On the selection of KPI-aware E2E paths for deterministic services in 6G networks

Salvatore Spadaro
Signal Theory and Communications
Universitat Politècnica de Catalunya
Barcelona, Spain
salvatore.spadaro@upc.edu

Albert Pagès
Signal Theory and Communications
Universitat Politècnica de Catalunya
Barcelona, Spain
albert.pages-cruz@upc.edu

Fernando Agraz
Signal Theory and Communications
Universitat Politècnica de Catalunya
Barcelona, Spain
fernando.agraz@upc.edu

Marta Blanco Caamaño
Telefónica Innovación Digital
Madrid, Spain
marta.blancocaamano@telefonica.com

Luis Miguel Contreras
Telefónica Innovación Digital
Madrid, Spain
luismiguel.contrerasmurillo@telefonica.com

***Abstract*—In the 6G ecosystem, the provisioning of deterministic services with KPI guarantees over multi-technological network domains requires the proper selection of the end-to-end (E2E) paths. Since each technological domain has its own capabilities and resources, in an E2E perspective, the proper selection of the domains to achieve the service KPIs is fundamental. The paper discusses the challenges of the E2E path/domain selection, provides an overall supporting control plane architecture and presents a proof of concept.**

Keywords—Path Computation, KPI, Deterministic services, 6G networks

I. Introduction

The emergence of innovative business models and use cases (e.g., Industry 4.0 and extended reality, XR), along with evolving operational paradigms such as edge and cloud computing, is driving the transformation of telecommunications infrastructure. These developments shape the expectations for network capabilities and performance. One critical requirement identified for future networks, such as 6G, is network determinism [1], which must be integrated into their design to meet the stringent demands of the mentioned use cases. Network determinism refers to Key Performance Indicators (KPIs), such as bandwidth, reliability, and end-to-end (E2E) latency, that must remain within very tight variance limits and be consistently guaranteed, regardless of network conditions. Among these KPIs, E2E latency and jitter are especially crucial for time-sensitive applications like industrial automation.

However, achieving and sustaining deterministic temporal KPIs throughout a service's lifetime is highly challenging. To ensure timely data delivery between endpoints, specialized hardware and data plane technologies must be developed and integrated. These technologies, combined with suitable communication protocols (for tasks like data forwarding, congestion control, and scheduling), enable determinism within data flows. Nonetheless, the data plane alone is not sufficient to enforce these KPIs.

Starting with the service requirements, it is necessary to map them to specific characteristics of the corresponding traffic flows, and configure the underlying hardware accordingly. Moreover, due to the tight KPI constraints, dynamic reconfiguration of network resources may be needed to adapt to current network conditions and maintain performance within acceptable limits. This necessity underscores the importance of a robust control layer. Software Defined Networking (SDN) has emerged as a promising control framework for managing deterministic networks [2]. Among the various SDN controller functions, Path Computation (PC) plays a pivotal role in establishing deterministic services with strict temporal KPIs. Determining the most suitable route for data to travel from source to destination, and configuring the associated network equipment, is essential in this context.

Hence, in deterministic networks, path selection is critical to meet target E2E latency and jitter values, which are influenced by factors such as network topology, data plane technology, and resource utilization (e.g., queue states at electronic switches). As a result, path computation must account for all these variables [3]. The challenge is further amplified in multi-domain scenarios, where each domain may have its own technologies and operational semantics. Thus, PC must not only compute routes within individual domains but also make intelligent decisions about domain selection to ensure E2E service performance (e.g., [4]).

In light of the above, in this paper we propose a control and management architecture and a set of mechanisms that allow for a KPI-aware path computation and selection to support E2E deterministic services across a multi-domain and multi-technology network infrastructure. The paper puts a special focus on addressing the main challenges that come from the explored scenario, mainly related to exposure and abstraction of topology and resources over multiple administrative domains. To this end, the rest of the paper is structured as follows: Section II describes the network scenario and the overall control and management architecture that is proposed. Section III highlights the challenges posed by the reference scenario and elaborates on the mechanism that are proposed to tackle them. Section IV provides a proof of concept to experimentally validate the proposed solution, and section V concludes the paper.

II. Reference Scenario

Fig. 1 depicts the scenario under consideration. A multi-domain multi-technology data plane is assumed, where each domain implements a specific data transport technology aligned with the current trends of the telecommunication network infrastructures that support 6G applications and services, such as 3GPP, Ethernet, IP or optical transport networks. Moreover, each technological domain implements its own control solution to configure and manage connectivity services at device level, that is, data forwarding and switching.

