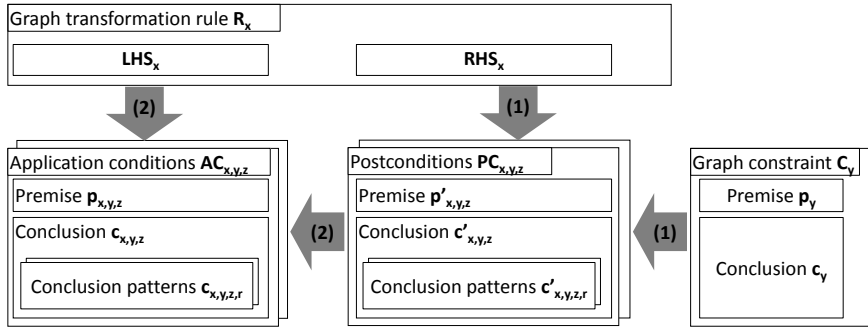
Fig. 16: Refinement procedure for one GT rule R_x and one graph constraint C_y

any application of R_x that would violate its corresponding postcondition. If applying R_x to a particular postcondition in reverse order is impossible, then R_x never violates this postcondition.

The following explanations are deliberately simplified and shown for the following rule-constraint pairs: $(R_a, C_a), (R_i, C_i)$. A detailed, formal description of the following steps can be found in [33, 42]. The presented steps are analogous for all other pairs of GT rules and graph constraints.

First, we define a concept that is crucial for the refinement algorithm: A *gluing* $g_{\ell,r}$ *of two patterns* p_ℓ *and* p_r is a pattern that represents a possible way of overlapping p_ℓ and p_r. We label node variables in gluings with uppercase letters and the original node variable(s) in p_ℓ and p_r. For instance, the node variable $A[x, y]$ of a gluing originates from node variables x of p_ℓ and y of p_r. In a valid gluing, (i) every node (link) variable has one or two original node (link) variables in p_ℓ and/or p_r, (ii) at least one node variable in $g_{\ell,r}$ has original node variables in both p_ℓ and p_r, (iii) the link variable mappings are compatible, i.e., if n_a and n_b are the original node variables of n_A and n_B, then e_{ab} is the original link variable of e_{AB}, and (iv) the attribute constraints of $g_{\ell,r}$ are the conjunction of the attribute constraints of p_ℓ and p_r.

5.1.3 Refinement of Activation Rule R_a Based on the Active-Link Constraint C_a

We begin with the refinement of the activation rule R_a based on the negative active-link constraint C_a. The refinement step can be performed faster for negative graph constraints because its empty conclusion maps to empty conclusions of the postcondition and the application conditions. We obtain the set of postconditions by calculating all twelve gluings $g'_{a,a,z}, 1 \leq z \leq 12$ of the RHS of the activation rule R_a and the premise p_a of the active-link constraint C_a (Figure 17). Each gluing represents a possible violation of C_a. We observe that the nine gluings $g'_{a,a,z}$ for $z \in \{3, 4, 5, 6, 7, 9, 10, 11, 12\}$ represent constraint violations that are *not* caused by activating e_{AB}. This means that any constraint violation corresponding to these gluings already existed before applying the activation rule R_a, which contradicts the assumption that weak consistency is fulfilled prior to invoking any GT rule. Therefore, we neglect the aforementioned gluings and only consider the remaining three gluings $g'_{a,a,1}$, $g'_{a,a,2}$, and $g'_{a,a,8}$ for the transformation from postconditions to application conditions. In step (2), we obtain three NACs, by changing the link-state attribute condition of e_{AB} from
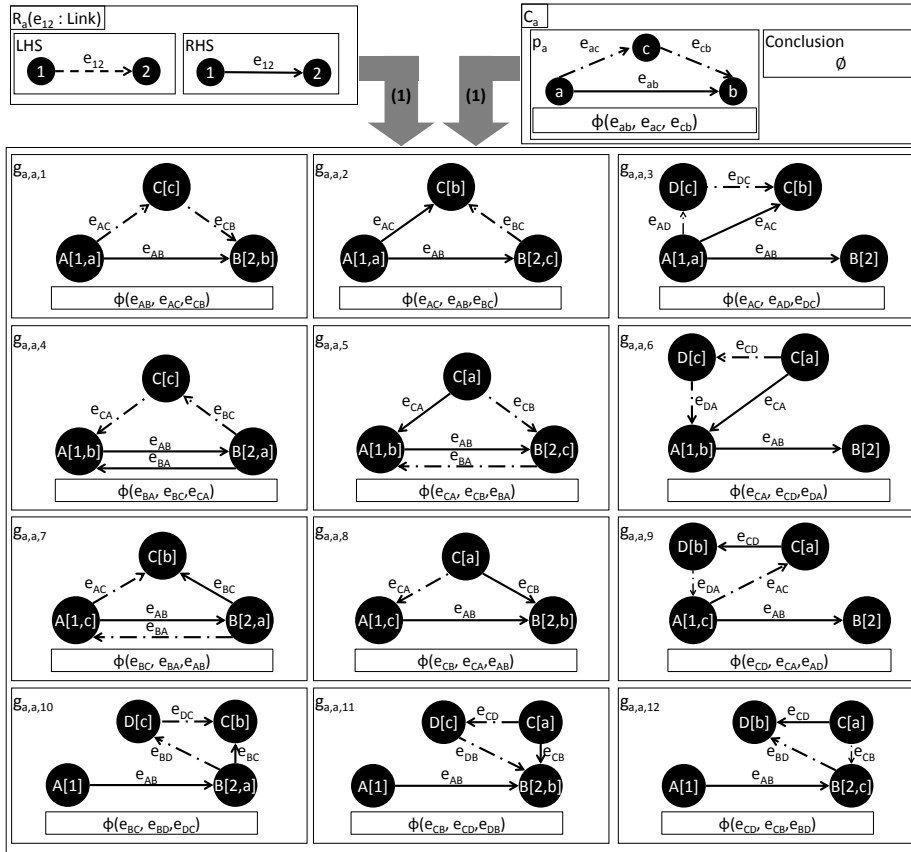
Fig. 17: All twelve possible gluings of RHS$_a$ and p_a

$s(e_{AB}) = \texttt{Active}$ to $s(e_{AB}) = \texttt{Unclassified}$. The resulting application conditions $\text{NAC}_{a,a,1}$, $\text{NAC}_{a,a,2}$, and $\text{NAC}_{a,a,8}$ are shown in Figure 18. Finally, we rename the node variables in the generated application conditions back to the original node variables of the activation rule R_a, as shown in Figure 21b.

Note that $\text{NAC}_{a,a,1}$ prevents the constraint violation shown in Figure 15a, and $\text{NAC}_{a,a,2}$ prevents the constraint violation shown in Figure 15b. Similarly, $\text{NAC}_{a,a,8}$ would prevent a constraint violation if e_{32} (and not e_{13}) were unclassified in Figure 15b.

5.1.4 Refinement of Inactivation Rule R_i Based on the Inactive-Link Constraint C_i

In this second example, we focus on the transformation of the constraint conclusion. The basic idea is that the gluings, which result from combining RHS$_x$ with the p_y, serve as a basis for deriving the conclusion of the postcondition. We obtain the conclusion pattern $c'_{x,y,z,1}$ of the postcondition PC$_{x,y,z}$ by first adding images of all elements in the conclusion of C_y that are not covered by the gluing of RHS$_x$ and p_y. In the original presentation of the constructive approach [33] contains an additional step: The conclusion pattern $c'_{x,y,z,1}$
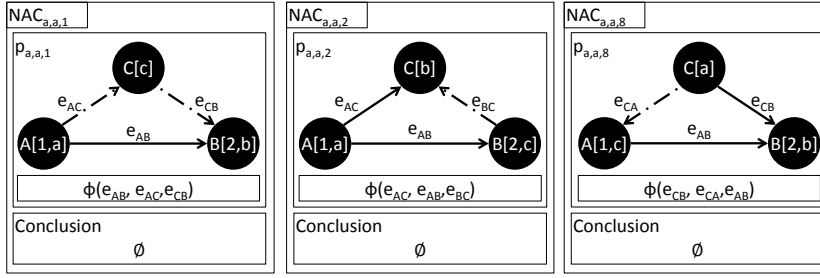
Fig. 18: Application conditions resulting from the refinement of R_a and C_a
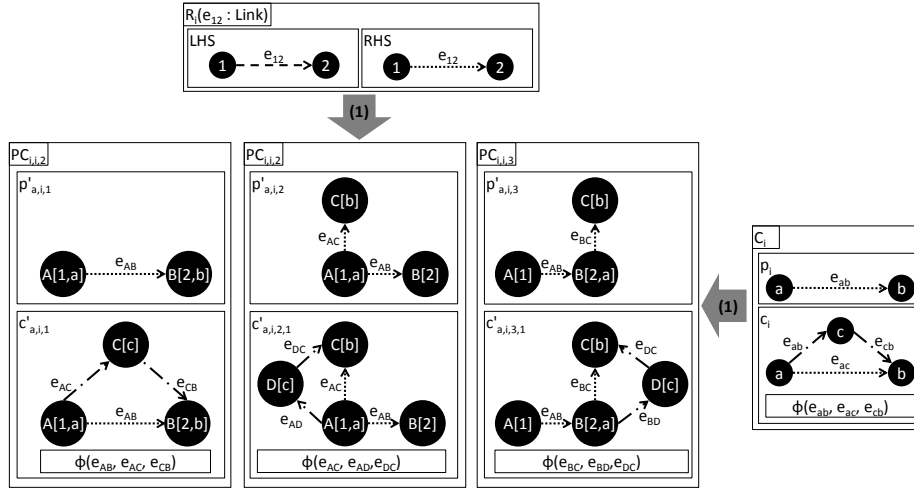
Fig. 19: Derivation of the postconditions $PC_{i,i,1}$, $PC_{i,i,2}$, and $PC_{i,i,3}$ for R_i based on C_i

serves as the basis to generate additional conclusion patterns $c'_{x,y,z,r}, r > 1$ by merging the freshly added node variables with the existing node variables. In our example, this step is not necessary because merging node variables results in loops or parallel links, which never occurs in the considered class of topologies.

As a concrete example, we refine the inactivation rule R_i based on the positive inactive-link constraint C_i. Figure 19 shows the three possible gluings $g'_{i,i,z}, z \in \{1,2,3\}$ of RHS_i the premise p_i of the inactive-link constraint C_i and the corresponding conclusions $c'_{i,i,1,1}$, $c'_{i,i,2,1}$ and $c'_{i,i,3,1}$. For $PC_{i,i,2}$ and $PC_{i,i,3}$, we may stop the refinement procedure here because the possible constraint violations represented by $PC_{i,i,2}$ and $PC_{i,i,3}$ are not caused by the inactivation of e_{AB}. The remaining postcondition $PC_{i,i,1}$ results in a new positive application condition $PAC_{i,i,1}$. Figure 21c shows the resulting application condition of R_a with appropriately renamed variables.
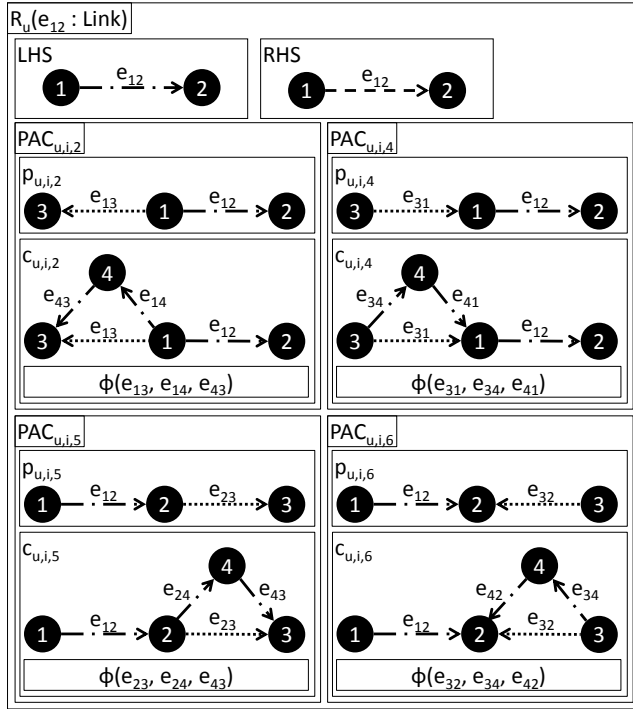
Fig. 20: Unclassification rule R_u after refinement

5.1.5 Refinement of the Remaining Combinations of GT Rules and Constraints

We only sketch the required remaining rule refinements instead of describing them in detail. Examples of additional rule refinement steps can be found in [42].

