MUDED: Integrating Networks with Applications through Multi-Domain Exposure and Discovery Mechanisms

Danny Alex Lachos Perez
University of Campinas (UNICAMP)
Sao Paulo, Brazil
dlachosp@dca.fee.unicamp.br

Christian Esteve Rothenberg
(Supervisor)
University of Campinas (UNICAMP)
Sao Paulo, Brazil
chesteve@dca.fee.unicamp.br

***Abstract*—The evolving Internet application landscape is envisioned to adopt technologies such as SDN, NFV, and MEC to provide softwarized services, requiring resource orchestration across multiple networks managed by different technological and administrative domains. In such multi-domain settings, the collaboration between networks and applications provides opportunities to both applications to improve their performances and network service providers to increase their business offerings. Although many systems are proposed to support such collaborations, they are point or incremental solutions. In this work, we propose the exploration of a more integrated architecture with huge possibilities taking a network-application integration (NAI) approach. Specifically, we explore the NAI possibilities in two concrete aspects: application-aware networking and network-aware applications. We review recent progress in these two aspects, and identify the key barriers in systematically realizing such a deep integration. To address these barriers, we propose a generic multi-domain NAI exposure and discovery framework, called MUDED. Through different systematic analysis and demonstrated prototypes, the MUDED's key components showcase: the maturing of the IETF ALTO protocol on the road to becoming a generic NAI possibilities discovery and exposure mechanism; and the optimized network view generation to simplify the network service placement and management.**

I. INTRODUCTION

Many novel multi-domain applications rely on a robust interaction with their deployment infrastructure. Data-intensive science applications (*e.g.*, the large hadron collider, telescopes, and light sources), for instance, rely on networks as one of the key components of their infrastructure for local and global interconnection of laboratories and data centres [26]. Another unexpected but evident example is the current COVID-19 pandemic, with many institutional applications taking advantage of the network infrastructure to share data quickly and support collaborative efforts from multiple communities and disciplines such as medicine, health, and disaster mitigation [39]. Flexible inter-domain routing [4, 29] and multi-domain service function chaining [2, 34] are also emerging applications that construct complex data flows between users in the network.

Clearly, the collaboration between networks and applications increases the quality and hence the business offering of the former, and the performance of the latter [21, 27]. Consequently, different systems and mechanisms have been proposed to support such collaboration. However, they are point or incremental solutions with various limitations. For example, network providers and applications have considered different nash equilibrium solutions [20, 36, 38]. However, such solutions are largely application/network-oblivious, making the interaction between them inefficient. Other solutions follow a "best-effort" [22, 30] or "blackbox-request" [1, 33] approach. Despite such proposals provide better network-application collaboration, they are typically implemented under the scenario of a single network, operated by a single commercial entity with their own bespoke applications (*e.g.*, Hadoop MapReduce, Google Search, Facebook).

At the crossroads, multi-domain applications are likely to take advantage of technologies such as Software Defined Networking (SDN), Network Function Virtualization (NFV), Multi-access Edge Computing (MEC) to provide softwarized network services, decomposed into Virtual Network Functions (VNFs), usually instantiated in distributed resources which are available across multiple domains with different technology and/or administration [35] (See Fig. 1). Such network services (or simply services) demand stringent requirements such as on-demand application deployment, information on deployment topology and network capabilities, scalability, and security. As a result, much of the existing work does not directly apply, and there is a need for two-way network-application interaction.

In this work, we propose to explore a more integrated and coherent architecture that takes a deep network-application integration (NAI) approach. Specifically, we explore the possibilities of NAI in two concrete aspects: application-aware networking and network-aware applications. The first one allows applications to specify diverse requirements for the network infrastructure. The second one allows networks to expose underlying network information available to applications.

Despite the huge possibilities of NAI, systematically realizing it is non-trivial. The key challenge is the lacking of generic and standard mechanisms for exposure and discovery of NAI possibilities. Existing solutions either fail to provide accurate resource sharing information [21, 31], or expose the complete information of the network [5, 6], raising scalability and security concerns.

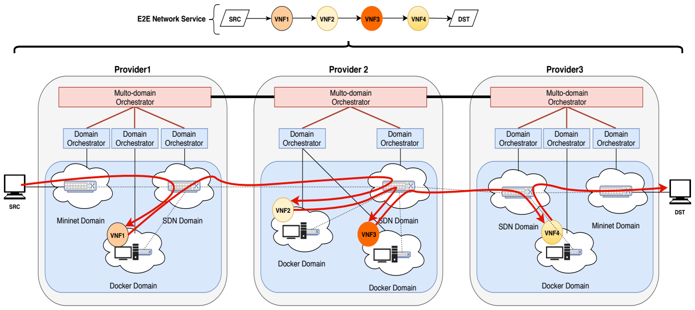

Fig. 1: An example for E2E Network Services across multiple domains (Multi-technology and Multi-administration)

To fill this gap, we propose MUDED, a **MU**lti-**D**omain generic NAI possibilities **E**xposure and **D**iscovery framework. The MUDED system resorts in two main mechanisms. First, a standardized discovery and exposure mechanism based on the IETF Application-layer Traffic Optimization (ALTO) protocol [31] to allow two-way network-application interaction. Second, a novel method for constructing abstracted network information to address the scalability challenges when processing large amounts of data in multi-vendor, heterogeneous technology environments.