In order to provide E2E services across the multi-domain multi-technology scenario, a two-layer control and management infrastructure is assumed. While the first layer is composed of a set of per-domain management functions (MFs), which reside on top of each technological domain, the second layer implements the set of MFs responsible for computing and orchestrating the E2E service provisioning.

In this paper, we put the focus on the MFs that are aimed to implement the functionalities to enable E2E KPI-aware path computation and the challenges they have to face. In particular, the Topology Exposure (TE) MFs, which are responsible for collecting and managing the information related to the resources, capabilities and topology of the infrastructure, and the Path Computation (PC) MFs, which take care of the KPI-aware path calculation at the technological domains, and perform the domain selection at E2E level, are discussed in the following sections.

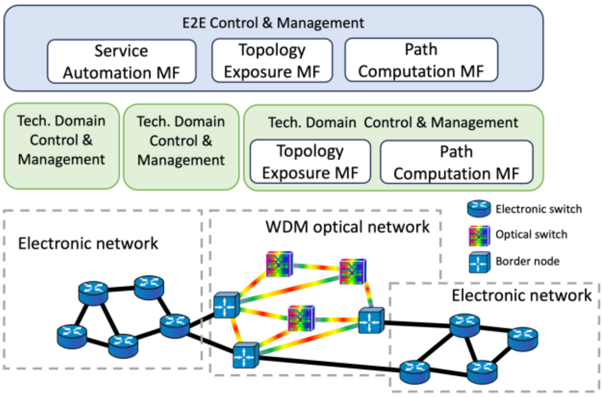

Fig. 1. Control and management architecture and network scenario.

III. Challenges on E2E KPI-aware Path Computation

The main input required by the PC entities running in the technological domains' control layer is related to the topology and the resources available on the underlying network infrastructure, as well as the capabilities offered by such infrastructure. In this regard, the TE MFs residing at the technological domain level are the responsible entities to collect the information about the resources of the data plane and their topological lay-out. Furthermore, information about additional functional capabilities, for example related to determinism such as specific forwarding protocols, priorities, replication, etc., can be collected. Lastly, status information having a direct impact on the targeted KPIs, such as current occupancy of network resources may be also collected. All this data collection is enabled by the Southbound interface (SBI) of the TE MFs. The gathered information can then be used by the PC MF of the domain to compute specific routes that accurately fulfil the connectivity requirements posed by the requested service. However, at the E2E level, topological information is not as rich as in the underlying technological domain mainly due to administrative and confidentiality limitations. Moreover, the heterogenicity of the multi-domain network infrastructure complicates the control and management tasks. For these reasons, some mechanism needs to be applied to allow for E2E path computation with accurate connectivity requirements fulfilment (i.e., the KPIs, such as E2E latency and jitter). In particular, an appropriated topology abstraction mechanism, which is able to provide homogeneous summarized (but precise enough) information related to the network infrastructures composing the different domains is crucial. To solve this, the E2E layer can rely on the TE MF of each technological domain, which can implement the functionalities for constructing an abstract view of its underlying data plane. Such view can be then consumed by the E2E TE MF to compose the E2E topology that will be used at this level to compute, provision and maintain multi-domain services with guaranteed deterministic performance.

From the provisioning perspective, upon the reception of an E2E service request, which conveys a set of connectivity requirements (e.g., latency and throughput) as well as a set of metrics stating the behavior of these requirements [5], the E2E PC MF has to compute an E2E path that fulfils such requirements. In a multi-domain scenario, while requirements like throughput or bandwidth stay the same across the E2E path and, thus, need to be equally fulfilled by all the traversed domains; time related requirements such as latency present a different behavior. Therefore, the latency requirement posed by an E2E service can be split in different latency budgets across the multiple domains traversed by the path supporting such service. For example, let us assume an E2E service that requests for a throughput of 100 Mb/s and 5 ms of E2E latency and has to traverse three network domains. While the throughput has to be equally satisfied by all three domains, the E2E latency can be achieved with different intra-domain paths combinations. For instance, if the first domain provides a path with 1ms latency, and the second and third ones provide 2ms latency paths, the E2E path will fulfil the overall latency requirement; but other path combinations would also be valid (e.g., 2ms latency paths in the first and second domains, and 1ms latency path in the third one). This flexibility, however, also poses some challenges. First of all, an inaccurate assignment of the latency budgets to the domains may impact on the number of services that the overall infrastructure will be able to support. Moreover, an E2E PC, which is aware of the KPIs (such as latency) offered by the underlying multi-domain network infrastructure is needed. The aforementioned topology abstraction plays a main role in this regard.