(i) The refinement of (R_i, C_a) results in two additional negative application conditions $NAC_{i,a,2}$ and $NAC_{i,a,8}$ of the inactivation rule R_i, as shown in Figure 21c, which are identical to the negative application conditions $NAC_{a,a,2}$ and $NAC_{a,a,8}$ of the activation rule R_a. The refined inactivation rule R_i is no longer applicable to the left topology in Figure 15c due to $NAC_{i,a,2}$.

(ii) The refinement of (R_u, C_i) results in four additional positive application conditions $PAC_{u,i,2}$, $PAC_{u,i,4}$, $PAC_{u,i,5}$, and $PAC_{u,i,6}$ (Figure 20). These application conditions require that a link e_{12} may only be unclassified if it is not part of the last match of the conclusion c_i for any incident link e_{13}, e_{31}, e_{23}, or e_{32} of its end nodes n_1 and n_2. The refined unclassification rule R_u is now no longer applicable to the topology in Figure 15d due to $PAC_{u,i,2}$.

(iii) The refinements of (R_{-e}, C_i), (R_{mod-w}, C_i), $(R_{mod-lat}, C_i)$, $(R_{mod-long}, C_i)$, (R_{mod-E}, C_i), and (R_{mod-h}, C_i) result in positive application conditions that are similar to the application conditions of the refined unclassification rule R_u and are not shown here for conciseness.

The specification of the TC algorithm after the rule refinement step is shown in Figure 21a.

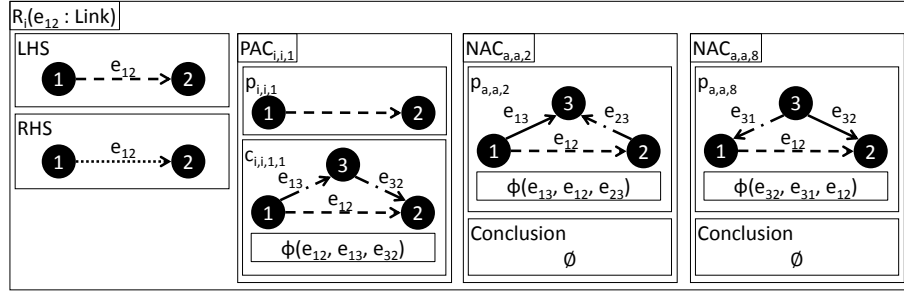(a) Control flow of TC algorithm after refinement (Same as Figure 13)

(b) Refined activation rule R_a

(c) Refined inactivation rule R_i

Fig. 21: Overview of TC algorithm after rule refinement

5.2 Deriving Context Event Handlers

The first refinement step resulted in additional application conditions for TC and context event rules. The refined rules R_u, R_{-e}, R_{mod-w}, $R_{mod-lat}$, $R_{mod-long}$, R_{mod-E}, and R_{mod-h} are now applicable in fewer situations than before. For the context event rules, this is problematic because context event rules represent unrestrictable modifications of the topology caused by the environment. For the unclassification rule R_u, this is problematic because its purpose is to deliberately unclassify links, which should always be possible. Therefore, we have to restore the original applicability of these GT rules *without* sacrificing their constraint-preserving behavior. We propose to transform each added application condition into a *handler operation*, which repairs any constraint violation that results from applying the original context event rule. We will next describe the general idea of deriving handler operations and then illustrate the algorithm for the unclassification rule R_u. The derivation of handler operations of the remaining GT rules is analogous.

Figure 22 shows the structure of a generic context event handler operation `handle-`R_x for the GT rule R_x. The fundamental idea is to (i) first apply the original GT rule R_x (i.e.,

Fig. 22: Structure of the handler operation `handle-`R_x for GT rule R_x

Fig. 23: Handler operation `handle-`R_u for the unclassification rule R_u

without the additional application conditions $AC_{x,1}$, $AC_{x,2}$, ...), and (ii) then to identify and resolve violations of $AC_{x,1}$, $AC_{x,2}$, ... The control flow of the handler operation ensures that it may only terminate if all violations have been resolved. The most important requirement is that the violation resolution strategy shall not produce new constraint violations. In our scenario, we propose to resolve any constraint violation by means of unclassifying links. This approach is valid because a topology consisting exclusively of unclassified links fulfills the inactive-link constraint C_i and the active-link constraint C_a. Therefore, a naïve violation resolution strategy could simply unclassify all links in the topology.

We now derive the concrete handler operation `handle-`R_u for the unclassification rule R_u. During the refinement step, four application conditions have been added to R_u, which are now translated into four identify-and-resolve loops. The invocation of the original, unrefined unclassification rule R_u and the violation resolution for $PAC_{u,i,2}$ are shown in Figure 23. The violation identification rule $R_{\text{find-violation-PAC-u,i,2}}$ identifies any link e_{13} that is *not* part of a triangle together with classified links e_{14} and e_{43} so that the predicate $\phi(e_{13}, e_{14}, e_{43})$ is fulfilled. We unclassify any such link e_{12} by invoking `handle-`R_u recursively. The violation-identifying and violation-resolving story nodes that correspond to $PAC_{u,i,4}$, $PAC_{u,i,5}$, and $PAC_{u,i,6}$ are analogous and omitted here for conciseness.
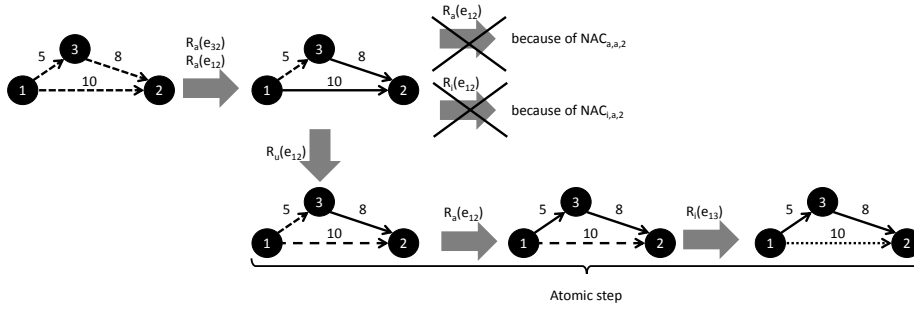
Fig. 24: Example of non-terminating execution of kTC with $k = 2$ and possible solution

5.3 Enforcing Termination

The refined activation rule R_a and inactivation rule R_i share two pairs of identical NACs. This means that whenever $NAC_{a,a,2}$ prevents the application of the activation rule R_a, $NAC_{i,a,2}$ also prevents the application of the inactivation rule R_i. The same holds for $NAC_{a,a,8}$ and $NAC_{i,a,8}$. In Figure 24, the top row of topology modifications shows a situation where a particular processing order of the unclassified links causes non-termination of the algorithm. In this case, kTC is executed with $k = 2$, and e_{32} and e_{12} are activated prior to processing e_{13}. Now, e_{13} can be neither activated nor inactivated due to the negative application conditions $NAC_{a,a,2}$ and $NAC_{i,a,2}$, respectively. This causes an infinite execution of the algorithm.

The situation can be solved by reverting link classifications. In this example, a possible solution can be derived from the only possible strongly consistent output topology: The link e_{13} and e_{32} are active, and e_{12} is inactive. This means that e_{12} should be unclassified again and e_{13} should be activated. The two rule applications have to happen atomically, i.e., the activation of e_{13} must follow immediately after the unclassification of e_{12}. This order can easily be implemented using programmed GT. Finally, e_{12} becomes inactive.

In fact, the solution for the example can be generalized, again using handler operations. We systematically transform the shared negative application conditions of the activation rule R_a and the inactivation rule R_i into an appropriate NAC-handling operation `handle-NAC`$_{aa2,ia2,aa8,ia8}$, which destroys all matches of the premises of the aforementioned four NACs (Figure 25). Contrary to the context event handler operations, we decided to place the NAC-handling operation in front of R_a and R_i because, otherwise, we would have to add invocations of the NAC-handling operation in both [Success]-branches.

Inside the NAC-handling operation, the first loop identifies all matches of the premise of $NAC_{a,a,2}$ and $NAC_{i,a,2}$ and unclassifies the link e_{13}, which is the largest link in the triangle w.r.t. $\prec_A$. Similarly, the second loop identifies all matches of the premises of $NAC_{a,a,8}$ and $NAC_{i,a,8}$ and unclassifies e_{32}. Finally, we drop $NAC_{a,a,2}$ and $NAC_{a,a,8}$ from R_a and $NAC_{i,a,2}$ and $NAC_{i,a,8}$ from R_i.

Proving termination Prior to introducing the NAC-handling operation, the number of iterations of the main loop of the TC algorithm was limited by the initial number of unclassified links. Now, additional links may become unclassified in each iteration, which requires us to prove that the current TC algorithm terminates.
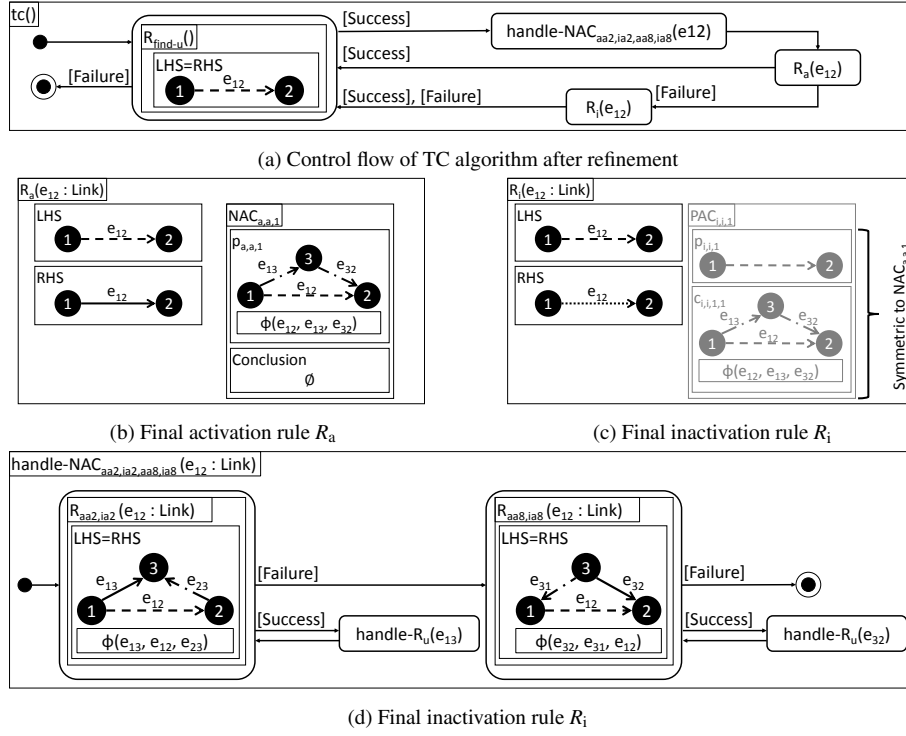
(a) Control flow of TC algorithm after refinement

(b) Final activation rule R_a

(c) Final inactivation rule R_i

(d) Final inactivation rule R_i

Fig. 25: Final specification of the TC algorithm

Theorem 3 *The TC algorithm A with NAC-handling operation (Figure 25) terminates for any input topology.*

Proof. Let E be the link set of the processed topology. Let A be the considered TC algorithm, represented by the predicate ϕ_A. We consider the sequence of all link states $s_x(e_1), \ldots, s_x(e_m)$ with $m := |E|$ after the x-th execution of the unclassified-link-identification rule R_{find-u}, where the links are ordered according to $\prec_A$. We compare two sequences of link states, s_x and $s_{x'}$, as follows: $s_x \sqsubset s_{x'}$ if and only if (i) some link e_y is unclassified in s_x and classified in $s_{x'}$, and (ii) the states of all links $e_{y'}$ with $e_{y'} \prec_A e_y$ are identical in s_x and $s_{x'}$:

$$s_x \sqsubset s_{x'} \; :\Leftrightarrow \exists y \in \mathbb{N}, y \le m : s_x(e_y) = \mathtt{Unclassified} \wedge s_{x'}(e_y) \in \{\mathtt{Active}, \mathtt{Inactive}\}$$
$$\wedge \; \forall y' \in \mathbb{N}, y' < y : s_x(e_{y'}) = s_{x'}(e_{y'})$$

Any sequence of active and inactive links is an upper bound for $\sqsubset$.

