Niklas Reusch 🖂

Technical University of Denmark Kongens Lyngby, Denmark

Silviu S. Craciunas 🖂

TTTech Computertechnik AG, Vienna, Austria

Paul Pop ⊠

Technical University of Denmark Kongens Lyngby, Denmark

— Abstract

Time-Sensitive Networking (TSN) extends IEEE 802.1 Ethernet for safety-critical and real-time applications in several areas, e.g., automotive, aerospace or industrial automation. However, many of these systems also have stringent security requirements, and security attacks may impair safety. Given a TSN-based distributed architecture, a set of applications with tasks and messages, as well as a set of security and redundancy requirements, we are interested to synthesize a system configuration such that the real-time, safety and security requirements are upheld. We use the Timed Efficient Stream Loss-Tolerant Authentication (TESLA) low-resource multicast authentication protocol to guarantee the security requirements, and redundant disjunct message routes to tolerate link failures. We consider that tasks are dispatched using a static cyclic schedule table and that the messages use the time-sensitive traffic class in TSN, which relies on schedule tables (called Gate Control Lists, GCLs) in the network switches. A configuration consists of the schedule tables for tasks as well as the disjoint routes and GCLs for messages. We propose a Constraint Programming-based formulation which can be used to find an optimal solution with respect to our cost function. Additionally, we propose a Simulated Annealing based metaheuristic, which can find good solution for large test cases. We evaluate both approaches on several test cases.

2012 ACM Subject Classification Computer systems organization \rightarrow Dependable and fault-tolerant systems and networks

Keywords and phrases TSN, real-time, scheduling

1 Introduction

Many modern safety-critical real-time systems are implemented on distributed architectures. They integrate functions with different security and safety requirements over the same deterministic communication network. For example, the network in a modern vehicle has to integrate high-bandwidth video and LIDAR data for Advanced Driver Assistance Systems (ADAS) functions with the highly critical but low bandwidth traffic of e.g. the powertrain functions, but also with the best-effort messages of the low-criticality diagnostic services. See Figure 1 for an example network architecture of a modern vehicle.

Time-Sensitive Networking (TSN) [19], which is becoming the standard for communication in several application areas, e.g. automotive to industrial control, is comprised of a set of amendments and additions to the IEEE 802.1 standard, equipping Ethernet with the capabilities to handle real-time mixed-criticality traffic with high bandwidth. A TSN network consists of several end-systems, that run mixed-criticality applications, interconnected via network switches and physical links. Available traffic types are Time-Triggered (TT) traffic for real-time applications, Audio-Video Bridging (AVB) for communication that requires less stringent bounded latency guarantees, and Best-Effort (BE) traffic for non-critical traffic.

We assume that safety-critical applications are scheduled using static cyclic scheduling and use the TT traffic type with a given *Redundancy Level* (RL) for communication. We consider that the task-level redundancy is addressed using solutions such as replication [21], and we instead focus



Figure 1 Example automotive TSN-based CPS with redundant routing

on the safety and security of the communication in TSN. The real-time safety requirements of critical traffic in TSN networks are enforced through offline-computed schedule tables, called Gate Control Lists (GCLs), that specify the sending and forwarding times of all critical frames in the network. Scheduling time-sensitive traffic in TSN is non-trivial (and fundamentally different from e.g. TTEthernet), because TSN does not schedule communication at the level of individual frames as is the case in TTEthernet. Instead, the static schedule tables (GCLs) governs the behavior of entire traffic classes (queues) which may lead to non-deterministic frame transmissions [9].

Since link and connector failures in TSN could result in fatal consequences, the network topology uses redundancy, e.g., derived with methods such as [12]. In TSN, IEEE 802.1CB Frame Replication and Elimination for Reliability (FRER) enables the transmission of duplicate frames over different (disjoint) routes, implementing merging of frames and discarding of duplicates.

Nowadays modern Cyber-Physical Systems (CPSs) are becoming increasingly more interconnected with the outside world opening new attack vectors [29,44] that may also compromise safety. Therefore, the security aspects should be equally important to the safety aspects. Timed Efficient Stream Loss-Tolerant Authentication (TESLA) [30] has been investigated as a low resource authentication protocol for several networks, such as FlexRay and TTEthernet [47] networks. Adding security mechanisms such as TESLA after the scheduling stage is oftentimes not possible without breaking real-time constraints, e.g. on end-to-end latency, and degrading the performance of the system [47]. Thus we consider TESLA and the overhead and constraints it imposes part of our configuration synthesis problem formulation.

1.1 Related Work

Scheduling for TSN networks is a well-researched problem. It has been solved for a variety of different traffic type combinations (TT, AVB, BE) and device capabilities using methods such as Integer Linear Programming (ILP), Satisfiability Modulo Theories (SMT) or various metaheuristics such as tabu search [9, 10, 13, 39].

Routing has also been extensively researched [14, 46]. The authors in [40] presented an ILP solution to solve the routing problem for safety-critical AFDX networks. In [45] the authors used a tabu search metaheuristic to solve the combined routing and scheduling problem for TT traffic in TTEthernet. In [33] the authors provide a simple set of constraints to solve a general multicast routing problem using constraint programming, which [12] builds on that to solve a combined topology and route synthesis problem. In [27] the authors use a load-balancing heuristic to distribute the bandwidth

usage over the network and achieve smaller latency for critical traffic.

Multiple authors have also looked at the combined routing and scheduling problem. The authors in [24] and [26] showed that they are able to significantly reduce the latency by solving the combined problem with an ILP formulation. In [28] the authors presented a heuristic for a more complex application model that allows multicast streams. They were able to solve problems that were infeasible to solve using ILP or separate routing and scheduling.

Recently authors have started to present security- and redundancy-aware problem formulations. The authors in [47] provided a security-aware scheduling formulation for TTEthernet using TESLA for authentication. In [25] the authors solve the combined routing and scheduling problem and considered authentication using block ciphers. The authors in [16] and [6], on the other hand, present a routing and scheduling formulation that is redundancy-aware but has no security considerations.

To the best of our knowledge our work is the first one to provide a formulation that is both security and redundancy-aware.

1.2 Contributions

In this paper, we address TSN-based distributed safety-critical systems and solve the problem of configuration synthesis such that both safety and security aspects are considered. Determining an optimized configuration means deciding on the schedule tables for tasks as well as the disjoint routes and GCLs for messages. Our contributions are the following:

- 1. We apply TESLA to TSN networks considering both the timing constraints imposed by TSN and the security constraints imposed by TESLA.
- 2. We formulate an optimization problem to determine: (i) the redundant routing of all messages; (ii) the schedule of all messages, encapsulated into Ethernet frames, represented by the GCLs in the network devices, and (iii) the schedule of all related tasks on end-systems.
- 3. We extend our Constraint Programming (CP) formulation from [37] and propose a new Simulated Annealing (SA)-based metaheuristic to tackle large scale networks that cannot be solved with CP
- 4. We evaluate the impact of adding the security from TESLA on the schedulability of applications and we evaluate the solution quality and scalability of the Constraint Programming (CP) and Simluated Annealing (SA) optimization approaches

We introduce the fundamental concepts of TSN in section 2 and of TESLA in section 3. In section 4 we present the model of our system, consisting of the architecture of the network, applications running on this architecture. Additionally we present a threat model and how it is addressed by TESLA with a security model. In section 5 we formulate the problem we are solving using the established models and present an example. In section 6 and section 7 we present the two different optimization approaches, CP and SA. Then, we evaluate these approaches using several test cases in section 8. section 9 concludes the paper.

2 Time-Sensitive Networking

Time-Sensitive Networking [19] has arisen out of the need to have more stringent real-time communication capabilities within standard Ethernet networks. Other technologies that offer real-time guarantees for distributed systems are TTEthernet (SAE AS6802 [20, 43]), PROFINET, and EtherCAT [35]. TSN comprises a set of (sub-)standards and amendments for the IEEE 802.1Q standard, introducing several new mechanisms for Ethernet bridges, extensions to the IEEE 802.3 media access control (MAC) layer, as well as other standards and protocols (e.g., 802.1ASrev).

The fundamental mechanisms that enable deterministic temporal behavior over Ethernet are, on the one hand, the clock synchronization protocol defined in IEEE 802.1ASrev [18], which provides a common clock reference with bounded deviation for all nodes in the network, and on the other hand,



Figure 2 Simplified TSN switch representation

the timed-gate functionality (IEEE 802.1Qbv [3]) enhancing the transmission selection on egress ports. The timed-gate functionality (IEEE 802.1Qbv [3]) enables the predictable transmission of communication streams according to the predefined times encoded in so-called Gate-Control Lists (GCL). A stream in TSN definition is a communication carrying a certain payload size from a talker (sender) to one or more listeners (receivers), which may or may not have timing requirements. In the case of critical streams, the communication has a defined period and a maximum allowed end-to-end latency.

Other amendments within TSN (c.f. [19]) provide additional mechanisms that can be used either in conjunction with 802.1Qbv or stand-alone. IEEE 802.1CB [5] enables stream identification, based on e.g., the destination MAC and VLAN-tag fields in the frame, as well as frame replication and elimination for redundant transmission. IEEE 802.1Qbu [2] enables preemption modes for mixedcriticality traffic, allowing express frames to preempt lower-priority traffic. IEEE 802.1Qci [4] defines frame metering, filtering, and time-based policing mechanisms on a per-stream basis using the stream identification function defined in 802.1CB.

We detail the Time-Aware Shaper (TAS) mechanism defined in IEEE 802.1Qbv [3] via the simplified representation of a TSN switch in Figure 2. The figure presents a scenario in which communication received on one of two available ingress ports (A and B) will be routed to an egress port C. The switching fabric will determine, based on internal routing tables and stream properties, to which egress port a frame belonging to the respective stream will be routed (in our logical representation, there is only one egress port). Each port will have a priority filter that determines which of the available 8 traffic classes (priorities) of that port the frame will be enqueued in. This selection will be made based on either the priority code point (PCP) contained in the VLAN-tag of 802.1Q frames or the stream gate instance table of 802.1Qci, which can be used to circumvent traffic class assignment of the PCP code. As opposed to regular 802.1Q bridges, where the transmission selection sends enqueued frames according to their respective priority, in 802.1Qbv bridges, there is a Time-Aware Shaper (TAS), also called timed-gate, associated with each traffic class queue and positioned before the transmission selection algorithm. A timed-gate can be either in an open (o) or closed (C) state. When the gate is open, traffic from the respected queue is allowed to be transmitted, while a closed gate will not allow the respective queue to be selected for transmission, even if the queue is not empty. The state of the queues is encoded in a local schedule called Gate-Control List (GCL). Each entry defines a time value and a state (o or C) for each of the 8 queues. Hence whenever the local clock reaches the specified time, the timed-gates will be changed to the respective open

N. Reusch et. al.

or closed state. If multiple non-empty queues are open at the same time, the transmission selection selects the queue with the highest priority for transmission.

The Time-Aware Shaper functionality of 802.1Qbv, together with the synchronization protocol defined in 802.1ASrev, enables a global communication schedule that orchestrates the transmission of frames across the network such that real-time constraints (usually end-to-end latencies) are fulfilled. The global schedule synthesis has been studied in [9, 10, 34, 39] focusing on enforcing deterministic transmission, temporal isolation, and compositional system design for critical streams with end-to-end latency requirements.