A. *Topology abstraction in support of KPI-aware Path Computation*

In the context of deterministic networks, a valid candidate to act as the gluing technology to provide a common view and control of the heterogenous multi-domain multi-technology scenario is the DetNet architecture [6] that is being proposed and standardized by the Internet Engineering Task Force (IETF). DetNet aims to provide a control plane to support deterministic connectivity services over multiple transport technologies, (e.g., IP, MPLS, OTN). However, in its current form, the DetNet model does not suffice to implement a topological abstraction able to provide accurate information that allows for an E2E path computation that considers the KPIs offered by the different domains. Aiming at standards compliance, we propose an IETF-based abstraction model that contains information related to the performance of the connectivity offered by the technological domains.

The proposed abstraction strategy follows the network slicing concept [7]. More specifically, the technological domains are abstracted as a set of border nodes interconnected by different slices that fulfil a set of performance requirements (i.e., KPIs). The border nodes are presented to the E2E layer as Provisioning Edges (PEs) and the exposed ports are translated into Service Demarcation Points (SDPs) [6]. The intra-domain connectivity (i.e., the network slices provisioned between borders) is abstracted in the form of Connectivity Constructs (CCs). The KPIs offered by each intra-domain connectivity element (i.e., CC) are modelled in the form of

Service Level Objectives (SLOs). Finally, the inter-domain connectivity existing between the PEs of different domains is modelled following the Attachment Circuit (AC, [8]) concept.

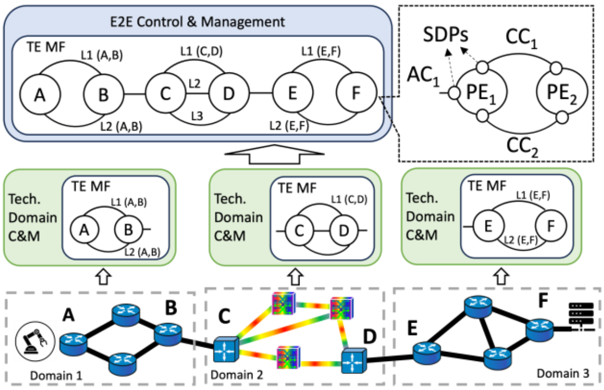

Fig. 2. Topology abstraction exposure mechanism and model.

Fig. 2 illustrates an example of the proposed abstraction computation and exposure mechanism. The figure depicts three domains implementing different data plane technologies. Each domain is governed by a single control and management (C&M) entity (green boxes) that implements the corresponding intra-domain MFs. From the exposure perspective, the TE MFs residing in the technological domain C&M collect the information from their underlying network infrastructures and compose the abstract view that will be afterwards exposed to the E2E. To construct the abstract topology, the TE MF of each technological domain requests to the corresponding PC MF to compute a set of paths between the border nodes of the domain also accounting for their expected performance (i.e., KPIs). Such paths are then modelled as CCs, and the border nodes are modelled as PEs. Moreover, the links that interconnect the domain to the neighboring ones are added to the model in the form of ACs. Finally, the end-points (i.e., ports) of the CCs and the ACs are modelled as SDPs. Hence, taking as an example Domain 1, the TE MF collects the complete topology, resources and capabilities from the infrastructure and requests to the PC MF to compute a set of paths between the border nodes A and B. The PC MF computes two paths with their associated achievable KPIs. For sake of simplicity, we just depict latency (L). With this information, the TE MF computes the abstract topology of the domain that will be then exposed to the E2E. The E2E TE MF, in turn, collects the abstract topologies of each domain it is in charge of and composes the complete abstract view of the scenario. Such E2E view will be used by the PC MF at the E2E level upon the reception of an E2E path computation request.

B. *E2E KPI-aware Path Computation mechanism*

The aforementioned topology abstraction enables the computation of KPI-aware paths between end-points in a multi-domain multi-technology infrastructure. However, this requires of dedicated logics and mechanisms that exploit the exposed information. In this paper, we propose a set of mechanisms that allow for an E2E path computation, which is aware of the KPIs offered by the underlying multi-domain network infrastructure and, thus, is able to provide quality assured connectivity to the requested deterministic services. In light of this, in this section we describe the E2E path computation operation that takes advantage of the previously described topology abstraction mechanism.