We now show that $s_{x-1} \sqsubset s_x$ for $x > 1$. Let e_y be the link that is bound by applying the unclassified-link-identification rule R_{find-u}. The NAC-handling operation unclassifies links $e_{y'}$ with $e_{y'} \prec_A e_y$ and thus $y' < y$. The activation rule R_a or the inactivation rule R_i activate or inactivate e_y, respectively.

Therefore, $s_{x-1} \sqsubset s_x$ because (i) the first $y-1$ elements of s_{x-1} and s_x are identical, and (ii) $s_{x-1}(e_y) = \mathtt{Unclassified}$ and $s_x(e_y) \in \{\mathtt{Active}, \mathtt{Inactive}\}$. The termination follows because, for a finite topology, any ordered sequence $s_1 \sqsubset s_2 \sqsubset \ldots$ has finite length.

$\square$

In this section, we dropped two of three application conditions of both the activation rule R_a and inactivation rule R_i. The remaining application conditions are $NAC_{a,a,1}$ for R_a and $PAC_{i,i,1}$ for R_i. $NAC_{a,a,1}$ and $PAC_{i,i,1}$ are complementary in the sense that $PAC_{i,i,1}$ is fulfilled if and only if $NAC_{a,a,1}$ is not fulfilled for a link e_{12}. This means that we may also drop $PAC_{i,i,1}$ from R_i without risking to compromise weak consistency. Removing $PAC_{i,i,1}$ is sensible because the resulting inactivation rule R_i is more efficient to evaluate and, arguably, easier to understand.

6 Evaluation

In this section, we present a comparative simulation-based evaluation study that serves as a proof-of-concept for our integration of the GT tool EMOFLON [47] and the network simulator SIMONSTRATOR [61]. In Section 6.1, we explain the research questions of this evaluation. In Section 6.2, we describe the technical setup and the configuration parameters. In Section 6.3, we describe the cost and utility metrics to answer the research questions. In Section 6.4, we present the measurement results, discuss them, and answer the research questions. In Section 6.5, we discuss threats to the validity of our findings.

6.1 Research Questions

Our goal is to investigate the benefits of the proposed approach for specifying families of TC algorithms by composing their corresponding predicates. For the sake of conciseness, we investigate the TC algorithms kTC and e-kTC and the effect of combining each of the algorithms with the minimum-weight predicate $\phi_{\text{min-weight}}$.

RQ1—Performance of e-kTC The first research question addresses the performance of e-kTC compared to kTC. As described in Section 2.4.10, the major objective of e-kTC is to extend the lifetime of the topology, which is achieved by establishing a fairer distribution of the per-node energy consumption. Our first research question is: Does e-kTC improve the network lifetime compared to kTC?

RQ2—Benefit of Minimum-Weight Predicate The second research question addresses the performance of applying the minimum-weight predicate $\phi_{\text{min-weight}}$. A motivation for introducing $\phi_{\text{min-weight}}$ in Section 2.4.11 was to reduce the cost of a TC algorithm in terms of memory and runtime while preserving the quality of the output topology. Our second research question is: How does w_{thres} influence the network lifetime and the resource consumption of kTC and e-kTC?

6.2 Evaluation Setup

Following best practices for simulation studies [35, 46], we rigorously document the evaluation setup to foster reproducibility of our results. The technical platforms of this evaluation are the GT tool EMOFLON [47] and the network evaluation platform SIMONSTRATOR [61] with its contained network simulator PEERFACTSIM.KOM [73], as shown in Figure 26 and as described in detail in the following. A SHARE virtual machine [29] is available for exploring our tool integration.[5]

---

[5] See https://github.com/Echtzeitsysteme/CorrectByConstructionTCFamilies-SoSyM17
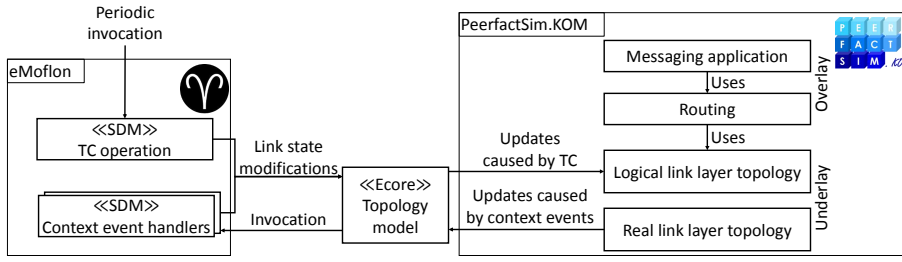
Fig. 26: Overview of the evaluation setup

GT tool All TC algorithms have been specified using a visual syntax of story diagrams in EMOFLON [47]. The specification is used to generate EMF-compliant Java code, which builds on an Ecore-based [72] topology metamodel. Whenever the topology model is modified by a context event in the simulator, the corresponding handler operation is triggered. The TC algorithm periodically updates the topology model via link state modifications.

Network simulator PEERFACTSIM.KOM [73] is a time-discrete network simulator that allows to simulate protocols on the underlay (i.e., the physical and link layer [87]) as well as on the overlay (i.e., network, transport and application layer [87]).

Node configuration The nodes are initially placed uniformly at random onto the simulation area. Each node moves according to the Gauss-Markov movement model [12]. As link weight, we use the Euclidean distance of its end nodes. A calibration curve estimates the required transmission power of each link based on its weight. To create the calibration curve, we used a topology consisting of a sender and a receiver node and set up a simple data transmission application that transfers data from the sender to the receiver at a rate of $1 \frac{\text{MB}}{\text{s}}$.

Underlay We implemented a generic topology control framework inside the simulator. This framework presents a *logical link layer topology* to overlay applications, which consists of all active links in the *real link layer topology*. The real link layer topology is updated by the movement model and the energy model, which describes the battery depletion of each node. In this evaluation, TC is invoked periodically every 10 min of simulated time. The set of context events is collected during this period and handled before invoking TC. A *TC run* consists of the context event handling and subsequent TC invocation. The parameter k of kTC and e-kTC is set to 1.41. This is a typical value [45, 69, 70, 71] because it is the largest value for k that allows to prove certain advanced properties (such as θ-separation [69]). The expected lifetime $\widehat{\text{L}}(e_{\text{ab}})$ of each link e_{ab} is calculated by dividing the energy level $\text{E}(n_{\text{a}})$ of n_{a} by the expected transmission power $\widehat{\text{P}}(e_{\text{ab}})$ of e_{ab}.

Overlay On the overlay, we run a messaging application, which emulates sporadic message exchange between random pairs of nodes. Every 30 s, each node sends a message of size 1 kB to another node that is selected randomly. We apply a routing algorithm that builds routing paths based on global knowledge.

Table 4: Parameters of the evaluation

Parameter	Value
Simulator	
Network simulator	SIMONSTRATOR/PEERFACTSIM.KOM 2.4 [61, 73]
Graph transformation tool	EMOFLON 2.16.0 [47]
Runs per configuration	5
Seeds	1...5
Simulation type	Terminating (after 25 h of simulated time)
Measurements per run	150
Code availability	As SHARE [29] virtual machine[a]
World size	*Variable* (see Table 5)
Node Configuration	
Node count	100
Transmission radius	130 m
Movement model	Gauss-Markov movement [12] ($\alpha = 0.2$, $v = 0.005 \frac{m}{s}$)
Placement model (initial)	Uniform at random
Battery level (initial)	Uniform at random betw. 30 % and 100 % of 130 J
Underlay	
Link count (initial)	*Variable* (see Table 5)
Link layer	IEEE 802.11 Ad-Hoc
TC algorithm	kTC, e-kTC
$w(e_{ab})$	Euclidean distance between n_a and n_b
$\widehat{L}_1(e_{ab})$	Based on calibration curve
k parameter of kTC	1.41
Interval of TC runs	10 min
TC runs per simulation run	150
Overlay	
Application	Messaging application
Communication pattern	Many-to-many
Message size	1 kB
Messaging interval	30 s

[a] See https://github.com/Echtzeitsysteme/CorrectByConstructionTCFamilies-SoSyM17

Summary of parameters Table 4 summarizes the parameters of the simulation setting that are fixed for all configurations. We run each simulation in *terminating mode*, i.e., we abort the simulation after 25 h, regardless of the remaining number of alive nodes. This means that, in each simulation run, TC is performed $\frac{25\,h}{10\,min} = 150$ times. All experiments were performed on a 64-bit machine with an Intel i7-2600 CPU (2 cores, 3.7 GHz) and 8 GB of RAM running Windows 7 Professional.

To obtain a set of representative *(topology) configurations*, we vary the side length of the quadratic simulation area. In total, we investigate 2 different evaluation configurations, summarized in Table 5, which correspond to world sizes $a \in \{750\,m, 500\,m\}$. Each configuration is simulated 5 times to obtain representative results. In the table, the configurations are ordered from the sparsest (n100w750) to the densest topology (n100w500). The node (out-)degree d_{out} is an indicator of the density of the topology: A high node degree indicates that a topology is dense, while a low node degree indicates a sparse topology.

Table 5: Varying parameters in the two topology configurations (n: initial node count, m: initial link count in real link layer, a: side length of quadratic simulation area, d_{out}: average out-degree, values of m, h, and d_{out} are averaged over 5 repetitions)

ID	n	a [m]	m	d_{out}
n100w750	100	750	812	7.7
n100w500	100	500	1 651	16.0

6.3 Metrics

The major goal of this evaluation is to assess cost and utility metrics of the output topology of a TC algorithm. With $G_{w_{\text{thres}}}$, we denote the output topology of the currently active TC algorithm, whose input topology is obtained by removing all links with a weight smaller than w_{thres} from the physical topology.

As utility metric, we consider the network lifetime (see also Section 2.4.10). We record for each simulation run the remaining 1-lifetime L_1, the remaining 50 %-lifetime $L_{50\%}$, and remaining 100 %-lifetime $L_{100\%}$. These values mark the starting point, the approximate middle, and the end of the degradation of the network.

As cost metric, we assess the required memory in terms of the size of the maintained input topology, and the required runtime in terms of the required CPU execution time as well as the number of link state modifications for performing the TC algorithm. The size $S(G_{w_{\text{thres}}})$ of the topology $G_{w_{\text{thres}}}$ is the sum of its node and edge counts. The *execution time* $T(G_{w_{thres}})$ is the real time (in contrast to the simulated time) that it takes to perform a TC run on $G_{w_{\text{thres}}}$. The *link state modification count* $N(G_{w_{thres}})$ is determined analogously by counting the link state modifications during a TC run on $G_{w_{\text{thres}}}$.

We compare results along two dimensions: Either we consider a fixed configuration (i.e., n100w500 or n100w750) and compare the performance of the TC algorithms, or we fix the world size, node count, TC algorithm and k-value and evaluate the influence of varying w_{thres}. For each dimension, we introduce relative metrics δ_M for each absolute metric M. When comparing a TC algorithm A with the baseline Maxpower TC algorithm, δ_M is defined as follows:

$$\delta_{L_1}(A) = \frac{L_1(A)}{L_1(\text{Maxpower TC})}, \qquad \delta_{L_{50\%}}(A) = \frac{L_{50\%}(A)}{L_{50\%}(\text{Maxpower TC})},$$

$$\delta_{L_{100\%}}(A) = \frac{L_{100\%}(A)}{L_{100\%}(\text{Maxpower TC})}, \qquad \delta_{\text{Size}}(A) = \frac{S(A)}{S(\text{Maxpower TC})},$$

$$\delta_T(A) = \frac{T(A)}{T(\text{Maxpower TC})}, \qquad \delta_N(A) = \frac{N(A)}{N(\text{Maxpower TC})}.$$

When evaluating the influence of varying w_{thres}, δ_M is defined as follows:

$$\delta_{\text{L}_1}(G_{w_{\text{thres}}}) = \frac{\text{L}_1(G_{w_{\text{thres}}})}{\text{L}_1(G_0)}, \qquad\qquad \delta_{\text{L}_{50\%}}(G_{w_{\text{thres}}}) = \frac{\text{L}_{50\%}(G_{w_{\text{thres}}})}{\text{L}_{50\%}(G_0)},$$

$$\delta_{\text{L}_{100\%}}(G_{w_{\text{thres}}}) = \frac{\text{L}_{100\%}(G_{w_{\text{thres}}})}{\text{L}_{100\%}(G_0)}, \qquad\qquad \delta_{\text{Size}}(G_{w_{\text{thres}}}) = \frac{S(G_{w_{\text{thres}}})}{S(G_0)},$$

$$\delta_T(G_{w_{\text{thres}}}) = \frac{T(G_{w_{\text{thres}}})}{T(G_0)}, \qquad\qquad \delta_N(G_{w_{\text{thres}}}) = \frac{N(G_{w_{\text{thres}}})}{N(G_0)}.$$

Setting w_{thres} to 0 is equivalent to disabling the minimum-weight optimization.

