

End-to-end Orchestration in Support of IIoT Applications over Optically Interconnected TSN Domains

Albert Pagès, Fernando Agraz, Salvatore Spadaro

Universitat Politècnica de Catalunya (UPC), Optical Communications Group (GCO), Barcelona (08034), Spain

Author e-mail address: albert.pages-cruz@upc.edu

Abstract: We demonstrate that joint orchestration of TSN and optical network domains in support of IIoT applications reduces the TSN blocking by four orders of magnitude and the usage of high priority queues by a 28-100%. © 2022 The Authors.

1. Introduction

One of the main requirements that future 6G network infrastructures impose is the concept of Time Sensitive Networks (TSNs) [1], which relates to services/applications with bounded latency times as well as the network infrastructures supporting them. One clear example is Industrial IoT (IIoT) applications, in which industrial assets (e.g., industrial assembly lines) require a coordinated and time-bounded operation to effectively execute their tasks. To this purpose, the development of network infrastructures with support for network determinism has been key in the recent years [2, 3], providing means for traffic scheduling and shaping based on the applications' time requirements [4].

In IIoT, control loops refer to a set of functions that coordinate the operation of the various industrial assets (e.g., mechanical arms). Several of them require very low communication latencies between the factory floor and the control application (e.g., few to tens of μ s), thus, needing to be deployed at in-house computational or edge computing premises to meet the stringent latency times. Others, however, although still requiring bounded latencies, have less stringent time requirements (in the order of ms). This opens the possibility to move such control applications to cloud premises (e.g., datacenters (DCs)) to reduce the related capital and operational expenditures. Nevertheless, low latency and deterministic network infrastructures are thus required between the TSN domains hosting the IIoT infrastructures and the cloud domains hosting the control applications. Optical networks are seen as prime candidates to handle this role. However, the assignment of resources at all involved domains (TSN/cloud and optical network) requires for a coordinated control and orchestration of the end-to-end (E2E) infrastructure to meet desired latency requirements, with the added benefits of optimizing the overall resource utilization.

2. End-to-end control and orchestration of optically-interconnected TSN domains

Several TSN domains here represent infrastructures in which both cloud premises and IIoT platforms may be present. The cloud domain hosting the IIoT control functions needs to be connected to the factory floor domain to implement the control loop. In this regard, TSN-able infrastructures support the associated data flows (TSN traffic). From a resource perspective, it translates in a set of queues for supporting data flows, with each queue characterized by a priority that allows for a bounded service time under certain threshold, and a queue capacity. Thus, at both source and destination domains, the assignment of suitable queues is required. Then, both domains are interconnected by means of optical connections (lightpaths (LPs)) across a Wavelength Division Multiplexing (WDM) optical network. Aside from TSN traffic, TSN domains also generate/receive best-effort (BE) traffic due to other non-time-sensitive operations (e.g., regular backups). As such, there is a resource competition in both network capacity and achievable latency. In a typical per-domain operation, due to the lack of cross-domain information, each domain is forced to assign resources at the best of its capabilities, since the achievable E2E performance is unknown beforehand. This results in that high priority queues, with lower service latencies, may be assigned to TSN traffic which does not really require such low service times to meet its E2E latencies. In addition, low latency LPs may be assigned to both BE and lower priority TSN traffic indistinctively, saturating such resources for TSN traffic with strict latency budgets. As a consequence, there may be a significant number of TSN traffic that is blocked due to E2E latency unfulfillments. Thus, an orchestrated approach to resource assignment is required. We define an orchestration strategy that pursues the minimization of blocking for TSN traffic due to latency reasons. Moreover, the use of low latency optical networks allows to redistribute the E2E latency budget to favor the assignment of low priority queues with larger service times, lowering the traffic scheduling complexity at TSN domains, which then will be using high priority resources (queues) to the minor extent possible. The orchestration strategy for all considered traffic is as follows:

- **TSN traffic:** the combinations of source and destination queues and LP that meet the latency requirements are computed, following K-Shortest Path (K-SP) using as metric the latency of the path and First Fit (FF), enforcing wavelength continuity, strategies for the routing and wavelength selection of the LP, respectively. Then, the combination that allows to use lower priority queues is selected, restricting it to also meet the data rate

requirements of the traffic and considering the queue occupancy with respect its capacity, that is, how many traffic flows are being supported over it. To do so, the harmonic mean of the service time of the queues is computed, selecting the combination with a higher value. Thus, TSN traffic is accommodated to resources that best fit its latency requirements, reducing the difference between them and the latency budget of the assigned queues.