The hypothesis of this work is that: *application and network integration is a key component for multi-domain settings, which for widely use should consider generic and standard mechanisms satisfying the ever so important features of NAI possibilities exposure and discovery, along with addressing the scalability and performance concerns*. The validation of such assertive takes place by pursuing three objectives: **(O1)** bridging the gap between networks and applications by showing the huge possibilities of NAI in settings traversing multiple domains; **(O2)** demonstrate the maturing of ALTO to be used as a standard mechanism for exposing and discovering NAI possibilities; and **(O3)** provide a novel abstraction mechanism to generate service-optimized network inventory views among different domains to deal with scalability requirements.

In order to achieve such objectives, a series of contributions were achieved by this work, among them: the systematic analysis of NAI possibilities, including the prototype implementation of a multi-domain NAI framework; several standardization proposals at the ALTO working group to extend the ALTO design to support multi-domain settings; and a novel abstraction mechanism to reduce the time to place services while optimizing the management of resources.

The rest of the paper is organized as follows. Section II presents the main research problem and motivation. Section III describes how the research proposal is novel to existing work. We point out main goals and contributions in Section IV before concluding the paper in Section V.

II. MOTIVATION & RESEARCH QUESTIONS

A fundamental problem in the Internet architecture is that applications and networks are designed with different objectives [21]. On the one hand, network applications (*i.e.*, network resource consumers) aim to optimize the application's utility (*e.g.*, maximize throughput, robustness, etc.). On the other hand, network providers (*i.e.*, Internet service providers or ISPs) aim to optimize the network's utility (*e.g.*, minimize inter-domain cost, minimize MLU - Maximum Link Utilization, etc.). With this inefficient interaction, getting a feedback between networks and applications is extremely limited.

In order to fill this gap, network providers and applications have considered different nash equilibrium solution approach (See Fig. 2a). ISPs, for example, attempt to improve the application issues through an infrastructure upgrade, usage-based charging model, rate limiting, or termination of services [38]. Meanwhile, applications attempt to improve the network efficiency having flexibility in shaping communications patterns as well as having flexibility to adapt to network topologies and conditions [20, 36]. However, the poor network-application cooperation in this approach does not allow to improve both network and application utility.

Other existing solutions adopt either a "best-effort" [22, 30] or "blackbox-request" [1, 33] approach (See Fig. 2b). In the first one, solutions allow applications to submit complete network requirements, and the network computes and enforces the optimal resource allocation for applications. In the second one, applications submit the amount of network resources needed, and the network returns success or failure based on the resource availability. Such solutions are quite good but nevertheless typically implemented for bespoke applications, such as big-data (Spark/Hadoop MapReduce), search (Google), social-network (Facebook). Besides, either "best-effort" or "blackbox-request" approach has limitations on the privacy/scalability, or has inefficiency in finding the optimal resource allocation for applications, respectively.

For softwarized multi-domain environments, instead of operating in the development of somewhat isolated and incremental solutions, we propose the exploration of a more integrated and coherent architecture. Such an architecture takes a multi-domain NAI approach (See Fig. 2c) to accommodate a variety of applications and domain sciences and a variety of administrations and virtualization technologies. The design of a single coherence system considers both application-aware information to the network and network awareness in applications. Such research perspective is addressed in subsection IV-A, "Deep Network & Application Integration", which poses the following research question:

Research Question #1: What are the possibilities of a much stronger network & application collaboration than the current mainstream networking?

In this subsection, we conduct a systematic review of the large variety of possibilities in designing and implementing NAI by application-aware networking and network-aware applications. We also present the design and evaluation of MUDED, a generic multi-domain NAI possibilities exposure and discovery framework.

Instead of building from scratch, we aim to design a multi-domain NAI framework that leverages the maturing of protocols and interfaces such as the IETF ALTO protocol.