6.4 Results and Discussion

The plots in Figures 27 and 28 show the development of the number of alive nodes for the dense topology (Figure 27) and the sparse topology (Figure 28) for different minimum-weight thresholds $w_{\text{thres}} \in \{0, 20, 40, 60, 80\}$. In each plot, the simulation runs of Maxpower TC, kTC, and e-kTC are shown.

While the plots provide a rather qualitative view of the performance of the TC algorithms, Tables 6 and 7 allow for a fine-grained analysis of the cost and utility metrics, as introduced in Section 6.3. We provide mean values of all 150 TC runs per simulation for the topology size ($\overline{S}$), the execution time ($\overline{T}$), and the link state modification count ($\overline{N}$).

Network lifetime We observe that kTC almost always performs worst and e-kTC performs best of all three TC algorithms w.r.t. the number of alive nodes. More precisely, for a fixed value of w_{thres}, the lifetime of $G_{w_{\text{thres}}}$ for kTC is always shorter than for e-kTC (see also Tables 6 and 7). Also, the remaining lifetime of kTC increases with increasing w_{thres}. For instance, all network lifetime values of the sparse topology are maximal when applying kTC with $w_{\text{thres}} = 80\,\text{m}$ ($\text{L}_1 = 556\,\text{min}$ compared to $\text{L}_1 = 532\,\text{min}$, 5 % improvement, Figure 28). A similar, yet not as strong, effect can be observed for e-kTC. For instance, in the sparse topology, $\text{L}_{100\%}$ for $w_{\text{thres}} = 20\,\text{m}$ is about 1 % larger than $\text{L}_{100\%}$ for $w_{\text{thres}} = 0\,\text{m}$. In the remaining cases, the network lifetime for e-kTC is maximal for $w_{\text{thres}} = 0\,\text{m}$.

Resource consumption We now focus on resulting topology size, real execution time, and required link state modification count. These cost metrics evaluate the resource consumption of a TC algorithm. In contrast to the network lifetime, all of the cost metrics strictly decrease with w_{thres} and reach their minimum for $w_{\text{thres}} = 80\,\text{m}$.

For the dense topology (`n100w500`, Figure 27), the size of the stored topology was reduced by up to 22 % (for kTC) and 33 % (for e-kTC). The savings in execution time are even more drastic: 80 % for kTC and 73 % for e-kTC. Finally, the savings w.r.t. link state modification count amount up to 72 % for kTC and 57 % for e-kTC.

For the sparser topology (`n100w750`, Figure 28), the savings w.r.t. topology size are comparable (20 % for kTC and 26 % for e-kTC). The savings w.r.t. execution time (62 % for kTC and 40 % for e-kTC) and link state modification count (59 % for kTC and 52 % for e-kTC) are smaller.

(a) $w_{\text{thres}} = 0.0\,\text{m}$

(b) $w_{\text{thres}} = 20.0\,\text{m}$

(c) $w_{\text{thres}} = 40.0\,\text{m}$

(d) $w_{\text{thres}} = 60.0\,\text{m}$

(e) $w_{\text{thres}} = 80.0\,\text{m}$

Fig. 27: Alive nodes over simulation time (World size: 500 m, node count: 100, k: 1.41)

(a) $w_{\text{thres}} = 0.0\,\text{m}$

(b) $w_{\text{thres}} = 20.0\,\text{m}$

(c) $w_{\text{thres}} = 40.0\,\text{m}$

(d) $w_{\text{thres}} = 60.0\,\text{m}$

(e) $w_{\text{thres}} = 80.0\,\text{m}$

Fig. 28: Alive nodes over simulation time (World size: 750 m, node count: 100, k: 1.41)

Table 6: Algorithm performance for different weight thresholds w_{thres} (World size: 500 m, node count: 100, k: 1.41). Lifetime values (L_1, $L_{50\%}$, $L_{100\%}$) are in simulated minutes. Execution time $\overline{T}$ is in CPU milliseconds. δ_M: relative metrix M. The best values for each combination of algorithm, w_{thres}, and metric M are highlighted in bold font.

Algo.	w_{thres} [m]	L_1 [min]	δ_{L_1}	$L_{50\%}$ [min]	$\delta_{L_{50\%}}$	$L_{100\%}$ [min]	$\delta_{L_{100\%}}$
Maxp.	0	564.0	1.00	700.0	1.00	854.0	1.00
kTC	0	544.0	1.00	664.0	1.00	862.0	1.00
kTC	20	544.0	1.00	664.0	1.00	864.0	1.00
kTC	40	570.0	1.05	684.0	1.03	**868.0**	**1.01**
kTC	60	**580.0**	**1.07**	698.0	1.05	862.0	1.00
kTC	80	564.0	1.04	**702.0**	**1.06**	860.0	1.00
e-kTC	0	**624.0**	**1.00**	**776.0**	**1.00**	**906.0**	**1.00**
e-kTC	20	**624.0**	**1.00**	770.0	0.99	902.0	1.00
e-kTC	40	616.0	0.99	752.0	0.97	**906.0**	**1.00**
e-kTC	60	600.0	0.96	726.0	0.94	882.0	0.97
e-kTC	80	590.0	0.95	712.0	0.92	864.0	0.95

Algo.	w_{thres} [m]	$\overline{S}$	δ_{Size}	$\overline{T}$ [ms]	δ_T	$\overline{N}$	δ_N
Maxp.	0	822.3	1.00	19.8	1.00	18.1	1.00
kTC	0	774.7	1.00	253.4	1.00	1340.3	1.00
kTC	20	775.0	1.00	227.6	0.90	1252.2	0.93
kTC	40	768.7	0.99	147.8	0.58	998.7	0.75
kTC	60	712.7	0.92	83.1	0.33	728.0	0.54
kTC	80	**603.2**	**0.78**	**49.9**	**0.20**	**504.7**	**0.38**
e-kTC	0	927.4	1.00	240.7	1.00	1303.1	1.00
e-kTC	20	912.9	0.98	211.6	0.88	1218.0	0.93
e-kTC	40	853.5	0.92	151.8	0.63	1031.2	0.79
e-kTC	60	747.3	0.81	115.8	0.48	791.2	0.61
e-kTC	80	**622.2**	**0.67**	**64.0**	**0.27**	**555.2**	**0.43**

Trade-off between cost and utility Surprisingly, the lifetime when applying kTC was generally very good for $w_{\text{thres}} = 80$ m. Therefore, we do not have to trade network lifetime against resource consumption in this situation. Instead, we can optimize resource consumption and network lifetime at the same time by choosing high values of w_{thres}. This observation underpins the intuition behind the minimum-weight predicate that considering short links does probably not help increasing the performance of the network. Our results even show that performing kTC also on short links may be harmful to the network lifetime. A possible explanation is that short links are probably helpful when used for direct message transfer (compared to a multi-hop forwarding of the same message). Despite this positive result, we have to remark that kTC is often not able to beat the baseline, i.e., Maxpower TC, w.r.t. network lifetime.

In contrast, e-kTC shows a rather monotonic behavior: In most cases, the network lifetime decreases with increasing w_{thres}. Here, we trade the reduced resource consumption against the shorter network lifetime. In the most extreme case ($w_{\text{thres}} = 80$ m), the lifetime drops to values between 92 % and 95 % (for the dense topology) and between 96 % and 99 % (for the sparse topology) of its maximum value. Choosing the best trade-off between network lifetime and resource consumption for e-kTC is certainly a task that depends on the available resources of the selected (simulated) target platform.

Table 7: Algorithm performance for different weight thresholds w_{thres} (world size: 750 m, node count: 100, k: 1.41). Lifetime values (L_1, $L_{50\%}$, $L_{100\%}$) are in simulated minutes. Execution time $\overline{T}$ is in CPU milliseconds. δ_M: relative M compared to same algorithm with $w_{\text{thres}} = 0$. The best values for each combination of algorithm, w_{thres}, and metric M are highlighted in bold font.

Algo.	w_{thres} [m]	L_1 [min]	δ_{L_1}	$L_{50\%}$ [min]	$\delta_{L_{50\%}}$	$L_{100\%}$ [min]	$\delta_{L_{100\%}}$
Maxp.	0	552.0	1.00	668.0	1.00	884.0	1.00
kTC	0	532.0	1.00	646.0	1.00	906.0	1.00
kTC	20	532.0	1.00	650.0	1.01	942.0	1.04
kTC	40	546.0	1.03	656.0	1.02	892.0	0.98
kTC	60	554.0	1.04	**670.0**	**1.04**	894.0	0.99
kTC	80	**556.0**	**1.05**	**670.0**	**1.04**	**956.0**	**1.06**
e-kTC	0	**592.0**	**1.00**	**710.0**	**1.00**	958.0	1.00
e-kTC	20	**592.0**	**1.00**	702.0	0.99	**966.0**	**1.01**
e-kTC	40	590.0	1.00	696.0	0.98	902.0	0.94
e-kTC	60	584.0	0.99	688.0	0.97	958.0	1.00
e-kTC	80	574.0	0.97	684.0	0.96	944.0	0.99

Algo.	w_{thres} [m]	$\overline{S}$	δ_{Size}	$\overline{T}$ [ms]	δ_T	$\overline{N}$	δ_N
Maxp.	0	412.9	1.00	8.7	1.00	9.0	1.00
kTC	0	401.5	1.00	36.5	1.00	604.9	1.00
kTC	20	399.9	1.00	34.3	0.94	567.2	0.94
kTC	40	390.8	0.97	27.7	0.76	473.4	0.78
kTC	60	366.4	0.91	22.0	0.60	361.3	0.60
kTC	80	**319.7**	**0.80**	**17.6**	**0.48**	**245.3**	**0.41**
e-kTC	0	443.6	1.00	32.8	1.00	563.0	1.00
e-kTC	20	435.9	0.98	31.5	0.96	532.9	0.95
e-kTC	40	418.7	0.94	30.7	0.93	470.1	0.83
e-kTC	60	383.8	0.87	25.8	0.79	383.2	0.68
e-kTC	80	**329.1**	**0.74**	**19.8**	**0.60**	**271.7**	**0.48**

Table 8 summarizes the best performance values of Maxpower TC, kTC, and e-kTC in Tables 6 and 7. In this table, δ_M denotes the relative value of metric M compared to the baseline Maxpower TC algorithm. e-kTC achieves network lifetime improvements between 6.0 % and 10.9 % for the dense topology (n100w500) and between 6.3 % and 9.3 % for the sparse topology (n100w750). At the same time, kTC only achieves improvements between 0.3 % and 1.6 % for the dense topology (n100w500) and between 0.3 % and 8.1 % for the sparse topology (n100w750). The execution times of kTC and e-kTC grow by a factor of two to three compared to the execution time of Maxpower TC, and the number of required link state modifications even increases by a factor of 25 to 30. This is not at all surprising because the logic behind Maxpower TC is simple compared to the rules of kTC and e-kTC. Especially handling context events results in additional required link state modifications. When comparing kTC with e-kTC w.r.t. cost metrics, kTC outperforms e-kTC. Its best execution time is 22.0 % (for n100w500) and 11.1 % (for n100w750) lower than the best execution time of e-kTC. Similarly, kTC outperforms e-kTC by 9.1 % (for n100w500) and 9.7 % (for n100w750) w.r.t. the best achieved link state modification count. However, the configuration that achieves the best execution time and link state modification count for e-kTC perform worse than the best-performing configuration of kTC w.r.t. network lifetime. Therefore, an

Table 8: Performance of kTC and e-kTC compared to Maxpower TC

Config.	Algo.	L_1 [min]	δ_{L_1}	$L_{50\%}$ [min]	$\delta_{L_{50\%}}$	$L_{100\%}$ [min]	$\delta_{L_{100\%}}$
n100w500	Maxp.	564.0	–	700.0	–	854.0	
n100w500	kTC	580.0	+2.8 %	702.0	+0.3 %	868.0	+1.6 %
n100w500	e-kTC	624.0	+10.6 %	776.0	+10.9 %	906.0	+6.0 %
n100w750	Maxp.	552.0	–	668.0	–	884.0	
n100w750	kTC	556.0	+0.7 %	670.0	+0.3 %	956.0	+8.1 %
n100w750	e-kTC	592.0	+7.2 %	710.0	+6.3 %	966.0	+9.3 %

Config.	Algo.	$\overline{S}$	δ_{Size}	$\overline{T}$ [ms]	δ_T	$\overline{N}$	δ_N
n100w500	Maxp.	822.3	–	19.8	–	18.1	
n100w500	kTC	603.2	−26.6 %	49.9	+152 %	504.7	+2688 %
n100w500	e-kTC	622.2	−24.3 %	64.0	+223 %	555.2	+2967 %
n100w750	Maxp.	412.9	–	8.7	–	9.0	
n100w750	kTC	319.7	−22.6 %	17.6	+102 %	245.3	+2626 %
n100w750	e-kTC	329.1	−20.3 %	19.8	+128 %	271.7	+2919 %

immediate comparison of kTC and e-kTC is not salient. Instead, a trade-off between cost and utility is necessary also here.