Craciunas et al. [9] define correctness conditions for generating GCL schedules, resulting in a strictly deterministic transmission of frames with 0 jitter. Apart from the technological constraints, e.g., only one frame transmitted on a link at a time, the deterministic behavior over TSN is enforced in [9] through isolation constraints. Since the TAS determines the temporal behavior of entire traffic classes (as opposed to individual frames like in TTEthernet [42]), the queue state always has to be deterministic. Hence, in [9], critical streams are isolated from each other either in the time or space domain by either allowing only one stream to be present in a queue at a time or by isolating streams that are received at the same time in different queues. This condition is called *frame/stream isolation* in [9]. In [39], critical streams are allowed to overlap to some degree (determined by a given jitter requirement) in the same queue in the time domain, thus relaxing the strict isolation.

Both approaches enforce that gate states of different scheduled queues are mutually exclusive, i.e., only one gate is open at any time, thus preventing the transmission selection from sending frames based on their assigned traffic class's priority. By circumventing the priority mechanism through the TAS, it is ensured that no additional delay is produced through streams of higher priorities, enforcing thus a highly deterministic temporal behavior.

3 Timed Efficient Stream Loss-Tolerant Authentication

TESLA provides a resource efficient way to do asymmetric authentication in a multicast setting [30]. It is described in detail in [30] and [32].

We are considering systems where one end-system wants to send a multicast-signal to multiple receiver end-systems, e.g., periodic sensor data. A message authentication code (MAC), which is appended to each signal, can guarantee authenticity, i.e., that the sender is who he claims to be, and integrity, i.e., that the message has not been altered. The MAC is generated and authenticated by a secret key that all end-systems share (i.e., symmetric authentication). The downside of this approach is that if any of the receiving end-systems is compromised, the attacker would be able to masquerade as the sender by knowing the secret key. In a multicast setting, an asymmetric approach, in which the receivers do not have to trust each other, is preferable.

The traditional asymmetric authentication approach is to use asymmetric cryptography with digital signatures (i.e., private and public keys); however, as stated in [31], the method is computationally intensive and not well suited for systems with limited resources and strict timing constraints.

TESLA, however, uses an approach where the source of asymmetry is a time-delayed key disclosure [30]. While this can be implemented with much less overhead, it requires time synchronization between the network nodes. For TSN, the time synchronization is given through the 802.1ASrev protocol.

Figure 3 visualizes the TESLA protocol. As described in [31], when using TESLA, time is divided into fixed intervals of length P_{int} . At startup a one-way chain of self authenticating keys K_i is generated using a hash function H, where $K_i = H(K_{i+1})$. Each key is assigned to one interval. The protocol is bootstrapped by creating this chain and securely distributing K_0 to all receivers [47].

Normally in TESLA, as described in [47], when a sender sends a message m in the *i*-th interval,



Figure 3 TESLA key chain (Adapted from [31])

it appends to that message: *i*, a keyed-MAC using the key of that interval K_i , and a previously used key K_{i-d} . Thus, a key remains secret for *d* intervals. When a receiver receives a message *m* in the interval *i* it can not yet authenticate it. It must wait until a message arrives in the interval i+d. This message discloses K_i , which can be used to decrypt the MAC of *m* and thus authenticate it. To ensure that K_i itself is valid, we can use any previously validated key. For example, we can check that $H(K_i) = K_{i-1}, H(H(K_i)) = K_{i-2}$ etc. This makes TESLA also robust to packet loss since any lost keys can be reconstructed from a later key, and any key can always be checked against K_0 .

Due to the deterministic nature of our schedule, we can make some modifications to the basic TESLA protocol without sacrificing security. The first modification is adopted from [47]. Since bandwidth is scarce, we do not release the key K_{i-d} with every message/stream. Instead, it will be released once in its own stream with an appropriate redundancy level. The second modification concerns the TESLA parameter *d*. This parameter is useful in a non-deterministic setting. Since the arrival time of a stream is uncertain, a high value for *d* makes it more likely that a stream can be authenticated, at the cost of an increased latency. [32] However, in our case, we know the exact time a stream will be sent and arrive. Thus, we assume that a stream's keyed-MAC will be generated using the key from the interval it *arrives at the last receiver*. Furthermore, we will release the key K_i in the interval, i + 1 minimizing the latency before a stream can be authenticated.

4 System Models

This section presents the architecture and application models, as well as the threat, security and fault models. Our application model is similar to the one used in related work [47], but we have extended it to consider TSN networks and the optimization of redundant routing in conjunction with scheduling.

4.1 Architecture Model

We model our TSN network as a directed graph consisting of a set of nodes \mathcal{N} and a set of edges \mathcal{L} . The nodes of the graph are either end-systems (ESs) or switches (SWs): $\mathcal{N} = \mathcal{ES} \cup \mathcal{SW}$. The edges \mathcal{L} of the graph represent the network links.

We assume that all of the nodes in the network are TSN-capable, specifically that they support the standards 802.1ASrev [18] and 802.1Qbv [3]. Thus we assume the whole network, including the end-systems, to be time-synchronized with a known bounded precision δ . All nodes use the time-aware shaper mechanism from 802.1Qbv to control the traffic flow.

Each end-system $e_i \in \mathcal{ES}$ features a real-time operating system with a periodic table-driven task scheduler. Hash computations, which will be necessary for TESLA operations on that end-system, take $e_i.H \ \mu s$.

A network link between nodes $n_a \in \mathcal{N}$ and $n_b \in \mathcal{N}$ is defined as $l_{a,b} \in \mathcal{L}$. Since in Ethernetcompliant networks all links are bi-directional and full-duplex, we have that for each $l_{a,b} \in \mathcal{L}$ there is also $l_{b,a} \in \mathcal{L}$. A link $l_{a,b} \in \mathcal{L}$ is defined by a link speed $l_{a,b}.s$.

Figure 4a shows a small example architecture with four end-systems, two switches, and full-duplex links.

Table 1 Notations

Description	Notation	Unit
Header overhead	ОН	Byte
Maximum transmission unit	MTU	Byte
TESLA key size	KS	Byte
TESLA MAC size	MAC	Byte
Hyperperiod	Н	μs
TSN Network Graph	$(\mathcal{N},\mathcal{L})$	
- Nodes	$\mathcal{N} = \mathcal{ES} \cup \mathcal{SW}$	
- End-system	$e_i \in \mathcal{ES}$	
- Hash computation time	$e_i.H$	μs
- Switch	$sw_j \in \mathcal{SW}$	
- Links	$\mathcal{L} \subseteq \mathcal{N} \times \mathcal{N}$	
- Network link	$l_{a,b}$	
- Link speed	$l_{a,b}.s$	μs
Application	$\lambda_l \in \Lambda$	
- Tuple	$(\Gamma_l, \mathcal{E}_l)$	
- Period	$\lambda_l.T$	μs
- Communication Depth	$\lambda_l.C$	
- Tasks	$t_m \in \mathcal{T}$	
- Execution end-system	$t_m.e$	
- Worst-case execution time	$t_m.w$	μs
- Period	$t_m.T$	μs
- Streams	$s_n \in \mathcal{S}$	
- Source task	$s_n.t_s$	
- Destination tasks	$s_n.T_d$	
- Size	$s_n.b$	Byte
- Period	$s_n.T$	μs
- Redundancy Level	$s_n.rl$	
- Security Level	$s_n.sl$	
- MAC generation task	$t_{s_n}^g$	
- MAC verification task	$t_{S_n}^{mv}$	
Security Application	$\lambda_l^s \in \Lambda^{sec}$	
- Key release task	t_m^r	
- Key verification task	t_m^{ν}	
- Key source end-system	$t_m^v.src$	
- Key stream	$s_n^k \in \mathcal{S}_k$	



Figure 4 Example architecture and application models

4.2 Application Model

An application $\lambda_l \in \Lambda$ is modeled as a directed, acyclic graph consisting of a set of nodes representing tasks Γ_l and a set of edges \mathcal{E}_l represents a data dependency between tasks.

A task is executed on a certain end-system $t_m.e$. The worst-case execution time (WCET) of a task is defined by $t_m.w \ \mu s$. A task needs all its incoming streams (incoming edges in the application graph) to arrive before it can be executed. It produces outgoing streams at the end of its execution time. Communication dependencies between tasks that run on the same end-system are usually done via, e.g., shared memory pools or message queues, where the overhead of reading/writing data is negligible and included in the WCET definition of the respective tasks. Dependencies between tasks on separate end-systems constitute communication requirements and are modeled by streams. A stream in the TSN context is a communication requirement between a sender and one (unicast) or multiple (multicast) receivers. An example application can be seen in Figure 4b. An application is periodic with a period $\lambda_l.T$, which is inherited by all its tasks and streams.

A stream s_n originates at a source task $s_n.t_s$ and travels to set of destination tasks $s_n.T_d$ (since we consider multicast streams). The stream size $s_n.b$ is assumed to be smaller than the maximum transmission unit (MTU) of the network. Each stream has a redundancy level $s_n.rl$, which determines the amount of required disjunct redundant routes for the stream to take. For each of these routes we model a sub-stream: $s_n^i \in S_{s_n}, 0 \le i < s_n.rl$ Hereby S_{s_n} is a set containing all sub-streams of s_n . This notation is useful to differentiate the different routes a stream takes through the network, and to make sure those routes do not overlap. A stream also has a binary security level $s_n.sl$ which determines if it is authenticated using TESLA (sl = 1) or not (sl = 0).

We define the hyperperiod *H* as the least-common multiple of all application periods: $H = lcm(\{\lambda_l, T | \lambda_l \in \Lambda\})$ We define the set \mathcal{T} to contain all tasks and the set \mathcal{S} to contain all streams (including redundant copies).

4.3 Fault Model

Reliability models discussed in [12] (e.g., Siemens SN 29500) indicate that the most common type of permanent hardware failures is due to link failures (especially physical connectors) and that ESs

and SWs are less likely to fail. These models are complementary to Mean Time to Failure (MTTF) targets established for various safety integrity levels in certification standards such as ISO 26262 for automotive [12]. As mentioned, we assume we know the required redundancy level to protect against permanent link failures. Our disjoint routing can guarantee the transmission of a stream of RL n despite any n - 1 link failures. For example, for the routing of s_2 with RL 2 in Figure 4a, any 1-link failure would still result in a successful transmission.

4.4 Threat Model

We use a similar threat model as [47] and assume that an attacker is capable of gaining access to some end-systems of our system, e.g., through an external gateway or physical access.

We consider that the attackers have the following abilities:

- They know about the network schedule and the content of the streams on the network;
- They can replay streams sent by other ES;
- They can attempt to masquerade as other ES by faking the source address of streams they send;
- They have access to all keys released and received by the ES they control;

4.5 Security Model

We use TESLA to address the threats identified in the previous section, which means that additional security-related models are required. These additional applications, tasks and streams can be automatically generated from a given architecture and application model.

First off, we need to generate, send, and verify a key in each interval for each set of communicating end-systems. We generate a key authentication application $\lambda_s \in \Lambda^{sec}$ for each sender end-system, which is modeled similarly to a normal application as a directed acyclic graph. The period $\lambda_s.T$ is equal to P_{int} (see section 3) and again inherited by tasks and streams. Each of these application consists of one key release task t_m^r scheduled on the sending end-system e_i . Additionally, it consists of key verification tasks t_j^v on each end-system e_j that receives a stream from e_i . The release task sends a multicast key-stream s_n^k to each of those verification tasks. The redundancy level of a key-stream $s_n^k.rl$ is set to the maximum redundancy level of all streams emitted by e_i . The size of a key stream $s_n^k.b$ is equal to the key size KS specified by the TESLA implementation. The security model for our example from Figure 4 can be seen in Figure 5.