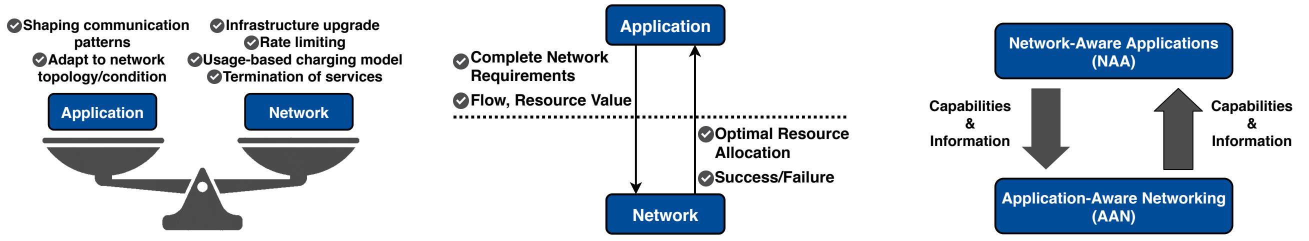

(a) (b) (c)

Fig. 2: Different approaches for the interaction of networks and applications: (a) nash equilibrium point, (b) best-effort/black-box approach, and (c) network-application integration (NAI) approach.

ALTO already introduces basic mechanisms (*e.g.*, modularity, dependency) and abstractions (*e.g.*, map services) for applications to improve their performance [10]. However, the current ALTO base protocol is not designed for a multi-domain setting of exposing network information. Towards such resolution, throughout subsection IV-B of this work, “Multi-domain Information Exposure & Discovery using ALTO”, we pose the following research question:

Research Question #2: How to expose and discover multi-domain NAI possibilities using ALTO?

In this subsection, we identify what network information the multi-domain applications need and the benefit of using it. We then discuss the current ALTO design issues for gathering such multi-domain information along with basic mechanisms to be considered to allow ALTO to expose network information across multiple domains.

On the other hand, as the deployment size and heterogeneity complexity of softwarized multi-domain networks increase, one of the foremost challenges for management systems is how to handle the scale of network service provisioning effectively. A common system engineering principle to deal with scalability requirements is to introduce proper abstraction mechanisms that reduce the discovery time of network resources while simplifying their management. As subsection IV-C showcases, aiming “ANI: Abstracted Network Inventory”, we pursue:

Research Question #3: How to effectively handle the scale and complexity of multi-domain environments to create proper abstract network views?

In this subsection, we propose the Abstracted Network Inventory (ANI) component to generate service-optimized network views over the same network inventory. ANI implements a novel abstraction method where network service requirements are used as an input to generate an optimized abstract network inventory representation, called Logical Network Inventory (LNI).

III. STATE OF THE ART

A. Network-Application Integration (NAI)

There are huge possibilities of NAI, however, there are still in place three major barriers to systematically realize such a deep integration.

(B1) Network information exposure. Applications are lacking of visibility of available and shared network resources (*e.g.*, bandwidth of shared resources for a set of flows), resulting in poor performance. Existing network resource exposure mechanisms, including graph-based [5, 6] and one-big-switch-based [21, 31] representations, either expose all sensitive information, or fail to capture the resource sharing between virtual flow requests.

(B2) Network information discovery. The lacking of a generic, flexible mechanism for applications to specify and discover the network information they need for NAI, from the network. Existing solutions (*e.g.*, [19, 21, 31]) provide application interfaces to discover E2E cost information of different packet spaces. However, this information is derived from the network’s fixed resource allocation (*e.g.*, fixed-route assignment) to the corresponding packet spaces, and applications are not provided the flexibility to discover additional network resources (*e.g.*, on-demand routing) that can satisfy their needs (*e.g.*, waypoint routing).

(B3) Optimized network view. Resource-filtering mechanisms may effectively reduce the redundancy in the network view. However, the number of available configurations would increase exponentially with both the network size and the number of services, which can be computationally expensive and time-consuming. To deal with scalability requirements, several solutions have been proposed for network view creation and abstraction [7, 24, 25, 28, 32]. However, none of the existing approaches take into account service requirements to generate an optimized network view representation.

B. ALTO in Multi-domain Scenarios

Emerging multi-domain applications (*e.g.*, collaborative data sciences [3, 37], flexible inter-domain routing [4, 29], E2E network services [2, 34]) require resource orchestration across multiple networks managed by different administrative domains. Such cross-domain applications can benefit substantially from network information exposure and discovery to make application-layer resource optimization and improve their performance

ALTO is a mature protocol that already provides a generic framework to discover and expose network information for applications to take optimized actions based on network information [10]. In particular, ALTO introduces *generic mechanisms* such as: (i) information resource directory (IRD), (ii) information consistency (tag, dependency, multi-info resources), and (iii) information update model (*e.g.*, incremental update with server-sent events). ALTO also introduces *abstractions* exposing network information to the applications: (i) network and cost maps, (ii) the path vector abstraction, and (iii) capability maps (*e.g.*, CDNI and unified property Map).