Answer to RQ1 Our answer to the first research question is positive: e-kTC is able to beat kTC with respect to network lifetime: For the dense and the sparse topology, the best configurations of e-kTC w.r.t. L_1, $L_{50\%}$, and $L_{100\%}$ outperform kTC in terms of execution time and link state modification count, but not in terms of topology size. Additionally, kTC is in general only able to beat the baseline Maxpower TC algorithm for large w_{thres} between 40 m and 80 m, while e-kTC generally performs better than Maxpower TC.

Answer to RQ2 Our answer to the second research question is that varying w_{thres} influences the lifetime to a certain extent (positively for kTC, negatively for e-kTC), while the savings w.r.t. resource consumption are considerable. A key observation is that the network lifetime for kTC tends to increase with increasing w_{thres}, while the network lifetime for e-kTC correlates negatively with w_{thres}

6.5 Threats to Validity

We first consider threats to *external validity*, i.e., the ability to generalize our findings in this evaluation to WSN topologies and TC algorithms in general. To mitigate this threat, we analyze different types of topologies: sparse and dense topologies. Our results (Section 6.4) are relatively homogeneous for all considered topologies (e.g., the relative performance of the algorithms). This increases our confidence in the validity of these findings. To strengthen our confidence, we plan to carry out additional experiments, e.g., with different movement models, placement models and node density to obtain additional representative data as future work. Managing the fast-growing number of possible configurations is a major challenge.

Another degree of freedom is the choice of k for kTC and e-kTC. The selection of $k = 1.41$ may be the unfortunate reason for the suboptimal performance of kTC. To mitigate this concern, we performed additional simulation runs with $k \in \{1.0, 1.2, 1.3, 2.0\}$ and for

a larger configuration with a world size of 2000 m and 1 000 sensor nodes, which are not discussed here for conciseness. The results of these preliminary experiments are comparable to the results discussed earlier. kTC generally tends to perform worse than e-kTC w.r.t. network lifetime, and the performance of kTC correlates positively with w_{thres}.

The limitation to only one overlay application, i.e., the random-pair communication application, is another threat to external validity. While this application reflects only one particular communication pattern, i.e., the many-to-many pattern, this pattern is relatively widespread in the communication systems domain. We plan to conduct experiments that cover communication patterns such as many-to-one, which is typical in data collection scenarios, or one-to-many, which is applied for data dissemination.

The last considered threat to external validity is that we model topologies as simple graphs (see Section 2.1). Modeling other types of topologies apart from WSN topologies will probably require to support loops and parallel links. To support parallel links and loops, we would drop the structural constraints $C_{no\text{-}loops}$ and $C_{no\text{-}par\text{-}links}$. This implies that the refinement step (see Section 5) may result in additional application conditions because less gluings may be dropped. Therefore, our approach is not limited to simple graphs.

A major threat to *internal validity* is the reproducibility of our results. To address this threat, we ran all experiments corresponding to each configuration (i.e., a fixed combination of world size, node count, algorithm, *k*-parameter, and minimum-weight threshold) at least 5 times with different random seeds and used the average values of these runs. We could further increase our confidence in the internal validity of our evaluation by increasing the number of random seeds and by performing an in-depth analysis of the variance within the data of each configuration.

7 Related Work

In this section, we survey related work of this paper with a focus on the development of correct algorithms, model-based software development and software product lines.

7.1 Correctness of Algorithms

One motivation of our work is the pertinent missing traceability between specification and implementation in the communication systems domain (see S1 in Section 1). The authors of [49] propose a calculus for WSN protocols that is the basis for a service-oriented middleware called *MufFIN* [78]. The idea behind their approach was to prove required properties based on a specification of a WSN protocol in terms of the calculus. Then, an equivalent byte code representation for the MufFin middleware should be generated. Unfortunately, it appears that the approach has not yet been fully implemented to showcase its applicability. In [51], the authors present a domain-specific language for WSN protocols. While their approach is not limited to TC algorithms, this generality makes it hard to constructively ensure required properties, which is a key objective of our work. In [20], the *ARESA* project is presented, an alliance of industrial and academic institutions that aims to tackle WSN research challenges jointly. One of the key objectives is formal analysis of WSN protocols. In contrast to this paper, the authors lay a focus on verifying properties based on a formal specification of WSN protocols. In contrast, our approach is to construct TC algorithms, i.e., specialized WSN protocols, that are correct by construction. In [57], the authors identify the automatic synthesis of implementations as a central research area because the verification of

protocols with large state space is still intractable today. Finally, a detailed anecdotal example of missing traceability between specification and implementation can be found in [42, Sec. 1]. There, we analyze the presentation of the CTCA algorithm [15] and highlight the existing gaps. This does *not* mean that the evaluated implementation is incorrect. The key message of this example is rather that it is at least difficult for the reader of [15] to understand the relation between the game-theoretic specification and the extensive pseudo code implementation, let alone the unpublished source code of the simulation study.

One of the key properties of the specified TC algorithms presented in this paper is that they are guaranteed to be correct by construction. In general, at least three major approaches exist for asserting or checking the correctness of software: First, *correct-by-construction approaches* integrate the properties during the construction of the software, which is the fundamental idea of this paper. Second, *verification-based approaches* examine an existing piece of software with respect to the required properties. Third, *testing* exercises the algorithm by executing a set of representative test cases, each consisting of input data and the expected result and, for each test case, comparing the actual result with the expected result. In [31], the authors highlight that combining these approaches is useful in realistic projects. Correct-by-construction approaches have been extensively studied in the context of hardware design and software development processes (e.g., [6, 7, 56]). However, to the best of our knowledge, no work (apart from ours [42, 43]) exists on constructing TC algorithms using the correct-by-construction paradigm. One technique for verifying correctness properties of software is model checking [60]. Many model checking approaches are only suitable for verifying properties for models of a finite size. Graph abstractions [5] are a formalism that alleviates this problem by introducing symbolic representations of whole classes of models. The benefit of the correct-by-construction approach, as applied in this paper, is that it ensures correctness for models of arbitrary size because each refined GT rule is guaranteed to preserve the specified correctness properties. In comparison to correct-by-construction and verification-based approaches, the goal of testing is to derive a finite number of representative test cases that cover a part of the space of possible input values of a piece of software that is as large as possible [52].

In the following, we survey formal frameworks for specifying consistency properties in the context of graph-based models. Graphical consistency constraints (for short graph constraints), as introduced in [33] and as used throughout this paper, express the requirement that particular combinations of nodes and edges should be present in or absent from a graph. This formalism has been generalized later to HLR categories in [22] and extended to cope with attributes in [16]. While graph constraints provide the benefit that they can be used to constructively refine GT rules, their expressiveness is relatively limited. For instance, global constraints such as connectivity cannot be expressed using graph constraints. For this reason, we proved the connectivity of weakly consistent topologies in this paper rather than integrating this property constructively during the development of the TC algorithms. In [30, 59], Habel and Radke present HR^* constraints, a new type of graph constraints that allow to express path-related properties, which may in principle also serve as input for the constructive approach. Future work should target the question in how far HR^* constraints are applicable in our application scenario.

In [32], the authors distinguish four situations in which a model transformation considers consistency conditions, including the preservation and enforcement of consistency constraints. The algorithm in this paper preserves the active-link constraint, and it enforces and preserves the inactive-link constraint.

7.2 Model-based Development of Communication Systems

Model-based techniques have shown to be suitable to describe [77] and construct [36] adaptive (communication) systems. Formal analysis of supposed properties of complex topology adaptation algorithms has already revealed special cases in which the implemented algorithms violate crucial topology constraints [86]. In [40], model checking is applied to detect bugs in the TC algorithm LMST, leading to an improved implementation thereof. This paper, in contrast, applies a constructive methodology [33] for GT to develop correct algorithms in the first place. In [45], variants of the TC algorithm kTC [69] are developed using GT, integrating the GT tool eMoflon[6] with a network simulator. While [45] focuses on the rapid prototyping of TC algorithms using programmed GT, this paper aims at devising a generic methodology to develop TC algorithms that fulfill a set of specified constraints by construction.

In the recent years, a number of model-based tools have been proposed for developing TC algorithms. The *Agilla* project[7] [24] provides a middleware platform that allows to use the same implementation for evaluating TC algorithms in a simulation and in a testbed environment. Agilla builds upon fUML, a subset of UML with formally defined semantics [50]. In [9], the authors present an extension of Agilla that is able to collect information w.r.t. the energy consumption of WSN nodes. The *ScatterClipse* project[8] [2] follows a similar goal and supports the TC algorithm developer with visualization and testing facilities. In contrast to this paper, Agilla and Scatterclipse focus on easing the development workflow of TC algorithms. To the best of our knowledge, integrating consistency properties constructively into this design process has not been targeted, yet.

7.3 Tackling Variability: Software Product Lines and Self-Adaptive Systems

One of the major contributions of this paper is to model commonalities and differences of TC algorithms (see S2 in Section 1). We begin with a short discussion of related work that identifies variability as a research challenge in the communication systems domain. In [4], the authors present an evaluation of five TC algorithms in the *WISELIB* algorithm library for WSNs. The authors focus on design decisions related to implementing reusable TC algorithms in WISELIB. In this paper, however, we focus on highlighting and formalizing the commonalities and variabilities of TC algorithms already at specification time.

Managing the commonalities and variability of (software) systems is a central topic in the software product lines (SPL) community. In the following, we provide a short introduction to SPLs. An SPL describes the possible configuration options of a software system, e.g., using feature models [38], whose syntax resembles the diagram shown in Figure 8. While feature models are typically used to model the possible configuration options of a system (often called the *problem space*), metamodeling is a technique for describing abstract representations of concrete systems, the *solution space* [48, 54]. In recent years, the expressiveness of feature models has been enhanced by introducing multiplicities, which further reduces the gap between problem space and solution space [58, 67, 82]. While traditional SPLs typically describe the configuration space of a software system at or before

---

[6] www.emoflon.org

[7] http://mobilab.wustl.edu/projects/agilla/

[8] http://www.mi.fu-berlin.de/inf/groups/ag-tech/projects/Z_Finished_Projects/ ScatterClipse/index.html

its deployment, a dynamic software product line (DSPL) models the possible reconfiguration at runtime [68]. For this purpose, the concept of binding times has been introduced to distinguish between features of a system that are bound, e.g., statically at compile time or dynamically at runtime [11].

Until today, there are only few contributions that connect SPLs with GT. For instance, in [75], the authors propose to model families of GT rules by merging multiple related GT rules into one GT rule whose variables are annotated with presence conditions, which specify for each annotated variable in which variant of the GT rule it is present. Based on this approach, the authors present tool support for automatically deriving variability-based GT rules from a set of traditional GT rules [74] and for editing the derived variability-aware rules [76]. In future work, this approach should be investigated w.r.t. its applicability to the WSN domain.

Several works apply SPL concepts to model adaptive communication systems. In [8], the authors specify the reconfiguration space of a flood warning WSN using SPL. In contrast to this paper, their focus lies on modeling the different communication interfaces (e.g., WiFi or Bluetooth) of a sensor node and the conditions under which the respective interfaces should be activated. In [53], the authors specify a product family of devices that act as environment monitoring and guidance system in a museum at the same time. In contrast to this paper, their focus lies not on TC but on modeling variability w.r.t. the following four dimensions: communication scope (i.e., unicast vs. anycast communication), measured metrics (e.g., humidity or temperature), actuation (e.g., visual or acoustic), and localization technology (e.g., RFID-based localization). In [18], the possible configuration dimensions of wireless sensor-actor network (WSAN) nodes are outlined using feature models. The authors propose a middleware that allows to instantiate the possible configurations in a memory- and energy-efficient way on WSAN nodes. One of the considered components reflects the topology of the WSAN nodes. In contrast to this paper, the authors of [18] focus on surveying the typical complexity of the WSN domain; especially, they treat the two considered TC algorithms (flat tree vs. hierarchical tree) as a black box. In [27] (an extension of [18]), an SPL engineering process is presented that allows to configure resource-efficient middleware systems for WSAN nodes with dedicated tasks based on a user-selected configuration (e.g., to build vehicular area networks [17] or intelligent living spaces [26]). The mapping between configuration space and the larger low-level solution space is performed via model transformation and code generation engine (e.g., for Java2 ME). Finally, [55] presents a variability-aware reference architecture that builds on [27].