As described in section II, the proposed control and management architecture is split into two layers. From the path computation perspective, the PC MFs that operate at the technological domain level are responsible for computing the paths between the border nodes of the domain and use detailed information of the network infrastructure to estimate the KPIs associated to such paths. For example, the PC MFs use the length and the type (e.g., copper or fiber) of the links to compute the propagation delay, and the occupancy of the buffers of the electrical switches. These kinds of values are collected by the TE MFs and used to estimate KPIs such as latency. In addition, as described in the previous sub-section, the paths computed between the border nodes are used to provide an abstract topology of the domain, which is enriched with the mentioned KPIs.

At the E2E level, the PC MF is responsible for computing the domain sequence that is more likely to fulfil the connectivity requirements of the requested service. To do this, the PC MF computes an E2E path over the complete abstract topology that has been previously collected by the E2E TE MF, which contains information about the KPIs offered by the multiple technological domains. Thanks to this information, once the abstract path has been computed, the E2E PC MF is able not only to decide the domain sequence of the E2E path, but also to assign a specific budget for the KPIs that, like latency, do not need to be equally fulfilled across all the domains. Hence, the E2E PC MF contacts the PC MFs of the involved technological domains (i.e., the ones that are part of the domain sequence) to request the actual paths between the border nodes of each one (per-domain paths). It is worth noting here that the per-domain path requests convey the budgets of the KPIs previously computed and assigned to each domain. Finally, once the E2E PC MF receives the per-domain paths, it checks that the composed KPIs satisfy the connectivity requirements of the E2E deterministic service and, if so, sends the E2E path for further provisioning.

IV. Proof of Concept

In this section, we provide the experimental assessment of the described mechanisms. To this end, we consider the three-domain scenario depicted in Fig. 2, with emulated data plane infrastructure. For the control and management, a virtualized environment was deployed where three Virtual Machines (VMs) executed each technological domain MFs, and a fourth one was dedicated to the E2E MFs.

With such a set-up, Fig. 3 illustrates the experimental validation of the topology exposure and abstraction process. In particular, the red dashed box labeled as (a) showcases the requests from the TE MFs of each domain to the corresponding infrastructure to collect the data plane resources and capabilities information to compose the topology of the domain. In addition, the TE MF computes the abstract topology that will be further exposed to the E2E C&M. To do this, each TE MF requests to the technological domain PC MF a set of paths between the border nodes of the infrastructure as well as their expected KPIs. The TE MF uses this information to build the topological abstraction model that the E2E control and management layer, the TE MF residing at such level contacts the TE MF of each of the underlying technological domains to request the mentioned abstraction model. Once the TE MF has collected the models from all the underlying technological domains, it composes the E2E abstract topology that will be used by the E2E PC MF to compute the E2E connectivity that will support the requested services. As said, given that the abstraction model includes information about the KPIs associated to the local interconnection of the domains, an E2E KPI-aware domain and path selection can be realized. The red dashed box labeled as (b) in Fig. 3 showcases the mentioned E2E topology exposure process. Finally, Fig. 3 (c) provides an extract of the information model that is collected by the E2E TE MF from the TE MF of the technological domain. In particular, the figure illustrates the information related to one of the CCs (i.e., paths) computed by the PC MF between the border nodes (C and D) of Domain 2. For sake of simplicity, the latency is used as the only KPI in this experiment. Such latency (delay) is modelled as an SLO associated to the CC. Furthermore, the figure depicts the ACs (i.e., inter-domain connectivity) that Domain 2 will be sent to the E2E. In the same line, during the boot-up of has to interconnect with Domains 1 and 3, respectively.