For a key verification task t_m^v .src is the end-system e_i whose key this task is verifying. Its execution time is equal to the length of one hash execution on its execution end-system: t_m^v . $w = (t_m^v.e).H$. A key release task's execution time is very short, since the key it releases has already been generated during bootstrapping. We model it to be last half the time of a hash execution: t_m^r . $w = (\frac{t_m^r.e).H}{2}$

Secondly, we need to append MACs to all non-key-streams with $s_n.sl = 1$. Thus, their length increases by the MAC length *MAC* specified by the TESLA implementation. For each stream s_n , a MAC generation task $t_{s_n}^g$ is added to the sender and a MAC validation task $t_{s_n}^{mv}$ to each receiver. Those tasks take the time of one MAC computation on the processing element to execute.

We define the set \mathcal{T}_{kr}^n to contain all key release tasks and \mathcal{T}_{kv}^n to contain all key verification tasks for a given node *n*. Furthermore let S_k contain all key streams.

Figure 4a shows key release and verification tasks in orange and MAC generation and validation tasks in red.

Figure 5 shows the security applications for our example.



Figure 5 Example security model for the applications in Figure 4

5 Problem Formulation

Given a set of applications running on TSN-capable end-systems that are interconnected in a TSN network as described in the architecture, application, and security models in section 4, we want to determine a system configuration consisting of:

- an interval duration *P_{int}* for TESLA operations,
- the routing of streams,
- the task schedule,
- the network schedule as 802.1Qbv Gate-Control Lists (GCLs),

such that:

- all deadline requirements of all applications are satisfied.
- the redundancy requirements of all streams and the security conditions of TESLA are fulfilled.
- the overall latency of applications is minimized.

5.1 Motivational Example

We illustrate the problem using the architecture and application from Figure 4. We have one application Figure 4b with 4 tasks, 2 streams and a period and deadline of 1000 μs . The tasks are mapped to the end-systems as indicated in the figure. Stream s_2 will be multicast. The size of both streams is 50 B. For TESLA's security requirements, i.e. $s_1 \cdot s_l = s_2 \cdot s_l = 1$, we generate two additional security applications (Figure 5).

We have a TSN network with a link speed of 10 Mbit/s and zero propagation delay. Our TESLA implementation uses keys that are 16 B and MACs that are 16 B. A hash computation takes 10 μ s on every ES.

A solution that does not consider the security and redundancy requirements is shown in Figure 6a. With the TSN stream isolation constraint outlined in section 2 taken into consideration, the GCLs



(a) Schedule without security & redundancy



(b) Schedule with security & redundancy



are equivalent to frame schedules. We depict in Figure 6a the GCLs as a Gantt chart, where the red rectangles show the transmission of streams s_1 and s_2 on network links, and the blue rectangles show the tasks' execution on the respective end-systems. To guarantee deterministic message transmission in TSN, we have to isolate the frames in the time (or space) domain, leading to the delay of s_1 and thus t_3 . We refer the reader to [9] for an in-depth discussion on the non-determinism problem and isolation solution in TSN.

In this paper, we are interested in solutions such as the one in Figure 6b, which considers both the redundancy and security requirements. The black dashed line in the figure separates the TESLA key release intervals, where P_{int} was determined to be 500 μs . Streams carrying keys are orange, key generation tasks pink, key verification tasks green, and the MAC generation/validation operations on ESs are shown in red. The routing of the non-key streams can be seen in Figure 4a. Note how the two redundant copies of s_2 , s_2^0 and s_2^1 use non-overlapping paths.

Of particular importance is the delay incurred by the time-delayed release of keys: tasks t_3 and t_4 can only be executed after the keys authenticating s_1 and s_2 have arrived in the second interval, and after key verification and MAC validation tasks have been run.

Scheduling problems like the one addressed in this paper are NP-hard as they can be reduced to the Bin-Packing problem [11] and may be intractable for large input sizes. In the following sections, we will propose a Constraint Programming (CP) formulation to solve the problem optimally for small test cases, and a heuristic to solve the problem for large test cases.

Х	s1	s2_0	s2_1	
ES1	ES1	nil	nil	
ES2	nil	ES2	ES2	
ES3	SW1	SW1	SW2	
ES4	nil	SW1	SW2	
SW1	ES1	ES2	nil	
SW2	nil	nil	ES2	

Table 2 Matrix X for example from subsection 5.1

6 Constraint Programming Formulation

Constraint Programming (CP) is a technique to solve combinatorial problems defined using sets of constraints on decision variables. For large scheduling problems it becomes intractable to use CP due to the exponential increase in the size of the solution space [38]. In order to achieve reasonable runtime performance, we split the problem into 3 sub-problems which we solve sequentially: (i) finding a route for all streams, (ii) finding P_{int} , and (iii) finding the network and task schedule.

6.1 Optimizing redundant routing

The first step of solving the proposed problem is to find a set of (partially) disjoint routes for each stream, depending on the stream's redundancy level. The constraints in this section are inspired by [12] and [33].

We model the stream routes with an integer matrix X, where the columns represent streams (including their redundant copies) and rows represent nodes of the network. An entry at the position of a stream s_n and a node n in this matrix referring to a node m, represents a link from m to n on the route of stream s_n . Alternatively the entry could be *nil*, in which case n is not part of the route.

Using the matrix X, we can construct the route for each stream bottom-up as a tree, by starting at the receiver nodes. See Table 2 for the matrix of our example.

To determine the route for each stream $s_n \in S$, for each node $n \in N$ we have the following optimization variables:

- $x(s_n,n)$ represents an entry of our matrix X. The domain of $x(s_n,n)$ is defined as: $D(x(s_n,n)) = \{m \in \mathcal{N} | l_{m,n} \in \mathcal{L}\} \cup \{n\} \cup \{nil\}$. We refer to $x(s_n,n)$ as the successor of *n* on the path to the stream sender node.
- $y(s_n, n)$ represents the length of the path from *n* to $s_n.t_s.e$, i.e. the length of the path from node *n* to the sender node of the stream. $D(y(s_n, n)) = \{i | 0 \le i \le |SW| + 1\}$

Furthermore, we define a few helper variables and functions. First off, we define S^d as the set of all distinct streams, i.e., excluding the redundant copies of streams with redundancy level (RL) greater than one. Additionally we define S_{s_d} as the set of all redundant copies (including the stream itself) of s_d . Then we define the following helper function:

$$xsum(s_d, n, m) = \sum_{s'_d \in \mathcal{S}_{s_d}} (x(s'_d, n) == m)$$
⁽¹⁾

This function allows us, for any given $s_d \in S^d$, to determine the number of redundant copies (including s_d itself) that use the link from *m* to *n* (nil is counted as zero).

Then we have the following constraint optimization problem:

$$Minimize: \sum_{s_n \in S} cost(s_n) \tag{RC1}$$

where

$$cost(s_n) = \sum_{n \in \mathcal{N} \setminus \{s_n, t_s, e\}} (x(s_n, n)! = nil)$$
(RC2)

s.t.

$$x(s_n, n) \neq nil \Rightarrow y(s_n, n) = y(f, x(s_n, n)) + 1,$$

$$\forall s_n \in \mathcal{S}, n \in \mathcal{N} \setminus \{s_n.t_s.e\}$$
(R1)

$$x(s_n,m) = nil \Leftrightarrow x(s_n,n) \neq m,$$

$$\forall s \in S, n, m \in N$$
(R2)

$$\forall s_n \in \mathcal{S}, \ n, m \in \mathcal{N}$$

 $x(s_n, n) \neq nil,$ (R3.1)

$$\forall s_n \in \mathcal{S}, \ n \in \{t_r.e | t_r \in s_n.T_d\}$$

$$x(s_n, s_n.t_s.e) = s_n.t_s.e, \tag{R3.2}$$

$$\forall s_n \in \mathcal{S}$$

$$x(s_n, n) = nil,$$

$$\forall s_n \in \mathcal{S}, \ n \in \mathcal{ES} \setminus \{t_r.e | t_r \in s_n.T_d\}$$
(R3.3)

$$y(s_n, s_n.t_s.e) = 0,$$

$$\forall s_n \in \mathcal{S}$$
(R4)

$$\sum_{s_d \in \mathcal{S}^d} \left((xsum(s_d, n, m) > 0) \times \frac{s_d \cdot b}{s_d \cdot T} \right) \le [m, n] \cdot s, \tag{R5}$$

$$n,m\in\mathcal{N}$$

$$\begin{aligned} x(s_n,n) &\neq x(s'_n,n), \\ \forall s_n \in \mathcal{S}, \ s'_n \in \mathcal{S}_{s_n} \setminus s_n, \ n \in \mathcal{N} \setminus \{s_n.t_s.e\}, \end{aligned}$$
(R6)

Please note that == and != are boolean expressions that evaluate to 1 if true and to 0 otherwise.

The cost function we are minimizing ((RC1),(RC2)) measures the length of the route of each stream. 1

The constraint (R1) prevents cycles in the route, as defined in [33]. The constraint (R2) disallows "loose ends", i.e., a node that has a successor/predecessor must have a predecessor/successor itself. Please note that we refer to the successor on the path from receiver to sender, i.e., the predecessor on the route. The constraint (R3.1) states that all receivers of a stream have to have a successor. The constraints (R3.2), (R3.3), and (R4) impose that the sender of the stream has itself as the successor, no other end-system has a successor, and the path length is 0 at the sender node, respectively. The constraint (R5) restricts the bandwidth usage of each link to be under 100%. If multiple copies of the same stream use the same link, only one of them is counted as consuming bandwidth since we assume that streams are intelligently split and merged using IEEE 802.1CB. The constraint (R6) forbids the routes of redundant copies of a stream to overlap at any point.

¹For some use cases, fully disjoint routes are not necessary. Refer to Appendix A for an updated formulation for this case

6.2 Optimizing P_{int}

To set up the TESLA protocol, we need to choose the parameter P_{int} . P_{int} is the duration of one key disclosure interval. It has a big influence on the latency of secure streams and thus on the feasibility/quality of the schedule.

When choosing P_{int} there is a trade-off between overhead and latency. A small P_{int} reduces the latency of secure streams but necessitates more key generation/verification tasks and key streams. Thus, we want to determine the maximum value of P_{int} for which the latency is still within all deadline bounds. To this end, we formulate constraints inspired by [47] for which we then determine the optimal solution. This value is used as a constant in the subsequent optimization of the schedule.

We introduce a new notation: For each application $\lambda_l \in \Lambda$ we define $\lambda_l . C$ to be the communication depth, i.e. the length of the longest path in the application graph where only edges with associated secure streams are counted (ES-internal dependencies and non-secure streams are ignored). This gives us a measure of the longest chain of secure communications within the application, which we can use to estimate the amount of necessary TESLA intervals.

Then we have the following formulation:

$$Maximize: P_{int} \tag{P0}$$

s.t.

$$\forall \lambda_l \in \Lambda, \quad P_{int} \cdot (\lambda_l . C + 1) \le \lambda_l . T \tag{P1}$$

$$H \mod P_{int} = 0 \tag{P2}$$

$$P_{int} \mod \gcd(\{\lambda_l, T | \lambda_l \in \Lambda\}) = 0 \quad or$$

$$P_{int} * n = gcd(\{\lambda_l, T | \lambda_l \in \Lambda\}), \quad n \in \mathbb{N}$$
(P3)

The constraint (P1) guarantees that P_{int} is small enough to accommodate the authentication of all secure streams for all applications. The communication depth λ_l .*C* of an application gives a lower bound of how many TESLA intervals are necessary to accommodate all these streams within the period of the application, since there have to be n + 1 intervals to accommodate the authentication of *n* secure streams.

The purpose of the constraints (P2) and (P3) is to align the TESLA intervals with the schedule. The (P2) makes P_{int} a divisor of the hyperperiod, while constraint (P3) makes P_{int} either a multiple or a divisor of the greatest common divisor of all application periods.