In fact, DSPLs are also applied to model self-adaptive systems, i.e., systems that monitor their environment, analyze the monitored data, plan appropriate measures and execute them to adapt to changing contextual environments. In [65], a framework for precalculating possible or probable configurations for resource-constraint devices is proposed. Such techniques are also useful in the WSN domain because the resources of WSN nodes are typically highly limited. In [64], an extension to traditional feature models is proposed that integrates a dedicated submodel for the context of the modeled system. In the context of WSN nodes, context feature models could be used to model the varying environmental conditions to which a sensor node may react by switching to a more appropriate TC algorithm (e.g., Maxpower TC in context with increased node dynamics). While traditional WSNs were typically configured at deployment time using a fixed TC algorithm (if any) and a fixed parameter set, self-adaptive WSNs can be suitable, e.g., for environmental monitoring to detect wildfires or floods. For instance, in [3], the authors present a framework that may be used, e.g., for reconfiguring parameters of WSN nodes in a wildfire detection network. In their case study, the framework is able to predict critical situations (e.g., imminent wildfires) and react appropriately by increasing the sampling rate. Use cases such as flood or wildfire

detection show that switching between TC algorithms is a sensible use case and should be one of our future lines of research.

8 Conclusion

In this paper, we proposed a model-driven methodology for designing families of TC algorithms using GT and graph constraints. The motivation of our research is that the state-of-the-art development of TC algorithms for WSNs exposes two major shortcomings: (i) While it is common to prove formal properties of the designed algorithms and to evaluate them extensively using simulators (and less frequently using hardware testbeds), it is often hard or even impossible to verify that the formal specification and the simulator (or testbed) implementation indeed represent the same algorithm (S1). (ii) Little effort is made to constructively reuse common substructures (e.g., structural constraints or tie breakers) that appear in many TC algorithms (S2).

In [43], we focused on S1 by proposing a systematic approach for developing individual correct-by-construction TC algorithms. In this approach, valid topologies are characterized using graph constraints [33], TC algorithms are specified using programmed GT [23], and an existing constructive approach [33] is applied to enrich GT rules with application conditions derived from the graph constraints. We illustrated the applicability of the proposed approach by re-engineering the TC algorithm kTC [69].

In this paper, we complement our work on S1 and tackle S2 by generalizing the results in [43] as follows: (i) We separate the (graph) constraints into common and algorithm-specific parts. We show the applicability of this step by specifying six existing TC algorithms as well as e-kTC, a novel variant of kTC. (ii) We adjust and extend the steps of the constructive approach proposed in [43] to be also applicable to families of TC algorithms. Finally, we present tool support that allows to immediately evaluate the TC algorithms, specified using the GT tool EMOFLON, in the SIMONSTRATOR network simulation environment.

Thanks to the proposed approach, it is now possible to rapidly specify and evaluate new TC algorithms. First, after specifying an appropriate algorithm-specific predicate, all proves of formal properties carry over to the new TC algorithm. Second, the tool integration between EMOFLON and SIMONSTRATOR mirrors the hierarchical, compositional structure of common and algorithm-specific predicates. This means that only the algorithm-specific predicate of a new TC algorithm needs to be implemented in EMOFLON—all other components may be reused immediately.

Outlook As future work, we aim at extending the proposed systematic approach to support the entire typical development workflow of TC algorithms, consisting of specification, simulation and testbed evaluation (Figure 29). In [42, 43], we have focused on the step-by-step specification and simulation-based evaluation of TC algorithms, which is shortly recapitulated here to lead over to the planned future work. For simplicity, we focus on the development of a single TC algorithm, even though the steps carry over to families of TC algorithms as well. In Step 1, we model relevant entities, relations, and attributes of the considered class of topologies (e.g., node energy level and link weight). In Step 2, we formalize the important properties of the considered TC algorithms in terms of first-order logic predicates. Properties that can be checked based on local knowledge are additionally transformed into graph constraints. For instance, the requirement that the topology is completely classified upon termination of the TC algorithm is translated into the no-unclassified-links constraint C_u. Properties that can only be checked globally are proved manually (e.g., A-connectivity of
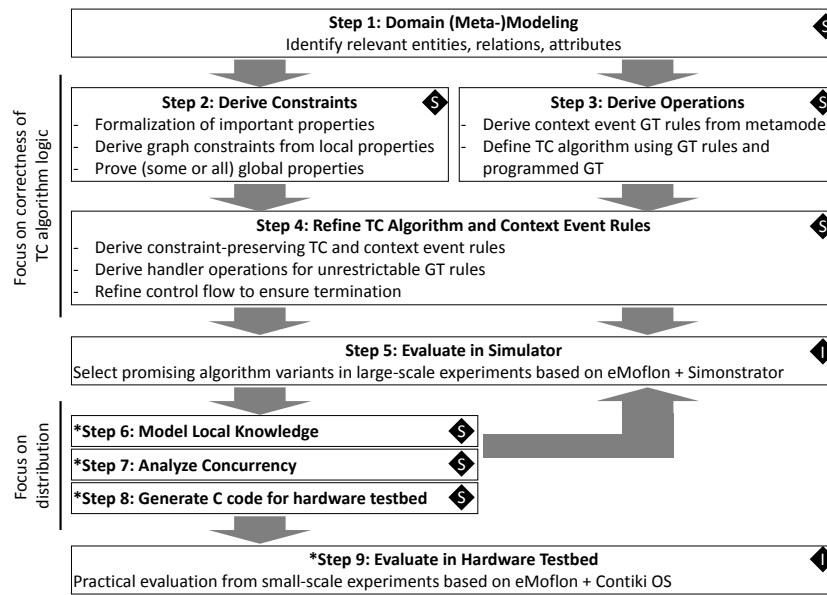
Fig. 29: Overview of systematic approach (***** marks future work; ◆I/◆S: Step in implementation/specification phase)

the output topology). In Step 3, we derive possible TC actions and context events from the metamodel and formalize them using TC and context event GT rules, respectively. For instance, link-weight modifications are represented by the weight modification rule $R_{\text{mod-w}}$. Additionally, we use programmed GT—in this case Story-Driven Modeling [23]—to specify the control flow of the TC algorithm. In Step 4, we combine the graph constraints from the second step and the GT rules from the third step to obtain refined GT rules that preserve all graph constraints. For all unrestrictable GT rules, we transform the generated application conditions into handler operations. At the end of the forth step, we analyze whether the generated application conditions of the TC rules may lead to a non-termination of the TC algorithm and refine its control flow to ensure termination. In Step 5, the TC algorithm is evaluated in a network simulator. For this paper, we decided to use an integration-based approach because the source code of the simulator is available. We could have generated code for the simulator as well, but this is typically a larger effort compared to a tool integration. To support the full development process of TC algorithms, we will address the following research questions in our future work.

– Step 6: How to model information about local knowledge? Currently, several of the patterns require knowledge about the 3-hop neighborhood of a node, which may not be available on real sensor nodes (due to memory limitations).
– Step 7: How can we analyze concurrency issues due to the parallel execution of the TC rules on real hardware? Inside the simulator, TC is executed sequentially, while, on real hardware nodes, the TC algorithm is executed concurrently on each node. We will have to analyze the GT-based specification for race conditions and other concurrency issues and resolve them, e.g., by implementing conflict resolution strategies.

– Step 8: How can we generate efficient code for the hardware testbed? We plan to add a second set of code generation templates to the GT tool EMOFLON to be able to generate embedded C code. Since testbed devices are typically highly resource-constraint, we will focus on optimizing the runtime and the memory footprint (of the generated code and at runtime) in this step.

We plan to use the SIMONSTRATOR [61] platform to evaluate the localized, parallel GT-based specification, which results from answering the first two questions, and we plan to use the Contiki operating system [21] as target platform for generating embedded C code (Step 9).

References

1. Agricola, I.: Elementary Geometry. AMS (2008) (cited on p. 12).
2. Al Saad, M., Fehr, E., Kamenzky, N., Schiller, J.: ScatterClipse: A model-driven tool-chain for developing, testing, and prototyping wireless sensor networks. In: Proc. of the Intl. Symposium on Parallel and Distributed Processing with Applications (ISPA 2008), pp. 871–885 (2008). URL `https://dx.doi.org/10.1109/ISPA.2008.22` (cited on p. 47).
3. Anaya, I.D.P., Simko, V., Bourcier, J., Plouzeau, N., Jézéquel, J.M.: A prediction-driven adaptation approach for self-adaptive sensor networks. In: Proc. of the 9th Intl. Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS 2014), pp. 145–154. ACM, New York, NY, USA (2014). URL `https://dx.doi.org/10.1145/2593929.2593941` (cited on pp. 3, 48).
4. Anguera, J., Blesa, M., Farré, J., López, V., Petit, J.: Topology Control Algorithms in WISELIB. In: Proc. of the ICSE Workshop on Software Engineering for Sensor Network Applications (SESENA 2010), pp. 14–19. ACM, New York, NY, USA (2010). URL `https://dx.acm.org/10.1145/1809111.1809118` (cited on pp. 3, 47).
5. Baldan, P., Corradini, A., König, B.: A framework for the verification of infinite-state graph transformation systems. Information and Computation **206**(7), 869–907 (2008). URL `https://dx.doi.org/10.1016/j.ic.2008.04.002` (cited on p. 46).
6. Baleani, M., Ferrari, A., Mangeruca, L., Sangiovanni-Vincentelli, A., Freund, U., Schlenker, E., Wolff, H.J.: Correct-by-construction transformations across design environments for model-based embedded software development. In: Proc. of Design, Automation and Test in Europe (DATE 2005), vol. 2, pp. 1044–1049 (2005). URL `https://dx.doi.org/10.1109/DATE.2005.105` (cited on p. 46).
7. Basu, A., Bensalem, B., Bozga, M., Combaz, J., Jaber, M., Nguyen, T.H., Sifakis, J.: Rigorous Component-Based System Design Using the BIP Framework. IEEE Software **28**(3), 41–48 (2011). URL `https://dx.doi.org/10.1109/MS.2011.27` (cited on p. 46).
8. Bencomo, N., Sawyer, P., Blair, G., Grace, P.: Dynamically adaptive systems are product lines too: Using model-driven techniques to capture dynamic variability of adaptive systems. In: Proc. of the Intl. Workshop on Dynamic Software Product Lines (DSPL 2008) (2008). URL `https://dx.doi.org/10.1109/SPLC.2008.69` (cited on pp. 3, 48).

9. Berardinelli, L., Di Marco, A., Pace, S., Pomante, L., Tiberti, W.: Energy consumption analysis and design of energy-aware wsn agents in fuml. In: Proc. of the European Conference on Modelling Foundations and Applications (ECMFA 2015), *LNCS*, vol. 9153, pp. 1–17. Springer (2015). URL `https://dx.doi.org/10.1007/978-3-319-21151-0_1` (cited on p. 47).

10. Beydeda, S., Book, M., Gruhn, V.: Model-driven software development, 15th edn. Springer (2005) (cited on p. 3).

11. Bürdek, J., Lity, S., Lochau, M., Berens, M., Goltz, U., Schürr, A.: Staged configuration of dynamic software product lines with complex binding time constraints. In: Proc. of the Intl. Workshop on Variability Modelling of Software-Intensive Systems (VaMoS 2014), pp. 16:1–16:8. ACM, New York, NY, USA (2013). URL `https://dx.doi.org/10.1145/2556624.2556627` (cited on p. 48).

12. Camp, T., Boleng, J., Davies, V.: A survey of mobility models for ad hoc network research. Wireless Communications and Mobile Computing **2**(5), 483–502 (2002). URL `https://dx.doi.org/10.1002/wcm.72` (cited on pp. 36, 37).