- **BE traffic:** at the TSN domains, BE traffic is assumed to be accommodated to a lowest priority queue with no stated latency requirements and of enough capacity to serve all incoming traffic flows. Then, the selection of LPs is done in the same K-SP FF fashion, but in this case, in reverse order, thus starting from the slowest LPs. This favors the use of low latency LPs by TSN traffic.

This strategy requires the gathering of cross-domain information and the coordination of all involved network segments. To this end, we also present an enabling Software Defined Networking (SDN)-based control and orchestration framework to support the provisioning of both TSN and BE E2E connectivity services. Fig. 1 (left) summarizes the overall assumed scenario. Client domains implement a TSN-enabled SDN controller (SDN-TSN), which is responsible for placing the data flows associated to a service to available queues according to the requested latency-requirements, following the model proposed in the IEEE 802.1Qcc [3] for the TSN configuration. An extended northbound interface (NBI) module allows for the communication with the orchestration plane, translating service requirements to TSN data plane ones. Specifically, the NBI implements an extended version of the Transport API (T-API) [5] that allows the orchestration layer, first, to collect the required information to compute the service allocation and, second, to configure the data plane to activate such service. With regards to data collection, the information related to the latency bounds that each TSN domain is able to provide is gathered by the orchestrator. This information is formatted as a set of priority queues with their occupancy and achievable service time (latency), as required by the previously explained orchestration strategy (see Fig. 1 (right)). For the configuration, the priority queue to be allocated for the service in that domain is received by the SDN controller through the T-API module, which processes and translates it into the configurations required by the Centralized User Configuration (CUC) module defined in the TSN standard. Then, the CUC contacts the Centralized Network Configuration (CNC), which computes the path in the TSN domain and assigns the data flows to the queues implementing the most suitable scheduler. Similarly, the optical transport network implements an SDN controller. The T-API-based NBI is extended for the SDN controller to send topological and resource information about the optical data plane, which contains the switching times for the nodes and the length of the links (see Fig. 1 (right)), along with the traditional parameters used for optical path computation such as wavelengths, etc. This allows for the latency-aware LP computation at the orchestration layer, which is then configured by the underlying SDN controller. The orchestrator implements the Service Manager (SM) module that receives the E2E service requirements. In addition, the Topology Abstraction module (TA) collects the data plane information of the different domains and builds a graph where each link and node contain the information needed to compute the resource assignment to be used by the requested service. The computation is done at the Service Mapping module that implements the described orchestration strategy.




Fig. 1. End-to-end latency-aware control and orchestration framework (left), and validation topologies (right).

3. Results and discussion

We analyze the performance of the proposed orchestration strategy using as benchmark the per-domain non-orchestrated approach to resource provisioning explained in previous section over an example scenario following the network topology depicted in Fig. 1 (right-top), with link distances in km. We assume that the optical network has 40 WDM channels per optical link, with each one having a maximum capacity of 40 Gb/s, allowing for traffic grooming