6.3 Optimizing scheduling

In this step, we want to find a schedule for all tasks and streams which minimizes the overall latency of streams while fulfilling all constraints imposed by deadlines, TESLA, and TSN. The routes for each stream and P_{int} are given by the previous scheduling steps and assumed constant here.

We define the following integer optimization variables:

- o_l^s as the offset of stream s on link or node l
- c_l^s as the transmission duration of stream s on link or node l
- a_l^s as the end-time of stream s on link or node l
- φ^s as the index of the earliest interval where stream s can be authenticated on any receiver
- \bullet o_n^t as the offset of task t (on node t.e)
- a_n^t as the end-time of task t (on node t.e)

As an example, let us assume a hyperperiod of 1000us and a stream *s* with a period of 500us. $o_l^s = 100, c_l^s = 50, a_l^s = 150$ would imply that the stream *s* is scheduled on link *l* in the following time intervals: (100, 150) and (600, 650).

Furthermore we define several helper variables. Let \mathcal{E}^s be the set containing all receiver endsystems of stream *s*:

$$\mathcal{E}^s = \{t.e \mid t \in s.T_d\}$$

Let \mathcal{R}^s be the set containing all links on the route of stream s as well as sender and receiver nodes:

$$\mathcal{R}^{s} = \{s.t_{s}.e\} \cup \mathcal{E}^{s} \cup \{l_{a,b} \mid x(s,b) = a, l_{a,b} \in \mathcal{L}\}$$

$$\tag{2}$$

Using these helper functions we define the following constraint-optimization problem for the task and network scheduling step:

$$Minimize: \sum_{\lambda_l \in \Lambda} cost(\lambda_l)$$
(CS1)

where

$$cost(\lambda_l) = max(\{a^t \mid t \in \Gamma_l\}) - min(\{o^t \mid t \in \Gamma_l\})$$
(CS2)

s.t.

$$cost(\lambda_l) \le \lambda_l.T$$

$$\forall \lambda_l \in \Lambda$$
(S1)

$$o_l^s = c_l^s = a_l^s = 0, (S2.1)$$

$$\forall s \in \mathcal{S}, \ l_{a,b} \in \mathcal{L}, \ l_{a,b} \notin \mathcal{R}^s$$

$$o^s = c^s = o^s = 0. \tag{S2.2}$$

$$\forall s \in \mathcal{S}, \ n \in \mathcal{N}, \ n \notin \mathcal{R}^{s}$$

$$e^{s} + e^{s} = e^{s}$$
(S3.1)

$$\forall s \in \mathcal{S}, \ l_{a,b} \in \mathcal{L}, \ l_{a,b} \in \mathcal{R}^s$$
(33.1)

$$o_n^s + c_n^s = a_n^s,$$

$$\forall s \in \mathcal{S}, \ n \in \mathcal{N}, \ n \in \mathcal{R}^s$$
(S3.2)

$$c_l^s = \left\lceil \frac{s.b}{l.s} \right\rceil,\tag{S4.1}$$

$$\forall s \in \mathcal{S}, \ l_{a,b} \in \mathcal{L}, \ l_{a,b} \in \mathcal{R}^s$$

$$c_n^s = n.H,$$
(S4.2)

$$\forall s \in \mathcal{S}, n \in \mathcal{N} \cap \mathcal{R}^s, s.secure == 1$$

The constraint (S1) sets the deadline for the completion of an application to its period. The constraints (S2.1) and (S2.2) set all optimization variables to zero for every stream, for all nodes and links not part of its route. For all other links and nodes constraints, (S3.1) and (S3.2) set the end-time to be the sum of offset a length. For each link on the route of a stream constraint (S4.1) sets the length to be the byte-size of the stream divided by the link-speed. In constraint (S4.2) the length of secure streams on end-systems is set to the length of one hash-computation on that end-system, approximating the duration of MAC generation/verification.

$$\varphi^{s} > \left\lfloor \frac{a_{n}^{s}}{P_{int}} \right\rfloor,$$

$$\forall s \in S, \ l_{n,k} \in f \cap \mathcal{R}^{s}, \ h \in \mathcal{E}^{s}, \ s \ secure = 1$$
(S5)

$$\sqrt{s} \in \mathcal{S}, \ t_{a,b} \in \mathcal{L} + \mathcal{K}, \ b \in \mathcal{C}, \ s.secure = -1$$

$$\sqrt{s} \geq a^{t_{key}} + \varphi^{s} * P_{int}$$

$$\forall s \in S, \ n \in \mathcal{E}^{s} \ s \ sacura = 1$$

$$(S6)$$

$$\forall s \in \mathcal{S}, n \in \mathcal{E}^{s}, s.secure = 1$$

 $\forall t_{key} \in T_{ky}^{n}$

$$a_{l_{a,b}}^{s} \le o_{l_{b,c}}^{s}$$

$$\forall s \in \mathcal{S}, \ l_{b,c} \in \mathcal{L} \cap \mathcal{R}^{s}$$
(S7.1)

$$a = x(s,b)$$

$$a_a^s \le o_{l_a,b}^s \tag{S7.2}$$

$$\forall s \in \mathcal{S}, \ s.secure == 1,$$

$$l_{a,b} \in \{l_{a,b} \mid l_{a,b} \in \mathcal{L} \cap \mathcal{R}^{s}, \ a = s.t_{s}.e\}$$

$$a_{l_{a,b}}^{s} \leq o_{b}^{s}$$

$$\forall s \in \mathcal{S}, \ s.secure == 1,$$

$$l_{a,b} \in \{l_{a,b} \mid l_{a,b} \in \mathcal{L} \cap \mathcal{R}^{s}, \ b \in \mathcal{E}^{s}\}$$

$$(S7.3)$$

In constraint (S5) the earliest authentication interval for a stream φ^s is bound to be after the latest interval where the stream is transmitted. In constraint (S6) the start time of the stream on any receiver end-system is then bound to be greater or equal to the start time of that interval plus the end-time of the necessary preceding key verification task. The constraints (S7.1), (S7.2) and (S7.3) make sure that every stream is scheduled consecutively along its route. Hereby constraint (S7.1) enforces the precedence among two links, (S7.2) among the MAC generation on the sender and the first link and (S7.3) among the last link and the following MAC verification.

$$\begin{aligned} (\alpha \times s_{1}.T + a_{l}^{s_{1}} \leq \beta \times s_{2}.T + o_{l}^{s_{2}}) &\lor (\beta \times s_{2}.T + a_{l}^{s_{2}} <= \alpha \times s_{1}.T + o_{l}^{s_{1}}) \\ &\forall s_{1}, s_{2} \in \mathcal{S}, s_{1} \neq s_{2}, \forall l \in \mathcal{R}_{1}^{s} \cap \mathcal{R}_{2}^{s}, \\ &\forall \alpha \in \{0, ..., lcm(s_{1}.T, s_{2}.T)/s_{1}.T\}, \quad \forall \beta \in \{0, ..., lcm(s_{1}.T, s_{2}.T)/s_{2}.T\} \\ &(\alpha \times s_{2}.T + o_{l_{b,c}}^{s_{2}} <= \beta \times s_{1}.T + o_{l_{a_{1},b}}^{s_{1}}) \lor (\beta \times s_{1}.T + o_{l_{b,c}}^{s_{1}} <= \alpha \times s_{2}.T + o_{l_{a_{2},b}}^{s_{2}}) \\ &\forall s_{1}, s_{2} \in \mathcal{S}, s_{1} \neq s_{2}, \forall l \in \mathcal{R}_{1}^{s} \cap \mathcal{R}_{2}^{s}, \\ &a_{1} = x(s_{1},b), \ a_{2} = x(s_{2},b), \\ &\forall \alpha \in \{0, ..., lcm(s_{1}.T, s_{2}.T)/s_{1}.T\}, \quad \forall \beta \in \{0, ..., lcm(s_{1}.T, s_{2}.T)/s_{2}.T\} \end{aligned}$$

The constraint (S8) prevents any streams from overlapping on any nodes or links. Furthermore, constraint (S9) guarantees that for each link connected to an output port of a switch, the frames arriving on all input ports of that switch that want to use this output port cannot overlap in the time domain. This is the frame isolation necessary for determinism in our TSN configuration, which is further explained in [9].

$o^t + t.w = a^t$	(T1)
$\forall t \in \mathcal{T}$	
$a^t \leq o^s_{t,e}$	(T2.1)
$\forall t \in \mathcal{T}, s \in \mathcal{S}, s.t_s = t, s.secure == 1$	
$a^t \leq o^s_{l_{a,b}}$	(T2.2)
$\forall t \in \mathcal{T}, s \in \mathcal{S}, s.t_s = t, s.secure == 0$	
$\forall l_{a,b} \in \mathcal{L} \cap \mathcal{R}^s, \ a == t.e$	
$a_{t.e}^s \leq o^t$	(T3.1)
$\forall t \in \mathcal{T}, s \in \mathcal{S}, t \in s.T_d, s.secure == 1$	
$a^s_{l_{a,b}} \leq o^t$	(T3.2)
$\forall t \in \mathcal{T}, s \in \mathcal{S}, t \in s.T_s, s.secure == 0$	
$orall l_{a,b} \in \mathcal{L} \cap \mathcal{R}^s, \ b \in \mathcal{E}^s$	
$(\alpha \times t_1.T + a^{t_1} \le \beta \times t_2.T + o^{t_2}) \lor$	(T4)
$(\boldsymbol{\beta} imes t_2.T + a^{t_2} \le \boldsymbol{\alpha} imes t_1.T + o^{t_1})$	
$\forall t_1, t_2 \in \mathcal{T}, \ t_1 \neq t_2,$	
$\forall \boldsymbol{\alpha} \in \{0,, lcm(t_1.T, t_2.T)/t_1.T\},\$	
$orall eta \in \{0,,lcm(t_1.T,t_2.T)/t_2.T\}$	
$(\alpha \times t.T + a^t \le \beta \times s.T + o^s_{t.e}) \lor$	(T5)
$(\boldsymbol{\beta} \times s.T + a_{t.e}^s \le \boldsymbol{\alpha} \times t.T + o^t)$	
$\forall t \in \mathcal{T}, s \in \mathcal{S}, s.secure == 1, t.e \in \mathcal{R}^{s}$	
$\forall \boldsymbol{\alpha} \in \{0,, lcm(t.T, s.T)/t.T\},\$	
$\forall \boldsymbol{\beta} \in \{0,, lcm(t.T, s.T) / s.T\}$	

The constraint (T1) sets the end-time of a task to be the sum of offset and length. The constraints (T2.1) and (T2.2) model the dependency between a task and all its outgoing streams: such streams may only start after the task has finished. Similarly, constraints (T3.1) and (T3.2) model the dependency between a task and its incoming streams: such a task may only start after all incoming streams have arrived. Finally, constraint (T4) prevents any two tasks from overlapping, while constraint (T5) prevents a task from overlapping with a MAC generation/verification operation.

7 Metaheuristic Formulation

As mentioned in section 6, the scheduling problem addressed in this paper is NP-hard. As a consequence, a pure CP formulation solved using a CP solver is not tractable for large problem sizes. Hence, in this section, we propose a metaheuristic-based strategy, which aims to find good solutions (without the guarantee of optimality) in a reasonable time, even for large test cases.