13. Chakeres, I., Belding-Royer, E.: AODV routing protocol implementation design. In: Proc. of the Intl. Conf. on Distributed Computing Systems Workshops (ICDCSW 2004), pp. 698–703 (2004). URL `https://dx.doi.org/10.1109/ICDCSW.2004.1284108` (cited on p. 5).

14. Chen, Y., Zhao, Q.: On the lifetime of wireless sensor networks. IEEE Communications Letters **9**(11), 976–978 (2005). URL `https://dx.doi.org/10.1109/LCOMM.2005.11010` (cited on p. 14).

15. Chu, X., Sethu, H.: Cooperative Topology Control with Adaptation for Improved Lifetime in Wireless Ad-Hoc Networks. In: Proc. of the IEEE International Conference on Computer Communications (INFOCOM 2012), pp. 262–270 (2012). URL `https://dx.doi.org/10.1109/INFCOM.2012.6195667` (cited on pp. 3, 14, 15, 46).

16. Deckwerth, F., Varró, G.: Generating Preconditions from Graph Constraints by Higher Order Graph Transformation. In: Proc. of the Intl. Workshop on Graph Transformation and Visual Modeling Techniques (GTVMT 2014), vol. 67, pp. 1–14. ECEASST (2014). URL `https://dx.doi.org/10.14279/tuj.eceasst.67.945` (cited on pp. 26, 46).

17. Delicato, F.C., Fuentes, L., Gámez, N., Pires, P.F.: A Middleware Family for VANETs. In: Proc. of the 8th Intl. Conf. on Ad-Hoc, Mobile and Wireless Networks (ADHOC-NOW 2009), pp. 379–384. Springer (2009). URL `https://dx.doi.org/10.1007/978-3-642-04383-3_31` (cited on p. 48).

18. Delicato, F.C., Fuentes, L., Gámez, N., Pires, P.F.: Variabilities of wireless and actuators sensor network middleware for ambient assisted living. In: Distributed Computing, Artificial Intelligence, Bioinformatics, Soft Computing, and Ambient Assisted Living (IWANN 2009 Workshops), *LNCS*, vol. 5518, pp. 851–858. Springer (2009). URL `https://dx.doi.org/10.1007/978-3-642-02481-8_129` (cited on pp. 3, 48).

19. Dijkstra, E.W.: A discipline of programming, vol. 1. Prentice Hall (1976) (cited on p. 26).

20. Dohler, M., Barthel, D., Maraninchi, F., Mounier, L., Aubert, S., Dugas, C., Buhrig, A., Paugnat, F., Renaudin, M., Duda, A., Heusse, M., Valois, F.: The ARESA Project: Facilitating Research, Development and Commercialization of WSNs. In: Proc .of the IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON 2007), pp. 590–599 (2007). URL `https://dx.doi.org/10.1109/SAHCN.2007.4292871` (cited on pp. 3, 45).

21. Dunkels, A., Gronvall, B., Voigt, T.: Contiki – A Lightweight and Flexible Operating System for Tiny Networked Sensors. In: Proc. of the Intl. Conf. on Local Computer

Networks (LCN 2004), pp. 455–462 (2004). URL `https://dx.doi.org/10.1109/LCN.2004.38` (cited on pp. 2, 51).

22. Ehrig, H., Ehrig, K., Prange, U., Taentzer, G.: Fundamentals of Algebraic Graph Transformation. Springer (2006). URL `https://dx.doi.org/10.1007/3-540-31188-2` (cited on pp. 21, 46).

23. Fischer, T., Niere, J., Torunski, L., Zündorf, A.: Story Diagrams: A New Graph Rewrite Language based on the Unified Modeling Language. In: Proc. of the Intl. Workshop on Theory and Application of Graph Transformation (TAGT 1998), pp. 296–309. Springer (1998). URL `https://dx.doi.org/10.1007/978-3-540-46464-8_21` (cited on pp. 21, 22, 49, 50).

24. Fok, C.L., Roman, G.C., Lu, C.: Agilla: A Mobile Agent Middleware for Self-adaptive Wireless Sensor Networks. ACM Trans. of Autonomous Adaptive Systems **4**(3), 16:1–16:26 (2009). URL `https://dx.doi.org/10.1145/1552297.1552299` (cited on p. 47).

25. Friis, H.T.: A note on a simple transmission formula. Proc. of the Institute of Radio Engineers **34**(5), 254–256 (1946). URL `https://dx.doi.org/10.1109/JRPROC.1946.234568` (cited on pp. 9, 14).

26. Fuentes, L., Gamez, N., Sanchez, P.: Variability in ambient intelligence a family of middleware solution. Ubiquitous Developments in Ambient Computing and Intelligence: Human-Centered Applications pp. 71–83 (2011). URL `https://dx.doi.org/10.4018/978-1-60960-549-0.ch006` (cited on pp. 3, 48).

27. Fuentes, L., Gámez, N.: Configuration process of a software product line for AmI middleware. Journal of Universal Computer Science **16**(12), 1592–1611 (2010). URL `https://dx.doi.org/10.3217/jucs-016-12-1592` (cited on p. 48).

28. Gabriel, K.R., Sokal, R.R.: A new statistical approach to geographic variation analysis. Systematic Biology **18**(3), 259–278 (1969). URL `https://dx.doi.org/10.2307/2412323` (cited on p. 12).

29. Gorp, P.V., Mazanek, S.: SHARE: a web portal for creating and sharing executable research papers. In: Proc. of the Intl. Conf. on Computational Science (ICCS 2011), vol. 4, pp. 589–597 (2011). URL `https://dx.doi.org/10.1016/j.procs.2011.04.062` (cited on pp. 35, 37).

30. Habel, A., Radke, H.: Expressiveness of graph conditions with variables. In: Proc. of the Intl. Colloquium on Graph and Model Transformation (GraMoT 2010), vol. 30. ECEASST (2010). URL `https://dx.doi.org/10.14279/tuj.eceasst.30.404` (cited on p. 46).

31. Hall, A., Chapman, R.: Correctness by construction: developing a commercial secure system. IEEE Software **19**(1), 18–25 (2002). URL `https://dx.doi.org/10.1109/52.976937` (cited on p. 46).

32. Hausmann, J.H., Heckel, R., Sauer, S.: Extended Model Relations with Graphical Consistency Conditions. In: Proc. of the Workshop on Consistency Problems in UML-based Software Development (UML 2002), Blekinge Institute of Technology, Research Report 2002:06, pp. 61–74. Department of Software Engineering and Computer Science, Blekinge Institute of Technology (2002). URL `http://www.db.informatik.uni-bremen.de/umlbib/conf/WRKUML2002CP.html` (cited on p. 46).

33. Heckel, R., Wagner, A.: Ensuring Consistency of Conditional Graph Rewriting – A Constructive Approach. In: Proc. of the Joint COMPUGRAPH/SEMAGRAPH Workshop, *ENTCS*, vol. 2, pp. 118–126. Elsevier (1995). URL `https://dx.doi.org/10.1016/S1571-0661(05)80188-4` (cited on pp. 3, 24, 26, 27, 28, 46, 47, 49).

34. Hermann, F., Gottmann, S., Nachtigall, N., Braatz, B., Morelli, G., Pierre, A., Engel, T.: Model Transformation, Proc. of the Intl. Conference on Model Transformation (ICMT 2013), chap. On an Automated Translation of Satellite Procedures Using Triple Graph Grammars, pp. 50–51. Springer (2013). URL `https://dx.doi.org/10.1007/978-3-642-38883-5_4` (cited on p. 3).

35. Hiranandani, D., Obraczka, K., Garcia-Luna-Aceves, J.J.: MANET protocol simulations considered harmful: the case for benchmarking. IEEE Wireless Communications **20**(4), 82–90 (2013). URL `https://dx.doi.org/10.1109/MWC.2013.6590054` (cited on p. 35).

36. Jacob, R., Richa, A., Scheideler, C., Schmid, S., Täubig, H.: A Distributed Polylogarithmic Time Algorithm for Self-stabilizing Skip Graphs. In: Proc. of the ACM Symposium on Principles of Distributed Computing (PODC 2009), pp. 131–140. ACM (2009). URL `https://dx.doi.org/10.1145/1582716.1582741` (cited on p. 47).

37. Jelasity, M.: Gossip. In: Self-organising Software: From Natural to Artificial Adaptation, pp. 139–162. Springer (2011). URL `https://dx.doi.org/10.1007/978-3-642-17348-6_7` (cited on p. 14).

38. Kang, K.C., Cohen, S.G., Hess, J.A., Novak, W.E., Peterson, S.A.: Feature-Oriented Domain Analysis (FODA) Feasibility Study. Tech. rep., Software Engineering Institute, Carnegie-Mellon University (1990). URL `https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=11231`. CMU/SEI-90-TR-21, ESD-90-TR-222 (cited on p. 47).

39. Karp, B., Kung, H.T.: GPSR: Greedy perimeter stateless routing for wireless networks. In: Proc. of the 6th Annual Intl. Conference on Mobile Computing and Networking (MobiCom 2000), pp. 243–254. ACM (2000). URL `https://dx.doi.org/10.1145/345910.345953` (cited on pp. 3, 12, 18, 21).

40. Katelman, M., Meseguer, J., Hou, J.: Redesign of the LMST Wireless Sensor Protocol through Formal Modeling and Statistical Model Checking. In: Proc. of the Intl. Conf. on Formal Methods for Open Object-Based Distributed Systems (FMOODS 2008), *LNCS*, vol. 5051, pp. 150–169. Springer (2008). URL `https://dx.doi.org/10.1007/978-3-540-68863-1_10` (cited on p. 47).

41. Khemapech, I., Miller, A., Duncan, I.: A survey of transmission power control in wireless sensor networks. In: Proc. of the 8th Annual Postgraduate Symposium on the Convergence of Telecommunications, Networking and Broadcasting (PGNet9s, pp. 15–20 (2007). URL `http://www.cms.livjm.ac.uk/pgnet2007/Proceedings/` (cited on p. 15).

42. Kluge, R., Stein, M., Varró, G., Schürr, A., Mühlhäuser, M., Hollick, M.: A systematic approach to constructing incremental topology control algorithms using graph transformation. Journal of Visual Languages & Computing (JVLC) (2016). URL `http://dx.doi.org/10.1016/j.jvlc.2016.10.003`. To appear (cited on pp. 3, 4, 27, 30, 46, 49).

43. Kluge, R., Varró, G., Schürr, A.: A Methodology for Designing Dynamic Topology Control Algorithms via Graph Transformation. In: Model Transformation, Proc. of the Intl. Conference on Model Transformation (ICMT 2015), *LNCS*, vol. 9152, pp. 199–213. Springer Intl. Publishing (2015). URL `https://dx.doi.org/10.1007/978-3-319-21155-8_15` (cited on pp. 3, 4, 46, 49).

44. Koch, M., Mancini, L.V., Parisi-Presicce, F.: A graph-based formalism for RBAC. ACM Trans. Inf. Syst. Secur. **5**(3), 332–365 (2002). URL `https://dx.doi.org/10.1145/545186.545191` (cited on p. 3).

45. Kulcsár, G., Stein, M., Schweizer, I., Varró, G., Mühlhäuser, M., Schürr, A.: Rapid prototyping of topology control algorithms by graph transformation. In: Proc. of the Intl. Workshop on Graph-Based Tools (GraBaTs 2014), *ECEASST*, vol. 68, pp. 1–15 (2014). URL `https://dx.doi.org/10.14279/tuj.eceasst.68.957` (cited on pp. 36, 47).

46. Kurkowski, S., Camp, T., Colagrosso, M.: MANET simulation studies: The incredibles. SIGMOBILE Mob. Comput. Commun. Rev. **9**(4), 50–61 (2005). URL `https://dx.doi.org/10.1145/1096166.1096174` (cited on p. 35).

47. Leblebici, E., Anjorin, A., Schürr, A.: Developing eMoflon with eMoflon. In: Model Transformation, Proc. of the Intl. Conference on Model Transformation (ICMT 2014), *LNCS*, vol. 8568, pp. 138–145. Springer (2014). URL `https://dx.doi.org/10.1007/978-3-319-08789-4_10` (cited on pp. 3, 35, 36, 37).

48. van der Linden, F., Schmid, K., Rommes, E.: Software Product Lines in Action, 1 edn. Springer (2007). URL `https://dx.doi.org/10.1007/978-3-540-71437-8` (cited on p. 47).

49. Martins, F., Lopes, L., Barros, J.a.: Towards the Safe Programming of Wireless Sensor Networks. In: Proc. of the Second International Workshop on Programming Language Approaches to Concurrency and Communication-cEntric Software (EPTCS 2009), vol. 17, pp. 49–62. Open Publishing Association (2010). URL `https://dx.doi.org/10.4204/EPTCS.17.5` (cited on pp. 3, 45).