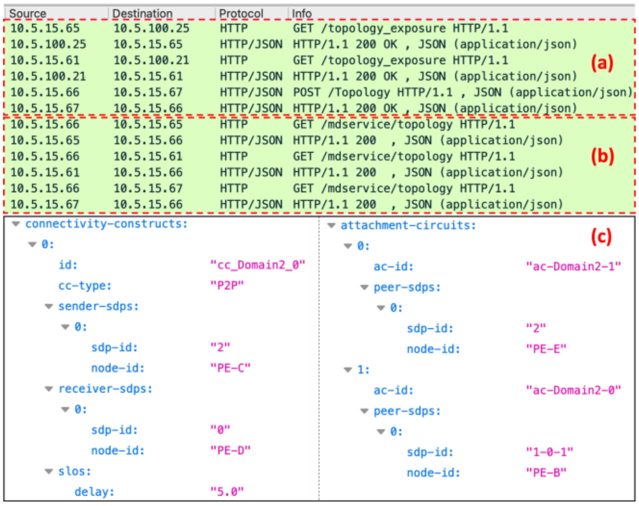

Fig. 3. Experimental validation of the TE MFs operation

Fig. 4 showcases the KPI-aware path computation process associated to the provisioning of an E2E deterministic service request, starting with the initial connectivity service request. As highlighted in the step (1) of the figure, such request conveys the requirements (i.e., KPIs) that the connectivity to be computed needs to fulfil for the service to run appropriately. Steps (2), (3) and (4) of the figure highlight the per-domain path computation process. As described in the previous section, the E2E PC MF uses the abstract topology composed by the E2E TE MF to realize the domain selection taking into account the KPIs offered at each technological domain. Once such domain selection has been done, the E2E PC MF requests to each domain's PC MF a path between the border nodes whose CCs provide the KPIs that fit best to the requirements posed by the service request. Once all the per-domain computations have been received, the E2E PC MF realizes a final E2E KPI check to assure the performance of the E2E path (step 5). If the check is successful, the provisioning of the connectivity will be finalized.

V. Conclusions

The provisioning of E2E deterministic services in the multi-domain scenario requires a careful path computation

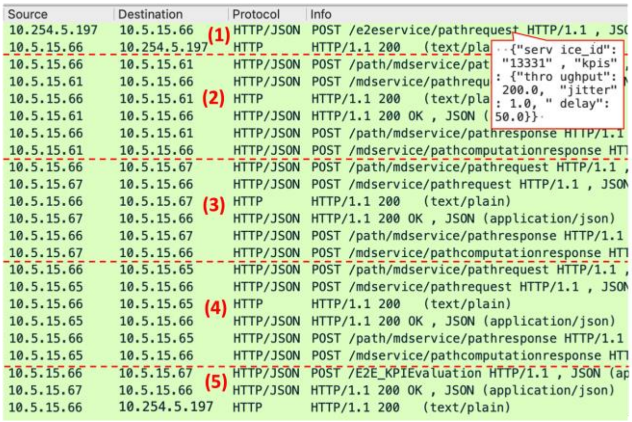

Fig. 4. E2E KPI-aware domain/path selection operation.

that allows for the fulfilment of the requested KPIs. In this paper, we reviewed the challenges posed to the KPI-aware path computation in such scenario, being the main ones related to topology exposure and coordinated KPI split between domains. Along the description of challenges, several mechanisms have been proposed to address them, with a special highlight on the control and management entities necessary to support and implement the associated operations. An experimental proof-of-concept has been presented to showcase the operation of the proposed solution, paving the way towards automated and KPI-aware connectivity provisioning in heterogeneous multi-domain multi-technology network infrastructures.

Acknowledgment

This publication is supported by the projects TRAINER-B (PID2020-118011GB-C22) funded by MCIN/AEI/10.13039/ 501100011033 and by the European Commission through Grant No. 101095890 (Predict-6G project).

References

[1] G. P. Sharma, D. Patel, J. Sachs, M. De Andrade, J. Farkas, J. Harmatos, B. Varga, H. P. Bernhard, R. Muzaffar, M. Ahmed, F. Dürr, D. Bruckner, E. Montes de Oca, D. Houatra, H. Zhang, and J. Gross, "Toward Deterministic Communications in 6G Networks: State of the Art, Open Challenges and the Way Forward," IEEE Access, vol. 11, pp. 106898–106923, September 2023.

[2] Y. Ryoo, T. Cheung, and H. Kim, "Implementation and field trial of SDN-based control plane for Deterministic Networking (DetNet)," Proc ICTC, Jeju Island, South Korea, October 2022.

[3] S. Spadaro, A. Pages, E. Guasch, and F. Agraz, "On the path computation for E2E deterministic services in future 6G networks," in 24th ICTON 2024 - International Conference on Transparent Optical Networks: July 14th-18th, 2024, Bari, Italy, 2024, pp. 1–4.

[4] A.Pagès, E. Guasch, F. Agraz and S. Spadaro, "KPI-aware service provisioning for remote industrial control systems management", Elsevier Computer Communications, vol. 238, June 2025.

[5] IETF Draft, "Path Computation Based on Precision Availability Metrics", March 2025.

[6] IETF RFC 8655, "Deterministic Networking Architecture," October 2019.

[7] IETF RFC 9543, "A Framework for Network Slices in Networks Built from IETF Technologies," March 2024.

[8] Internet Draft, "YANG Data Models for Bearers and Attachment Circuits-as-a-Service (ACaaS)," draft-ietf-opsawg-teas-attachment-circuit-20 (work in progress), January 2025.