An overview of our strategy is presented in algorithm 1. We use a Simulated Annealing (SA) metaheuristic [23] to find solutions $\Phi = (\mathcal{R}, \Sigma)$, consisting of a set of routes \mathcal{R} and a schedule Σ . As an input, we provide our architecture model $(\mathcal{N}, \mathcal{L})$ and the application model Λ . SA randomly explorers the solution space in each iteration by generating "neighbors" of the current solution using design transformations (or "moves"). We consider both routing and scheduling-related moves, and

Algorithm 1 Simulated Annealing Metaheuristic **Function** *heuristic*($\mathcal{N}, \mathcal{L}, \Lambda, T_{start}, \alpha, k, p_{rmv}, a, b, w$) 1 2 $\Phi_{best} = \Phi = \text{InitialSolution}(\mathcal{N}, \mathcal{L}, \Lambda, k);$ $c_{best} = c = Cost(\Phi, a, b);$ 3 $t = T_{start};$ 4 while stopping-criterion not True do 5 $\Phi_{new} = \text{RandomNeighbour}(\Phi, p_{rmv});$ 6 $c_{new} = \text{Cost}(\Phi_{new}, a, b);$ 7 $\delta = c_{new} - c;$ 8 if $\delta < 0$ or random $(0,1) < e^{-\frac{\delta}{t}}$ then 9 $\Phi = \Phi_{new};$ 10 $c = c_{new};$ 11 if $c_{new} < c_{best}$ then 12 $\Phi_{best} = \Phi_{new};$ 13 $c_{best} = c_{new}$ 14 $t = t * \alpha;$ 15 end 16 return Φ_{best} ; 17

the choice is controlled by a p_{rmv} parameter that gives the probability of a routing move. To measure the quality of a solution we use a cost function with two parameters *a* and *b* which are factors for punishing overlap of redundant streams and missed deadlines for applications, respectively. While we always accept better solutions, the central idea of Simulated Annealing is to also accept worse solutions with a certain probability in order to not get stuck in local optima [7].

algorithm 1 shows the main loop of the heuristic. We start out with an initial solution, a cost value, and a positive temperature. (line 2-4). Then, we repeat the steps described below until a stopping criterion like a time- or iteration-limit is met. We create a slight permutation of the current solution Φ by using the *RandomNeighbour* function (line 6). We calculate the cost of the new solution (line 7) and a delta of the new and old cost (line 8). Now, if the delta is smaller than 0, i.e., if Φ_{new} is a better solution than Φ , we choose Φ_{new} as the current solution (line 10-12). Alternatively, the new solution is also accepted if a random chosen value between 0 and 1 is smaller than the value of the acceptance probability function $e^{-\frac{\delta}{t}}$. This acceptance probability will decrease with the temperature over time and is also influenced by δ , which gives a measure of how much worse the new solution is. Finally, since we will occasionally accept worse solutions, we keep track of the best cost achieved overall and adjust it if necessary (line 12-14).

7.1 Precedence graph

We introduce a helper data structure in the form of a precedence graph. A precedence graph is a collection of special DAGs, one for each application. These DAGs are expanded versions of the DAGs from the application model. Here, streams are modeled as nodes instead of edges, and each redundant copy of a stream has its own node. See Figure 7 for an example. This data structure helps to model all the dependencies between tasks and streams in the scheduling algorithm. Additionally, we will use the set of all topological orders of this graph as our solution space for the scheduling step. An order can be seen as a scheduling priority assignment that respects all precedence constraints.

7.2 Initial solution

In the beginning, we create an initial solution Φ from the given architecture and application model. A solution is a tuple (\mathcal{R}, Σ) consisting of a set of routes \mathcal{R} and a schedule Σ . algorithm 2 details the function to find the initial solution.

To find an initial set of routes, we iterate through all streams and all pairs of sender and receiver ES (lines 2-3). For each such pair, we calculate and store k shortest paths for the given topology (line 5). For each redundant copy of a stream beyond the first, we calculate the shortest path in a weighted graph, where we weight all link used by previous copies with w instead of 1 (line 7). For the initial solution, we choose the shortest path for each pair (line 8). Note that our k-shortest-path algorithm only generates paths without repeated nodes that do not traverse any end-system.

To find an initial schedule, we have to create the precedence graph P (line 11) and decide an order O of this graph.

For the initial solution, we construct an order on the level of applications, i.e., we avoid interleaving nodes of different applications. We prioritize key applications (lines 12-14) before other (normal) applications (lines 15-17). This order is consequently used to create a schedule (line 19). See Figure 7 for an example order.

7.3 Neighbourhood function

The neighbourhood function *RandomNeighbour*(Λ , p_{rmv}) is detailed in algorithm 3. It is used during Simulated Annealing to create a slight permutation of a given solution/candidate Φ . It contains two fundamental moves: Changing the routing Λ . \mathcal{R} or changing the schedule Λ . Σ . Which move is taken is decided randomly (line 3). The parameter p_{rmv} influences how likely it is that the routing move is taken, e.g., $p_{rmv} = 0.5$ would result in a probability of 50%.

A routing move consists of choosing a random stream *s* out of the set of all streams (line 4), choosing a random receiver e_r out of all receivers of that stream (line 5) and then assigning a random path out of the set of k-shortest-paths calculated during the creation of the initial solution (line 6).

A scheduling move consists of choosing two random normal (non-key) applications d_1 and d_2 (lines 8,9), switching their order O in the precedence graph P (line 10) and recalculating the schedule (line 11). Whenever a new schedule is calculated, we also optimize its latency (line 12). This is further explained in subsection 7.6.



Figure 7 Example precedence graph with associated order

Algorithm 2 InitialSolution

1 F	Sunction InitialSolution($\mathcal{N}, \mathcal{L}, \Lambda, k, w$)
	// routing
2	foreach $s \in S$ do
3	foreach $e_r \in \{t_r.e t_r \in s.T_d\}$ do
4	if IsFirstCopyOfStream(s) then
5	$K_s^{e_r} = \text{ShortestPaths}(s.t_s.e, e_r, k, \mathcal{N}, \mathcal{L});$
6	else
7	$K_s^{e_r}$ = ShortestPathsWeighted(<i>s.t_s.e, e_r, k, N, L, w</i>);
8	$\Phi . R^s = \text{ShortestPath}(K_s^{e_r});$
9	end
10	end
	// schedule
11	$P = CreatePrecedenceGraph(\Lambda);$
12	foreach $\lambda_s \in \Lambda^{sec}$ do
13	$O = O \cup \text{TopologicalOrder}(\lambda_s, P);$
14	end
15	foreach $\lambda_n \in \Lambda \setminus \Lambda^{sec}$ do
16	$O = O \cup \text{TopologicalOrder}(\lambda_n, P);$
17	end
18	$\Phi K = K; \Phi P = P; \Phi O = O;$
19	$\Phi.\Sigma = \text{Schedule}(O, \Phi.\mathcal{R});$
20	return Φ;

Algorithm 3 RandomNeighbour

1 F	Sunction RandomNeighbour(Φ , p_{rmv})
2	p = random[0, 1];
3	if $p < p_{rmv}$ then
4	$s = RandomStream(\Phi);$
5	$e_r = \text{RandomReceiver}(s);$
6	$\Phi.\mathcal{R}^s = \text{RandomPath}(\Phi.K_s^{e_r});$
7	else
8	$d_1 = \text{RandomNormalApplication}(\Phi);$
9	$d_2 = \text{RandomNormalApplication}(\Phi);$
10	$\Phi.O = $ SwitchSchedulingOrder($d_1, d_2, \Phi.O$);
11	$\Phi.\Sigma = \text{Schedule}(\Phi.O, \Phi.\mathcal{R});$
12	$\Phi.\Sigma = \text{OptimizeLatency}(\Phi.\Sigma, \Phi.P);$
13	end
14	return Φ;

7.4 Cost function

The cost function is used in the simulated annealing metaheuristic to evaluate the quality of a solution. A lower cost means a better solution. algorithm 4 shows how our cost function is calculated. It consists of two components: a routing cost c_{route} and a schedule cost c_{sched} . The routing cost is the sum of the number of overlaps of redundant stream (one for each stream for each link) which is punished with a factor *a* and the total accrued length of all routes. The schedule cost is the sum of the number of infeasible applications, which is punished with a factor *b*, and the total sum of all application latencies (distance between start-time of first task and end-time of the last task). The factors *a* and *b* should be sufficiently high such that solutions with less overlap and infeasible applications are preferred.

Algorithm 4 Cost

1 Function $Cost(\Phi, a, b)$ 2 $| c_{route} = a * Overlaps(\Phi.\mathcal{R}) + Length(\Phi.\mathcal{R});$

3 $c_{sched} = b * \text{Infeasible}(\Phi, \Sigma) + \text{Latency}(\Phi, \Sigma);$

4 **return** $c_{route} + c_{sched}$;

7.5 ASAP list scheduling

To calculate a schedule for a given precedence graph with associated order and routing, we use an ASAP list-scheduling heuristic [41], which schedules each node of the precedence graph in the given order.

The algorithm, presented in algorithm 5, starts by iterating through each entry n of the given order O (line 2). An entry may either be a task or a stream. For each entry, we determine where it will be scheduled and create an indexable list L with all these locations (line 3). For a task, that set would contain just one end-system, while for a stream, it may contain many links (which are synonymous to an output port of a switch/ES) and also multiple end-systems, if the stream is secure, thus requiring MAC generation/verification.

Using these locations we also create a set of blocks (line 4). A block *b* is a tuple $(e, l, o, \underline{o}, \overline{o}, prev, next)$ which is associated to an entry *e* (task/stream) and a location *l* (node/link). *o* represents the block offset. \underline{o} and \overline{o} are parameters representing a lower and upper bound on the offset, which are used during the algorithm. The set *B* is implemented as a linked list, where *prev* and *next* are references to neighboring blocks on the route *L*. Note that in the case of multicast streams *next* could contain references to multiple blocks.

We now iterate over all these blocks (lines 7-8). For each block we begin by calculating the lower bound on the offset (line 9)². Usually, this lower bound is going to be the end-time (offset+length) of the block on the previous link, making sure that a stream is scheduled consecutively along its route. The first block is the maximum of all end-times of the last blocks of the predecessors of the current entry *n* in the precedence graph. For example for application λ_1 in Figure 7, the lower bound of the offset of the block of t_3 would be set to the maximum of the end-times of the last blocks of s_1 , s_2^0 and s_2^1 .

Also, for a secure stream, for all blocks on receiver ESs (i.e., MAC validation tasks), the lower bound is set to the end-time of the corresponding key verification task in the TESLA interval after the stream was received on the ES, since, according to the TESLA security condition, the stream can only be authenticated from that point on.

²The algorithm can be found in Appendix B

Algorithm 5 Scheduling - ASAP Heuristic

1 F	Function Schedule(P, \mathcal{R})
2	foreach $n \in O$ do
3	$L = GetRoute(n, \mathcal{R});$
4	B = CreateBlocks(n, L);
5	1 = L[0];
6	i = 0;
7	while true do
8	b = B[1];
9	$b.\underline{o} = \text{CalculateLowerBound}(n, b, P, \mathcal{R});$
10	o = EarliestOffset(b, 1);
11	if $o == \infty$ then
12	return false;
13	else if $o \leq b.\overline{o}$ then
14	b.o = o;
15	foreach $g \in b.next$ do
16	if IsBlockOnLink(g) then
17	$g.\overline{o}$ = LatestQueueAvailableTime(g, o);
18	end
19	i = i + 1;
20	if $i < len(L)$ then
21	1 = L[i];
22	else
23	break;
24	end
25	else
26	g = b.prev;
27	$g.\underline{o} = \text{EarliestQueueAvailableTime}(b, o);$
28	l = b.prev.l;
29	i = L.indexOf(l);
30	end end
31	end
32	$\Sigma = \text{UpdateSchedule(B)};$
33	end
34	return Σ;

N. Reusch et. al.



(c) Step 3

In the next step, the earliest possible offset for the current block is calculated (line 10). This function returns the earliest offset greater or equal to the lower bound within the feasible region. For more detail see subsubsection 7.5.1.