50. Mayerhofer, T., Langer, P., Kappel, G.: A Runtime Model for fUML. In: Proc. of the Workshop on Models@Run.Time (MRT 2012), pp. 53–58. ACM, New York, NY, USA (2012). URL `https://dx.doi.org/10.1145/2422518.2422527` (cited on p. 47).

51. Mori, S., Umedu, T., Hiromori, A., Yamaguchi, H., Higashino, T.: Data-centric programming environment for cooperative applications in WSN. In: Proc. of the IFIP/IEEE Intl. Symposium on Integrated Network Management (IM 2013), pp. 856–859 (2013) (cited on pp. 3, 45).

52. Myers, G.J., Sandler, C., Badgett, T.: The art of software testing. John Wiley & Sons (2011) (cited on p. 46).

53. Ortiz, O., García, A.B., Capilla, R., Bosch, J., Hinchey, M.: Runtime variability for dynamic reconfiguration in wireless sensor network product lines. In: Proc. of the 16th Intl. Software Product Line Conference - Volume 2 (SPLC 2012), pp. 143–150. ACM, New York, NY, USA (2012). URL `https://dx.doi.org/10.1145/2364412.2364436` (cited on pp. 3, 48).

54. Pohl, K., Böckle, G., van der Linden, F.: Software Product Line Engineering, 1st edn. Springer (2005). URL `https://dx.doi.org/10.1007/3-540-28901-1` (cited on p. 47).

55. Portocarrero, J.M.T., Delicato, F.C., Pires, P.F., Batista, T.V.: Reference architecture for self-adaptive management in wireless sensor networks. In: Proc. of the Intl. Conf. on Adaptive and Intelligent Systems (ICAIS 2014), pp. 110–120. Springer, Cham (2014). URL `http://dx.doi.org/10.1007/978-3-319-11298-5_12` (cited on pp. 3, 48).

56. Potop-Butucaru, D., Caillaud, B.: Correct-by-Construction Asynchronous Implementation of Modular Synchronous Specifications. Proc. of the Intl. Conf. on Application of Concurrency to System Design (ACSD 2005) pp. 48–57 (2005). URL `https://dx.doi.org/10.1109/ACSD.2005.10` (cited on p. 46).

57. Qadir, J., Hasan, O.: Applying Formal Methods to Networking: Theory, Techniques, and Applications. IEEE Communications Surveys Tutorials **17**(1), 256–291 (2015). URL `https://dx.doi.org/10.1109/COMST.2014.2345792` (cited on pp. 3, 45).

58. Quinton, C., Romero, D., Duchien, L.: Cardinality-based Feature Models with Constraints: A Pragmatic Approach. In: Proc. of the Intl. Software Product Line Confer-

ence (SPLC 2013), pp. 162–166. ACM, New York, NY, USA (2013). URL `http://dx.doi.org/10.1145/2491627.2491638` (cited on p. 47).

59. Radke, H.: Weakest Liberal Preconditions relative to HR* Graph Conditions. In: Pre-proceedings of the Intl. Workshop on Graph Computation Models (GCM 2010), pp. 165–178 (2010). URL `http://formale-sprachen.informatik.uni-oldenburg.de/~skript/fs-pub/Radk10b.pdf` (cited on p. 46).

60. Rensink, A., Schmidt, A., Varró, D.: Model checking graph transformations: A comparison of two approaches. In: Graph Transformations, Proc. of the Intl. Conference on Graph Transformation (ICGT 2004), *LNCS*, vol. 3256, pp. 226–241. Springer (2004). URL `https://dx.doi.org/10.1007/978-3-540-30203-2_17` (cited on p. 46).

61. Richerzhagen, B., Stingl, D., Rückert, J., Steinmetz, R.: Simonstrator: Simulation and Prototyping Platform for Distributed Mobile Applications. In: Proc. of the Intl. Conf. on Simulation Tools and Techniques (SIMUTools '15), pp. 99–108. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering) (2015). URL `https://dx.doi.org/10.4108/eai.24-8-2015.2261064` (cited on pp. 2, 3, 35, 37, 51).

62. Rodoplu, V., Meng, T.H.: Minimum energy mobile wireless networks. IEEE Journal on Selected Areas in Communications **17**(8), 1333–1344 (1999). URL `https://dx.doi.org/10.1109/49.779917` (cited on pp. 12, 21).

63. Rozenberg, G. (ed.): Handbook of Graph Grammars and Computing by Graph Transformation, vol. 1: Foundations. World Scientific (1997). URL `https://dx.doi.org/10.1142/3303` (cited on pp. 3, 21).

64. Saller, K., Lochau, M., Reimund, I.: Context-aware DSPLs: Model-based Runtime Adaptation for Resource-constrained Systems. In: Proc. of the Intl. fSoftware Product Line Conference Co-located Workshops (SPLC 2013 Workshops), pp. 106–113. ACM, New York, NY, USA (2013). URL `https://dx.doi.org/10.1145/2499777.2500716` (cited on pp. 3, 48).

65. Saller, K., Oster, S., Schürr, A., Schroeter, J., Lochau, M.: Reducing feature models to improve runtime adaptivity on resource limited devices. In: Proc. of the Intl. Software Product Line Conference - Volume 2 (SPLC 2012), pp. 135–142. ACM, New York, NY, USA (2012). URL `https://dx.doi.org/10.1145/2364412.2364435` (cited on pp. 3, 48).

66. Santi, P.: Topology control in wireless ad hoc and sensor networks. ACM computing surveys (CSUR) **37**(2), 164–194 (2005). URL `https://dx.doi.org/10.1145/1089733.1089736` (cited on pp. 2, 5, 6).

67. Schnabel, T., Weckesser, M., Kluge, R., Lochau, M., Schürr, A.: CardyGAn: Tool support for cardinality-based feature models. In: Proc. of the Intl. Workshop on Variability Modelling of Software-intensive Systems (VaMoS 2016), pp. 33–40 (2016). URL `https://dx.doi.org/10.1145/2866614.2866619` (cited on p. 47).

68. Schroeter, J., Mucha, P., Muth, M., Jugel, K., Lochau, M.: Dynamic configuration management of cloud-based applications. In: Proc. of the Intl. Software Product Line Conference (SPLC 2012), pp. 171–178. ACM, New York, NY, USA (2012). URL `https://dx.doi.org/10.1145/2364412.2364441` (cited on p. 48).

69. Schweizer, I., Wagner, M., Bradler, D., Mühlhäuser, M., Strufe, T.: kTC - Robust and Adaptive Wireless Ad-Hoc Topology Control. In: Proc. of the Intl. Conf. on Computer Communications and Networks (ICCCN 2012), pp. 1–9 (2012). URL `https://dx.doi.org/10.1109/ICCCN.2012.6289318` (cited on pp. 3, 9, 18, 21, 36, 47, 49).

70. Stein, M., Kulcsár, G., Schweizer, I., Varró, G., Schürr, A., Mühlhäuser, M.: Topology Control with Application Constraints. In: Proc. of the Intl. Conf. on Local Computer

Networks (LCN 2015), pp. 438–441 (2015). URL `https://dx.doi.org/10.1109/LCN.2015.7366313` (cited on pp. 3, 36).

71. Stein, M., Petry, T., Schweizer, I., Bachmann, M., Mühlhäuser, M.: Topology control in wireless sensor networks: What blocks the breakthrough? In: Proc. of the Intl. Conf. on Local Computer Networks (LCN 2016), pp. 1–9 (2016). To appear (cited on pp. 3, 12, 13, 18, 21, 36).

72. Steinberg, D., Budinsky, F., Merks, E., Paternostro, M.: EMF: Eclipse Modeling Framework. Pearson Education (2008) (cited on p. 36).

73. Stingl, D., Gross, C., Rückert, J., Nobach, L., Kovacevic, A., Steinmetz, R.: Peerfact-Sim.KOM: A simulation framework for peer-to-peer systems. In: Proc. of the Intl. Conf. on High Performance Computing and Simulation (HPCS 2011), pp. 577–584. IEEE (2011). URL `https://dx.doi.org/10.1109/HPCSim.2011.5999877` (cited on pp. 35, 36, 37).

74. Strüber, D., Rubin, J., Arendt, T., Chechik, M., Taentzer, G., Plöger, J.: RuleMerger: Automatic construction of variability-based model transformation rules. In: Proc. of Fundamental Approaches to Software Engineering (FASE 2016), pp. 122–140. Springer (2016). URL `https://dx.doi.org/10.1007/978-3-662-49665-7_8` (cited on p. 48).

75. Strüber, D., Rubin, J., Chechik, M., Taentzer, G.: A variability-based approach to reusable and efficient model transformations. In: Proc. of Fundamental Approaches to Software Engineering (FASE 2015), pp. 283–298. Springer (2015). URL `https://dx.doi.org/10.1007/978-3-662-46675-9_19` (cited on p. 48).

76. Strüber, D., Schulz, S.: A tool environment for managing families of model transformation rules. In: Graph Transformations, Proc. of the Intl. Conf. on Graph Transformation (ICGT 2016), pp. 89–101. Springer (2016). URL `https://dx.doi.org/10.1007/978-3-319-40530-8_6` (cited on p. 48).

77. Taentzer, G., Goedicke, M., Meyer, T.: Dynamic change management by distributed graph transformation: Towards configurable distributed systems. In: Proc. of the Intl. Workshop on Theory and Application of Graph Transformations (TAGT 2000), pp. 179–193. Springer (2000). URL `https://dx.doi.org/10.1007/978-3-540-46464-8_13` (cited on p. 47).

78. Valente, B., Martins, F.: A Middleware Framework for the Internet of Things. In: Proc. of the Intl. Conf. on Advances in Future Internet, pp. 139–144. ThinkMind Digital Library (2011) (cited on p. 45).

79. Völter, M., Stahl, T., Bettin, J., Haase, A., Helsen, S.: Model-Driven Software Development: Technology, Engineering, Management. John Wiley & Sons (2013) (cited on p. 3).

80. Wang, Y.: Topology control for wireless sensor networks. In: Wireless Sensor Networks and Applications, Signals and Communication Technology, pp. 113–147. Springer (2008). URL `https://dx.doi.org/10.1007/978-0-387-49592-7_5` (cited on pp. 2, 3, 12, 18, 21).

81. Wattenhofer, R., Zollinger, A.: XTC: a practical topology control algorithm for ad-hoc networks. In: Proc. of the Intl. Parallel and Distributed Processing Symposium (IPDPS 2004), pp. 216–223. IEEE (2004). URL `https://dx.doi.org/10.1109/IPDPS.2004.1303248` (cited on pp. 3, 12, 18, 21).

82. Weckesser, M., Lochau, M., Schnabel, T., Richerzhagen, B., Schürr, A.: Mind the Gap! Automated Anomaly Detection for Potentially Unbounded Cardinality-Based Feature Models. In: Proc. of Fundamental Approaches to Software Engineering (FASE 2016), pp. 158–175. Springer (2016). URL `http://dx.doi.org/10.1007/`

`978-3-662-49665-7_10` (cited on p. 47).

83. Winter, T.: RPL: IPv6 routing protocol for low-power and lossy networks. IETF RFC 6550 (2012). URL `https://tools.ietf.org/html/rfc6550` (cited on p. 5).

84. Yao, A.C.C.: On constructing minimum spanning trees in k-dimensional spaces and related problems. SIAM Journal on Computing **11**(4), 721–736 (1982). URL `https://dx.doi.org/10.1137/0211059` (cited on pp. 3, 13, 18, 21).

85. Zave, P.: Understanding SIP through Model-Checking. In: Principles, Systems and Applications of IP Telecommunications. Services and Security for Next Generation Networks, *LNCS*, vol. 5310, pp. 256–279. Springer (2008). URL `https://dx.doi.org/10.1007/978-3-540-89054-6_13` (cited on p. 3).

86. Zave, P.: Using Lightweight Modeling to Understand Chord. SIGCOMM Comput. Commun. Rev. **42**(2), 49–57 (2012). URL `https://dx.doi.org/10.1145/2185376.2185383` (cited on pp. 3, 47).

87. Zimmermann, H.: OSI Reference Model–The ISO Model of Architecture for Open Systems Interconnection. IEEE Transactions on Communications **28**(4), 425–432 (1980). URL `https://dx.doi.org/10.1109/TCOM.1980.1094702` (cited on p. 36).