between same source/destination pairs. The propagation delay of the fiber links is assumed to be 5 μ s/km, while the switching time of intermediate optical nodes is assumed to be 1 ms. Each TSN domain can be source or destination of both BE and TSN traffic. In regards of TSN traffic, each domain is modelled as being able to support four priorities of traffic, from most to least priority (4,3,2,1), with a maximum service time of 0.5, 1, 2 and 5 ms, respectively, and a capacity of 15 traffic flows per priority. The table in Fig. 2 (left) summarizes the traffic profile of the connectivity services, based on the requirements presented in [6], with BE representing an 80% of the connection requests and the remaining 20% equally distributed across all TSN service types. Connectivity services' arrivals follow a Poisson process, with exponentially distributed mean inter-arrival and holding times (IAT and HT). We put the focus on the performance experienced by TSN traffic. Fig. 2 (middle) depicts the blocking probability (BP) of TSN traffic connectivity requests due to latency unfulfillments, that is, the E2E resource assignment (queues and LP) is unable to satisfy the latency stated by the service type, as a function of the mean HT, with a mean IAT equal to one time unit and 10^8 arrivals per data point. The experienced total BP for the overall traffic is lower than 10% in the worst case.




Fig. 2. Summary of considered traffic and service types (left), TSN BP due to latency requirements unfulfillment (middle), and average queue assignment per priority (right).

It can be appreciated how an orchestrated approach is able to substantially reduce the experienced BP of TSN traffic, up to 4 orders of magnitude. This is because resources are assigned following an informed approach, adjusting the achieved E2E latency to the requirements of the connections. Thanks to this, low latency LPs are used to carry TSN traffic with low latency requirements or distant end-points, while higher latency LPs are allocated for the BE one or less demanding TSN traffic. This allows for a relaxation of the requirements of the TSN domain, which favors the assignment of low priority queues that require less stringent packet scheduling, while still meeting the E2E latency requirements. Such outcome can be observed in Fig. 2 (right), which depicts the average percentage of queue assignment per priority in the TSN domains. The joint orchestration of queues and LP assignment enables the reduction of high priority queue (priorities 4 and 3) assignment by around a 30-100%. Given the obtained results, we explored larger network scenarios, employing the DT topology depicted in Fig. 1 (right-bottom). The obtained results are also depicted in the plots of Fig. 2. The same conclusions hold here: the joint orchestration enables the reduction of BP due to latency causes by more than 4 orders of magnitude, reducing the burden of high priority queues by a 28-80%. An interesting conclusion can also be extracted: the joint orchestration of TSN and optical network domains enables the placement of IIoT control functions at further cloud sites, without loss of performance for the associated traffic flows. This is thanks for sure to the low latency performance of optical networks but an optimized assignment of TSN queues enabled by the joint resource orchestration plays a key role.

4. Conclusions and acknowledgments

We have demonstrated that by jointly orchestrating the TSN domains and the optical network interconnecting them, reductions in BP due to latency unfulfillments of up to four orders of magnitude are achieved, even in geographically large network scenarios. The main benefit is the relaxation of the resource assignment at TSN domains, enabling the use of lower priority queues for traffic flows, reducing the use of high priority queues up to a 98%. *This publication is supported by the projects TRAINER-B (PID2020-118011GB-C22) funded by MCIN/AEI/10.13039/501100011033 and TIMING (TSI-063000-2021-145) funded by MCIN and NextGenerationEU/PRTR.*

References

- [1] Y. Seol et al., "Timely Survey of Time-Sensitive Networking: Past and Future Directions", IEEE Access, **9**, pp. 142506-142527, Oct. 2021.
- [2] IETF, "RFC 8655 - Deterministic Networking Architecture", Oct. 2019.
- [3] B. Hadj et al., "SDN-based configuration solution for IEEE 802.1 time sensitive networking (TSN)", ACM SIGBED Review, **16** (1), pp. 27-32, Feb. 2019.
- [4] K. Christodoulopoulos et al., "Enabling the Scalability of Industrial Networks by Independent Scheduling Domains", OFC 2020, paper Th2A.24.
- [5] ONF Transport API (T-API), [online]: <https://wiki.opennetworking.org/display/OTCC/TAPI>
- [6] S. Bigo et al., "Dynamic Deterministic Digital Infrastructure for Time-Sensitive Applications in Factory Floors", IEEE JSTQE, **27** (6), pp. 1-14, Nov.-Dec. 2021.