If such an offset is found and it is smaller than or equal to the upper bound, we can assign it to the block (line 14). We then iterate through each of the following blocks and set their upper bound to the latest point in time when their node is available and has been since the offset (line 15-18). This is done to fulfill the TSN constraint which forbids different streams to interleave within a queue (c.f. [36], [9] for a more detailed explanation).

If such an offset is found but it is larger than the upper bound, it is impossible to schedule the block while the port is still available, i.e., without it interleaving with other streams (line 25). Consequently, we have to backtrack and schedule the previous block at a later time. Therefore we set the lower bound *of the previous block* to the earliest time when the current port is available and remains so until the offset (line 26-27).

Figure 8 gives an example of this process. In step 1, s_1 has already been scheduled, and we are in the process of scheduling s_2 . We have scheduled the first block on l_{e_2,sw_1} and are now trying to schedule the second one on l_{sw_1,e_3} . The lower bound of our offset \underline{o} is set to the end-time of the first block. The upper bound \overline{o} is set to the latest time after which l_{sw_1,e_3} is still available after the offset of the first block, i.e., the start time of s_1 on that link. Finally, we find the earliest offset o to be only after

Figure 8 Backtrack example: Scheduling *s*₂

the end time of s_1 . It cannot be earlier since then the blocks of s_2 and s_1 would overlap. However, scheduling s_2 at that time is not possible since it would mean that the two streams interleave at the same port. Consequently, in step 2, we backtrack and reschedule the first block of s_2 by setting the lower bound on its offset to the earliest time when its port is available and remains so until o. In step 3, we are able to schedule the second block of s_2 without problems.

Once we have successfully found an offset for each block, we can update the schedule (line 32). This will remove the found blocks *B* from the feasible region.

7.5.1 Calculating the earliest offset

Calculating the earliest offset (algorithm 6 shows the function) for a given block is an important part of the heuristic. It takes a block b as an input and calculates the feasible region for that block (line 2). It then returns the lowest possible time that is within the feasible region and greater or equal than the lower bound (lines 3-6).

Algorithm 6 ASAP Heuristic - EarliestOffset

1 Function <i>EarliestOffset(b)</i>								
	<pre>/* ordered set of intervals */</pre>							
2	I = GetFeasibleRegion(b);							
3	foreach $i \in I$ do							
4	$o = \max(b.\underline{o}, i.begin);$							
5	if <i>i.contains(o)</i> then							
6	return <i>o</i> ;							
7	end							

The function to calculate the feasible regions for a given block *b* is detailed in algorithm 7. We start by getting all free intervals on the node/link *b.l* for the period *b.e.T* of the block (line 3). This ensures that the feasible region does not include any previously scheduled blocks on that node/link. The function then proceeds to fill the data structure R_{feas} with the free intervals, while cutting of a piece with the length of the block *b* from the end of each such interval (lines 4-7). This makes the feasible region represent all feasible values for the *offset* of the block.

7.6 Optimizing the latency for secure streams

If the block is assigned to a link, we have to cut down the feasible region further. Due to the TSN isolation constraint, it is not allowed to transmit two different streams on the same port at the same time. Thus, we iterate here over all the subsequent blocks b_{next} of the current block b, i.e., the blocks on the next links/ES on the route of the stream associated with the block (line 9). If the next block is also assigned to a link (not to an ES), we iterate through all already scheduled blocks b_{other} on that link b_{next} . I. These are blocks from other streams with whose predecessors, wherever they are scheduled, we are not allowed to overlap. Thus, we cut the interval (b_{other} . prev.o, b_{other} .o) from the feasible region (line 13).

Figure 9 provides two examples of feasible regions, shown in green, for a stream s_2 on two different routes. Looking at Figure 9(a), note the free space at the end of the period and before s_1 on l_{sw_1,e_3} . Choosing an offset anywhere in this space would result in s_2 being scheduled outside its period or overlapping with s_1 . Choosing an offset in the first free space on l_{e_2,sw_1} would result in s_1 and s_2 being transmitted to the same port at the same time, breaking the TSN isolation constraint. Note how in Figure 9(b) this is not the case, since s_2 is transmitted to a different port (l_{sw_1,e_4}) than s_1 .



(b) Route 2

```
Figure 9 Feasible region example
```

Algorithm 7 ASAP Heuristic - GetFeasibleRegion

```
1 Function GetFeasibleRegion(b)
 2
       R_{feas} = \emptyset;
       B_{free} = \text{GetFreeIntervals}(b.l, b.e.T);
 3
       /* (i) Add all free intervals that could contain block b */
       foreach iv \in B_{free} do
 4
           if iv.end - Length(b) \ge iv.begin then
 5
 6
             R_{feas} = \text{AddToFeasibleRegion}(R_{feas}, (\text{iv.begin, iv.end - Length}(b)));
 7
       end
       if IsLink(b.l) then
 8
           /* (ii) Cut out the interval blocked by other streams on the
               next port (TSN Stream Isolation) */
           foreach b_{next} \in b.next do
 9
               if IsLink(b_{next}.l) then
10
                   foreach b_{other} \in GetAllBlocksForLink(b_{next}.l) do
11
12
                       if b_{other} \neq b_{next} then
                            R_{feas} = \text{CutFromFeasibleRegion}(R_{feas}, (b_{other}.prev.o, b_{other}.o));
13
                   end
14
           end
15
       return R<sub>feas</sub>;
16
```

After we have created a new schedule, we can apply some post-processing to minimize its latency. Since TESLA requires a separation of sending and receiving tasks into separate intervals and since we are using an ASAP heuristic, there can be a significant gap between those tasks, as can be seen in Figure 10(a), resulting in an increased latency. To minimize the latency, the algorithm in algorithm 8 will go through each secure stream of each application (line 4). It will use the *OptimizeLatencyForSecureStream* function in algorithm 9 to optimize each stream individually. This function shifts all instances of the given stream as close to the instances on the receiver end-system as possible, without breaking the TESLA constraint. It also has an optional boolean parameter. If that is set, it also shifts the sending task of the given stream (otherwise there would be no latency gain). However, when we are optimizing a redundant stream, i.e. a stream where multiple copies originate at the same task, said task should only be moved when the last copy is optimized (lines 6-11). Otherwise, we can shift it immediately (line 13)

	Algorithm	8	ASAP Heuristic -	0	ptimizeLatency
--	-----------	---	------------------	---	----------------

1 F	unction O	$ptimizeLatency(\Sigma, P)$
2	foreach	λ in Λ do
3	forea	ach $n \in TopologicalOrder(\lambda, P)$ do
4	i	f IsStream(n.e) and n.e.secure and $n.e \in S^d$ then
5		if $n.e.rl > 1$ then
6		foreach $s_r \in S_{n.e}$ do
7		if $s_r \neq n.e$ then
8		$n_r = \text{GetNode}(s_r, P);$
9		OptimizeLatencyForStream(n_r , False);
10		end
11		OptimizeLatencyForStream(<i>n.e</i> , True);
12		else
13		OptimizeLatencyForStream(<i>n</i> , True);
14	end	
15	end	

The *OptimizeLatencyForSecureStream* function in algorithm 9 works internally by looping through the list of receivers of the given stream (line 2, multiple in case of a multicast stream). It goes backwards through the linked list of blocks for the stream, starting with the block on the last link before the current receiver (line 4). For each block, it will increase the offset as much as possible (move them as far as possible to the right) (line 10). After changing the offset we update the schedule (line 11). Then we continue iterating through the linked list (lines 17-19). If we arrive at the last block and the *move_task* boolean is set, we finish by the offset of the sender task (lines 12-16).

8 Experimental Results

In this section, we evaluate our two solutions to the formulated problem: The Constraint Programming formulation (referred to as *CP*, described in section 6) and the Simulated Annealing metaheuristic (referred to as *SA*, described in section 7). We analyze their scalability, runtime and solution quality and evaluate the impact of added redundancy and security.

Both solutions were implemented in Python 3.9. We developed a software tool with a web-based interactive user interface to display the models and solutions, including a routing graph and the

N. Reusch et. al.

```
Algorithm 9 ASAP Heuristic - OptimizeLatencyForStream
1 Function OptimizeLatencyForStream(n, move_task)
         foreach es_{recv} \in receivers(n) do
 2
             b = BlockOnLink(es<sub>recv</sub>, n);
  3
             b_{prev} = b.prev;
  4
             t<sub>kv</sub> = GetKeyVerificationTask(n.src, e);
  5
             b_{kv} = \text{GetBlock}(t_{kv});
  6
             i = GetTESLAIntervalForBlock(b_{kv});
  7
             ub = i * t_{kv}.T - b_{prev}.L;
  8
             while b_{prev} \neq \emptyset do
  9
                  b_{prev}.o = \min(ub, b_{prev}.\overline{o});
 10
                  UpdateSchedule(b<sub>prev</sub>);
 11
                  if b_{prev}. prev == 0 and move_task and IsLastReceiver(es<sub>recv</sub>) then
 12
                      /* Also move the sender task closer to the first block of
                           the stream */
                      t<sub>sender</sub> = GetSenderTask(n);
 13
                      b_{sender} = \text{GetBlock}(t_{sender});
 14
                      ub = b_{prev}.o - b_{sender}.L;
 15
 16
                      b_{prev} = t_{sender};
                  else if b_{prev}. prev \neq \emptyset then
 17
                      ub = b_{prev}.o - b_{prev}.prev.L;
 18
 19
                      b_{prev} = b_{prev}.prev;
             end
20
21
         end
```



(b) Optimized stream

Figure 10 Latency optimization for secure streams

schedule.³ For solving the CP formulation we use the CP-SAT solver from Google OR-Tools [1]. For calculating k-shortest-paths in the metaheuristic we use the *shortest_simple_paths* function from the NetworkX [15] library. All evaluations were run on a High Performance Computing (HPC) cluster, with each node configured with 2xIntel Xeon Processor 2660v3 (10 cores, 2.60GHz) and 16 GB memory. Both CP and SA run on one node at a time

8.1 Test cases used for the evaluation

For the scalability evaluation we used the following test cases, see Table 3: the example presented in section 5 (*example*), a realistic automotive test case from a large automotive manufacturer (*auto*) [12], a medium-sized automotive case study from [22] (*case_study*) and 16 synthetic test cases of increasing size and complexity. The topology of the auto test case was adjusted to allow disjunct redundant routes.

For the redundancy/security impact evaluation, we used a further set of 100 synthetic test cases grouped into four batches.

We created the synthetic test cases to be as realistic as possible: They all feature secure streams, redundancy levels between 1 and 3, applications with complex dependencies and a realistic network topology that allows disjunct redundant paths.

To create realistic topologies, we developed a custom algorithm, as follows. For a given number of switches and end-systems, we create that many random points in 2D space. Then we connect each switch to its closest neighbor until every switch is connected to 4 other switches. Afterwards, we connect each end-system to the closest 3 switches.

To create realistic application DAGs, we used the GGen tool presented in [8] and the layer-bylayer method with a depth of 3 and a connection probability of 50%. If a DAG contains separate subgraphs these are split into separate applications. The application period is chosen randomly among the set {10, 15, 20, 50ms}. Nodes of the generated DAG are interpreted as tasks with a random WCET, upper bound at 6% of the period. Tasks are divided randomly between ES. All outgoing edges of a node in the DAG combined are interpreted as a stream, with the source node as sender task and the destination nodes as receiver tasks. The stream has a random size below or equal to 1.500 Bytes, with a random RL between 1-3 and a 30% probability to be considered security critical.

We used a global link speed of 1000 Mbit/s. TESLA uses 16 B keys and MACs. A hash computation takes 10 μ s on every ES.

8.2 Scalability evaluation

To evaluate the scalability, we ran both the CP and the SA solutions on the same test cases with the same computing resources. Table 3 shows the results for each test case for both solutions. The columns **# ES**, **# SW**, **# Streams**, **# Tasks** give the total number of ES, SW, streams and tasks respectively. **# Receiver Tasks** gives the total sum of stream receiver task (since we consider multicast streams, one stream can have multiple). The **Cost** column gives the total cost of the found solution following the cost function in algorithm 4. The **T** column shows the total runtime of the solver.

The CP solution was given a timeout of 60min. If CP failed to find an optimal solution in time, or ran out of memory, we reported Cost and T as empty "/". The SA solution was given a timeout of 10min (20min for the largest test case, giant1). We used ParamILS [17] to optimize the following parameters for the SA heuristic: T_{start} , α , k, p_{rmv} and w. a was set to 50000, b to 10000.

³The tool including the obtained results is available on GitHub: https://github.com/nreusch/TSNConf

Test case	Method	# ES	# SW	# Streams	# Recv. Tasks	# Tasks	Cost	Т
example	СР	4	2	6	10	9	467	1 s
example	SA	4	2	6	10	9	477	10 m
auto	СР	20	32	84	102	74	/	/
auto	SA	20	32	84	102	74	38031	10 m
case_study	СР	6	2	29	31	28	3771	130 s
case_study	SA	6	2	29	31	28	6114	10 m
tiny1	СР	4	2	2	2	6	1 708	0.2 s
tiny1	SA	4	2	2	2	6	1 708	10 m
tiny2	СР	4	2	3	4	6	1732	0.2 s
tiny2	SA	4	2	3	4	6	1732	10 m
tiny3	СР	4	2	11	13	15	7450	14 m
tiny3	SA	4	2	11	13	15	18088	10 m
small1	СР	8	4	10	16	20	5421	2 s
small1	SA	8	4	10	16	20	13 303	10 m
small2	СР	8	4	14	20	23	9110	60 m
small2	SA	8	4	14	20	23	13794	10 m
small3	СР	8	4	29	48	35	7 7 0 5	17.5 m
small3	SA	8	4	29	48	35	13781	10 m
medium1	СР	16	8	23	34	37	12991	4.5 m
medium1	SA	16	8	23	34	37	22883	10 m
medium2	СР	16	8	30	47	43	6552	5.2 m
medium2	SA	16	8	30	47	43	19455	10 m
medium3	СР	16	8	36	53	47	15515	60 m
medium3	SA	16	8	36	53	47	26486	10 m
large1	СР	32	16	47	86	73	/	/
large1	SA	32	16	47	86	73	43872	10 m
large2	СР	32	16	33	65	72	24953	25 m
large2	SA	32	16	33	65	72	41 0 26	10 m
large3	СР	32	16	69	170	104	/	/
large3	SA	32	16	69	170	104	34860	10 m
huge1	СР	64	32	84	183	133	/	/
huge1	SA	64	32	84	183	133	73070	10 m
huge2	СР	64	32	99	213	161	/	/
huge2	SA	64	32	99	213	161	57246	10 m
huge3	СР	64	32	99	197	169	/	/
huge3	SA	64	32	99	197	169	93357	10 m
giant1	СР	128	64	144	347	261	/	/
giant1	SA	128	64	144	347	261	101799	20 m

Table 3 Scalability tests

Note that the CP solver will return once the optimal solution is found, while the SA solver will always run until the timeout and return the best feasible (i.e. no missed deadlines or overlap) solution found up to that point. However, SA is able to find a first feasible solution in a very quick time. For all test cases in Table 3 it could find one in less than 10 s.

The table shows that CP is able to find solutions up to medium-sized test cases within the given

timeout, but it does not scale to the larger test cases. SA is scalable; it is able to find solutions even for the the largest test cases. This scalability comes at an increase in cost by 67% on average, which can be reduced by giving a longer timeout. This increase is mostly caused by increased application latencies (scheduling cost), which are still within the deadlines, while the routing cost is usually close to or equal to the optimal routing cost from the CP solution. The conclusion is that SA can be successfully used to route and schedule large realistic test cases, and its quality is comparable to the optimal solutions obtained by CP.

	Batch name	Security	Redundancy	Cost	Bandwidth	CPU
0	batch0 - large streams, small tasks	no	no	3822.08	0.09	1.44
1		no	yes	+1.34%	+25.45%	+0.00%
2		yes	no	+256.25%	+5.49%	+12.99%
3		yes	yes	+258.10%	+36.64%	+15.93%
4	batch1 - large streams, large tasks	no	no	17721.32	0.08	7.06
5		no	yes	+-0.00%	+15.91%	+0.00%
6		yes	no	+64.31%	+3.18%	+1.39%
7		yes	yes	+64.98%	+22.16%	+1.75%
8	batch2 - small streams, large tasks	no	no	18547.6	0.01	7.51
9		no	yes	+0.06%	+31.29%	+0.00%
10		yes	no	+52.86%	+33.92%	+0.21%
11		yes	yes	+53.12%	+97.96%	+0.70%
12	batch3 - small streams, small tasks	no	no	3713.32	0.01	1.55
13		no	yes	+0.37%	+26.62%	+0.00%
14		yes	no	+238.77%	+27.73%	+8.54%
15		yes	yes	+245.32%	+85.57%	+10.57%

8.3 Impact of adding redundancy and security to a test case

Table 4 Impact of security and redundancy measures

The feasible solutions fulfil the security and redundancy requirements of streams. These requirements introduce extra tasks and streams that need to be routed and scheduled, leading to an overhead compared to the case when we would ignore the security and redundancy requirements of an application. In this set of experiments, we were interested to evaluate the overhead fulfilling the redundancy and security requirements compared to the case these are ignored. These overheads were measured on solution cost, available bandwidth, and CPU resources. Hence, we created four batches of 25 synthetic test cases each. Each test case has a random topology with 8 switches and 16 end-systems, 24 tasks and multiple applications with random DAGs. Streams have a random RL between 1 and 3 and 30% probability to be considered security critical.

Each batch features either large (1000-1500 B) or small (1-250B) streams and either large ($\leq 10\%$ of period) or small ($\leq 2\%$ of period) tasks. Each batch was run 4 times using the first feasible SA solution with different combinations of enabled/disabled security and redundancy requirements. Disabled security means that all streams are set to a security level of 0, which disabled redundancy means that all streams are set to a redundancy level of 1.

Table 4 shows the results. We always take the results for the no-security, no-redundancy run as a baseline and note the percentual increase in total cost, total bandwidth occupation percentage and total CPU utilization percentage in the following rows. Bandwidth and CPU utilization are measured as the mean of the total utilization over all links and ESs respectively.

As can be seen, the impact of adding security and redundancy differs significantly, depending on the size of initial streams and tasks. Note that an increase in overhead is expected with an increase in the number and difficulty of the security and redundancy requirements.

Adding redundancy has a negligible impact on cost and CPU utilization but always has a significant impact on bandwidth. Adding security always has a significant impact on cost, as each application with secure streams has to be split into multiple TESLA intervals. The impact of adding security on bandwidth and CPU utilization depends largely on the relative size of streams and WCET of tasks compared to the TESLA overhead. For example, 16 bytes of overhead for a MAC a much more significant for a 100 B stream than for a 1000 B stream.

8.4 Discussion

Our proposed SA implementation is able to determine good solutions in a reasonable time, even for large test cases. In addition, it can find feasible solutions (where all timing, safety and redundancy requirements are satisfied) extremely quickly, within 10 s even for large test cases. This can be useful, e.g., for evaluating several architectures in terms of their monetary costs and redundancy allowed by the physical topology, prototyping or for rapid runtime reconfiguration in case of failures or changes in traffic patterns. Although CP can find optimal solutions, it does not scale for large test cases, and it is not flexible, that is, it will not report solutions which are not feasible. An advantage of SA is its ability to find return near-feasible solutions for those test cases that cannot be solved, i.e. solutions with some infeasible applications or overlapping streams. SA is able to point out the names of the offending apps/tasks and streams, which can give a good indication on where the configuration has to be improved to become feasible, e.g., by increasing the redundancy in the physical topology or by changing the mapping of tasks to ESs.

9 Conclusion

In this paper, we addressed the combined TSN routing and scheduling problem for complex applications with redundancy and security requirements.

We proposed TESLA as an efficient authentication protocol for use-cases with low-powered devices and multicast communication. We proposed a modification to the protocol that makes it more lightweight, made possible by the real-time guarantees of our network.

We developed two methods to solve the combined routing and scheduling problem: A Constraint Programming solution which can solve small and medium-sized test cases optimally and a solution that combines a Simulated Annealing metaheuristic and an ASAP list scheduling which can solve very large test cases.

We formalized the constraints governing our problem and came up with novel ways to handle the complexities introduced by TESLA and redundancy while calculating correct solutions in the heuristic. Furthermore, we developed and shared a useful tool for reuse of our solutions and interactive visualization of routes and schedules.

We evaluated the impact of adding security and redundancy to existing applications and showed that much the overheads depend on the size of existing tasks and streams.

— References –

¹ CP-SAT Solver Guide.

IEEE Standard for Local and metropolitan area networks – Bridges and Bridged Networks – Amendment 26: Frame Preemption. *IEEE Std 802.1Qbu-2016 (Amendment to IEEE Std 802.1Q-2014)* (Aug 2016), 1–52.

- 3 IEEE Standard for Local and metropolitan area networks Bridges and Bridged Networks Amendment 25: Enhancements for Scheduled Traffic. *IEEE Std 802.1Qbv-2015 (Amendment to IEEE Std 802.1Q— as amended by IEEE Std 802.1Qca-2015, IEEE Std 802.1Qcd-2015, and IEEE Std 802.1Q—/Cor 1-2015)* (March 2016), 1–57.
- 4 IEEE Standard for Local and metropolitan area networks–Bridges and Bridged Networks–Amendment 28: Per-Stream Filtering and Policing. IEEE Std 802.1Qci-2017 (Amendment to IEEE Std 802.1Q-2014 as amended by IEEE Std 802.1Qca-2015, IEEE Std 802.1Qcd-2015, IEEE Std 802.1Q-2014/Cor 1-2015, IEEE Std 802.1Qbv-2015, IEEE Std 802.1Qbu-2016, and IEEE Std 802.1Qbz-2016) (Sep. 2017), 1–65.
- 5 IEEE standard for local and metropolitan area networks–frame replication and elimination for reliability. *IEEE Std* 802.1CB-2017 (2017), 1–102.
- 6 ATALLAH, A. A., HAMAD, G. B., AND MOHAMED, O. A. Routing and scheduling of time-triggered traffic in time-sensitive networks. *IEEE Transactions on Industrial Informatics 16*, 7 (2020), 4525–4534.
- 7 BURKE, E. K., KENDALL, G., ET AL. Search methodologies. Springer, 2005.
- 8 CORDEIRO, D., MOUNIÉ, G., PERARNAU, S., TRYSTRAM, D., VINCENT, J. M., AND WAGNER, F. Random graph generation for scheduling simulations. *Simutools 2010 - 3rd International Icst Conference on Simulation Tools and Techniques* (2010).
- 9 CRACIUNAS, S. S., SERNA OLIVER, R., CHMELIK, M., AND STEINER, W. Scheduling real-time communication in IEEE 802.1Qbv Time Sensitive Networks. In 24th International Conference on Real-Time Networks and Systems (RTNS) (2016), ACM.
- 10 DÜRR, F., AND NAYAK, N. G. No-wait Packet Scheduling for IEEE Time-sensitive Networks (TSN). In Proc. RTNS (2016), ACM.
- 11 FALK, J., DÜRR, F., AND ROTHERMEL, K. Exploring practical limitations of joint routing and scheduling for tsn with ilp. In *Proc. RTCSA* (2018), pp. 136–146.
- 12 GAVRILUT, V., ZARRIN, B., POP, P., AND SAMII, S. Fault-tolerant topology and routing synthesis for ieee time-sensitive networking. In *Proceedings of the 25th International Conference on Real-Time Networks and Systems* (2017), Association for Computing Machinery, p. 267–276.
- 13 GAVRILUȚ, V., ZHAO, L., RAAGAARD, M. L., AND POP, P. Avb-aware routing and scheduling of time-triggered traffic for tsn. *IEEE Access* 6 (2018), 75229–75243.
- 14 GRAMMATIKAKIS, M. D., HSU, D., KRAETZL, M., AND SIBEYN, J. F. Packet routing in fixedconnection networks: A survey. *Journal of Parallel and Distributed Computing* 54, 2 (1998), 77–132.
- 15 HAGBERG, A. A., SCHULT, D. A., AND SWART, P. J. Exploring network structure, dynamics, and function using networkx. In *Proceedings of the 7th Python in Science Conference* (Pasadena, CA USA, 2008), G. Varoquaux, T. Vaught, and J. Millman, Eds., pp. 11 – 15.
- 16 HUANG, K., WAN, X., WANG, K., JIANG, X., CHEN, J., DENG, Q., XU, W., PENG, Y., AND LIU, Z. Reliability-aware multipath routing of time-triggered traffic in time-sensitive networks. *Electronics 10*, 2 (2021).
- 17 HUTTER, F., HOOS, H. H., LEYTON-BROWN, K., AND STÜTZLE, T. ParamILS: an automatic algorithm configuration framework. *Journal of Artificial Intelligence Research 36* (October 2009), 267–306.
- 18 IEEE. 802.1AS-Rev Timing and Synchronization for Time-Sensitive Applications. http://www. ieee802.org/1/pages/802.1AS-rev.html, 2016. Accessed: 11.06.2019.
- 19 IEEE. Official Website of the 802.1 Time-Sensitive Networking Task Group. http://www.ieee802. org/1/pages/tsn.html, 2016. Accessed: 11.06.2019.
- 20 ISSUING COMMITTEE: AS-2D2 DETERMINISTIC ETHERNET AND UNIFIED NETWORKING. SAE AS6802 Time-Triggered Ethernet. http://standards.sae.org/as6802/, 2011. retrieved 20-May-2014.
- 21 IZOSIMOV, V., POP, P., ELES, P., AND PENG, Z. Design optimization of time- and cost-constrained faulttolerant distributed embedded systems. In *Design, Automation and Test in Europe* (2005), pp. 864–869 Vol. 2.
- 22 KANDASAMY, N., HAYES, J. P., AND MURRAY, B. T. Dependable communication synthesis for distributed embedded systems. *Reliability Engineering and System Safety* 89, 1 (2005), 81–92.
- 23 KIRKPATRICK, S., GELATT, C. D., JR., AND VECCHI, M. P. Optimization by simulated annealing. *Science* 220 (1983), 671–680.

- 24 LAURSEN, S. M., POP, P., AND STEINER, W. Routing optimization of avb streams in tsn networks. *SIGBED Rev. 13*, 4 (Nov. 2016), 43–48.
- 25 MAHFOUZI, R., AMINIFAR, A., SAMII, S., ELES, P., AND PENG, Z. Security-aware routing and scheduling for control applications on ethernet tsn networks.
- 26 NAYAK, N. G., DÜRR, F., AND ROTHERMEL, K. Routing algorithms for ieee802.1qbv networks. *SIGBED Rev. 15*, 3 (Aug. 2018), 13–18.
- 27 OJEWALE, M. A., AND YOMSI, P. M. Routing heuristics for load-balanced transmission in tsn-based networks. *SIGBED Rev. 16*, 4 (Jan. 2020), 20–25.
- 28 PAHLEVAN, M., TABASSAM, N., AND OBERMAISSER, R. Heuristic list scheduler for time triggered traffic in time sensitive networks. *SIGBED Rev. 16*, 1 (Feb. 2019), 15–20.
- 29 PEREIRA, T., BARRETO, L., AND AMARAL, A. Network and information security challenges within industry 4.0 paradigm. *Procedia Manufacturing 13* (2017), 1253 1260.
- 30 PERRIG, A., CANETTI, R., SONG, D., AND TYGAR, J. D. Efficient and secure source authentication for multicast. In *Network and Distributed System Security Symposium*, NDSS (2001), vol. 1, pp. 35–46.
- 31 PERRIG, A., CANETTI, R., TYGAR, J. D., AND SONG, D. The tesla broadcast authentication protocol. *RSA CRYPTOBYTES* (2002), 2002.
- 32 PERRIG, A., SONG, D., CANETTI, R., TYGAR, J. D., AND BRISCOE, B. Timed efficient stream loss-tolerant authentication (tesla): Multicast source authentication transform introduction. RFC 4082, RFC Editor, June 2005.
- 33 PHAM, Q. D., AND DEVILLE, Y. Solving the quorumcast routing problem by constraint programming. Constraints 17, 4 (2012), 409–431.
- 34 POP, P., LANDER RAAGAARD, M., CRACIUNAS, S. S., AND STEINER, W. Design optimization of cyber-physical distributed systems using IEEE time-sensitive networks (TSN). *IET Cyber-Physical Systems: Theory and Applications 1*, 1 (2016), 86–94.
- 35 PRYTZ, G. A performance analysis of EtherCAT and PROFINET IRT. In *Proc. ETFA* (2008), IEEE Computer Society.
- **36** RAAGAARD, M. L., AND POP, P. Optimization algorithms for the scheduling of IEEE 802.1 Time-Sensitive Networking (TSN). Tech. rep., DTU Compute, Technical University of Denmark, 2017.
- 37 REUSCH, N., POP, P., AND CRACIUNAS, S. S. Work-in-progress: Safe and secure configuration synthesis for tsn using constraint programming. In 2020 IEEE Real-Time Systems Symposium (RTSS) (2020), pp. 387–390.
- 38 ROSSI, F., VAN BEEK, P., AND WALSH, T. Handbook of constraint programming. Elsevier, 2006.
- **39** SERNA OLIVER, R., CRACIUNAS, S. S., AND STEINER, W. IEEE 802.1Qbv Gate Control List Synthesis using Array Theory Encoding. In *Proc. Real-Time and Embedded Technology and Applications Symposium (RTAS)* (2018), IEEE.
- 40 SHEIKH, A., BRUN, O., CHÉRAMY, M., AND HLADIK, P.-E. Optimal design of virtual links in afdx networks. *Real-Time Systems* 49 (05 2013), 308–336.
- 41 SINNEN, O. Fundamental Heuristics. John Wiley & Sons, Ltd, 2007, ch. 5, pp. 108–144.
- 42 STEINER, W. An evaluation of SMT-based schedule synthesis for time-triggered multi-hop networks. In *Proc. RTSS* (2010), IEEE.
- 43 STEINER, W., BAUER, G., HALL, B., AND PAULITSCH, M. TTEthernet: Time-Triggered Ethernet. In *Time-Triggered Communication*, R. Obermaisser, Ed. CRC Press, Aug 2011.
- 44 STUDNIA, I., NICOMETTE, V., ALATA, E., DESWARTE, Y., KAANICHE, M., AND LAAROUCHI, Y. Survey on security threats and protection mechanisms in embedded automotive networks. *Proceedings of the International Conference on Dependable Systems and Networks* (2013), 6615528.
- 45 TĂMAŞ-SELICEAN, D., POP, P., AND STEINER, W. Design optimization of ttethernet-based distributed real-time systems. *Real-Time Syst. 51*, 1 (Jan. 2015), 1–35.
- **46** WANG, B., AND HOU, J. Multicast routing and its qos extension: problems, algorithms, and protocols. *IEEE Network 14*, 1 (2000), 22–36.
- 47 ZHAO, R., QIN, G., LYU, Y., AND YAN, J. Security-aware scheduling for ttethernet-based real-time automotive systems. *IEEE Access* 7 (2019), 85971–85984.

A Routing constraint formulation for forbidden overlap

To achieve a constraint formulation in which overlap is possible do the following: Replace (RC2):

$$cost(s_n) = length_cost(s_n) + 100 * overlap_cost(s_n)$$
 (RC2)

Introduce the following:

$$length_cost(s_n) = \sum_{n \in \mathcal{N} \setminus \{s_n.t_s.e\}} (x(s_n, n)! = nil)$$
(RC3)

$$overlap_cost(s_n) = \sum_{n \in \mathcal{N} \setminus \{s_n, t_s, e\}} \sum_{m \in \mathcal{N} \setminus \{n\}} link_cost(s_n, n, m)$$
(RC4)

$$link_cost(s_n, n, m) = (xsum(s_n, n, m) - 1) * (x(s_n, n) = m)$$
(RC5)

Remove (R6)

B Additional functions from metaheuristic formulation

B.1 CalculateLowerBound

Algorithm 10 ASAP Heuristic - CalculateLowerBound

```
1 Function CalculateLowerBound(n, b, P, \mathcal{R})
       lb = 0;
2
       if b.prev == \emptyset then
3
            /* If n is a task or the first stream instance */
            for
each n_{prev} \in Predecessors(n, P) do
 4
                b = LastBlock(n_{prev});
 5
                lb = max(lb, b.o + Length(b));
 6
            end
 7
       else if IsLink(b.l) then
 8
            /* If n is a stream and b.l is a link */
            foreach l_{prev} \in PredecessorLinks(b.l, n, \mathcal{R}) do
 9
                b_{prev} = BlockOnLink(b.l, n);
10
                lb = max(lb, b_{prev}.o + Length(b_{prev}));
11
            end
12
       else
13
            /* If n is a stream and l is a receiver end-system */
            t_{key}^{verify} = \text{GetKeyVerificationTask}(n, l);
14
           b_{key}^{verify} = \text{GetBlockForEntry}(t_{key}^{verify});
15
           i = GetTESLAIntervalForBlock(b_{key}^{verify});
16
           \mathbf{lb} = b_{key}^{verify}.o + i * b_{key}^{verify}.e.T + Length(b_{key}^{verify});
17
       end
18
19
       return max(lb, b.<u>o</u>);
```

B.2 BlockQueues

Algorithm 11 ASAP Heuristic - BlockQueues

```
1 Function BlockQueues(n, B, L)
      for
each l \in L do
2
         /* Calculate blocks for each frame of n */
         i = 0;
3
         foreach b \in B do
 4
             offsets[i] = b.o + i * b.e.T;
 5
             endtimes[i] = b.o + i * b.e.T + Length(b);
 6
 7
             i = i + 1;
         end
8
         /* Block queues/end-systems */
         foreach T \in Periods do
9
             for i = 0 to len(offsets)-1 do
10
                 o = offset[i];
11
                 e = endtimes[i];
12
                 if e\%T < o\%T then
13
                    /* Handle wrap around period border */
                    CutFromFeasibleRegion(1, o\%T, T);
14
                    CutFromFeasibleRegion(l, 0, e\%T);
15
                 else
16
                   CutFromFeasibleRegion(l, o\%T, e\%T);
17
                 end
18
             end
19
         end
20
      end
21
```