However, the current ALTO base protocol is not designed for a multi-domain setting of exposing and discovering network information. This work identifies the benefits of using multi-domain information in applications traversing multiple domains. It also elaborates key design requirements of ALTO for exposing multi-domain information along with a set of mechanisms to design a multi-domain ALTO framework, thus removing barriers **(B1)** and **(B2)**.

C. *Network Inventory Abstractions*

To discover qualified configurations for applications, the network requires up-to-date and accurate resource availability representation, in the form of a network inventory that can become very large in multi-domain environments. As such, different solutions are proposed to create an optimized representation of the network inventory [7, 24, 25, 28, 32].

For example, UNIFY [7] provides an abstraction of type big-switch and big-software that includes computing and network resources. Solution in [28] uses linear inequalities to represent network resources availability in terms of bandwidth. However, all these solutions do not consider service requirements to generate a network representation. Likewise, [24, 25] provide different abstraction models (big-switch, virtual link with single weights, and virtual link with multiple weights) of optical transport networks, focusing especially on centralized radio access networks (C-RANs). However, both solutions also do not consider the use of service requirements as an input to generate an abstract network representation. Work in [32] supports grouping by joining entities within hypernodes or hyperedges. Our proposal has two main differences (i) filtering instead of aggregation as abstraction mechanism and (ii) service requirements as an additional input.

In order to remove barrier **(B3)**, we propose the ANI component that implements a novel method (i) receiving services requirements from a catalog, (ii) receiving a network representation from a network inventory, and (iii) processing those inputs to generate an optimized network view (*i.e.*, LNI).

IV. GOALS AND CONTRIBUTIONS

In general, the main goal of this research involves the systematic study of NAI in the context of multi-domain environments that, for widely use it needs to consider standard mechanisms for exposing and discovering NAI possibilities, including mechanisms for providing an optimized network view to address the scalability concerns. We attain such a goal by the approaches and the contributions described in the next subsections, each defining an activity in the development of this work. Such activities are mapped as illustrated in Figure 3.

A. *Deep Network & Application Integration*

We review huge possibilities in designing and implementing NAI. Suggesting answers for the research question #1.

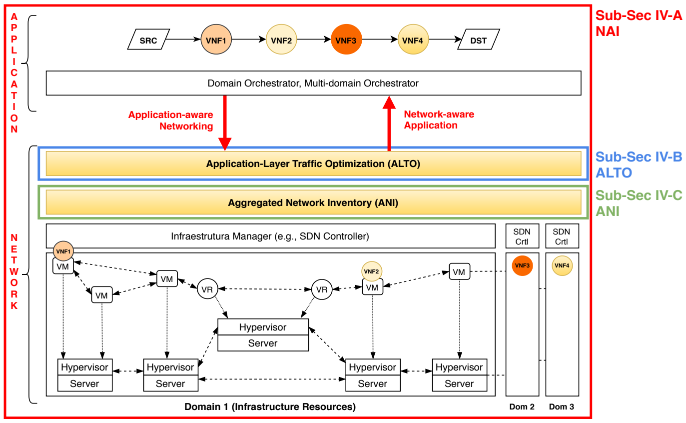

Fig. 3: Scope of the thesis towards NAI thought multi-domain exposure and discovery mechanisms.

On the one hand, applications have varying needs for network latency, bandwidth, packet loss, etc. However, such applications' requirements are often unknown to the network due to applications and networks are decoupled. Thus, one concrete aspect of NAI is adding application knowledge to the network so that applications can express finer granularity requirements. There are substantial possibilities in designing and implementing NAI by application-aware networking. For example, the network infrastructure can provide better support for applications introducing different capabilities such as transport differentiation and in-network storage/compute.

On the other hand, applications running over networks face challenges due to the lack of network state and information. Applications can benefit from network information exposure to make them more flexible in terms of rate adaptation, transmission time, server/path selection, among others. Therefore, the other side of designing and implementing NAI is network-aware applications, and there are many possibilities as well. For instance, applications have possibilities to conduct transport selection capabilities based on network state (*e.g.*, packet loss), performance metrics (*e.g.*, throughput, max reservable Bandwidth), capability information (*e.g.*, delivery/acquisition protocol), and locality (*e.g.*, servers location and paths).

In addition, we also propose the design of MUDED, a multi-domain generic framework for NAI possibilities exposure and discovery. MUDED resorts in two mechanisms. First, an ALTO-based discovery and exposure mechanism to allow two-way network and application interaction. One, in the "top-down" direction, is related to the discovery expression by the application of the required properties needed to be supported by the network. The other one, in the "bottom-up" direction, is related to the network information exposure that can be processed and consumed by the application. Second, a novel abstraction method directed to the ANI component to generate service-optimized network views or LNIs. Altogether, the following contributions have been achieved:

1) A systematic review of the large variety of possibilities in designing and implementing NAI by application-aware

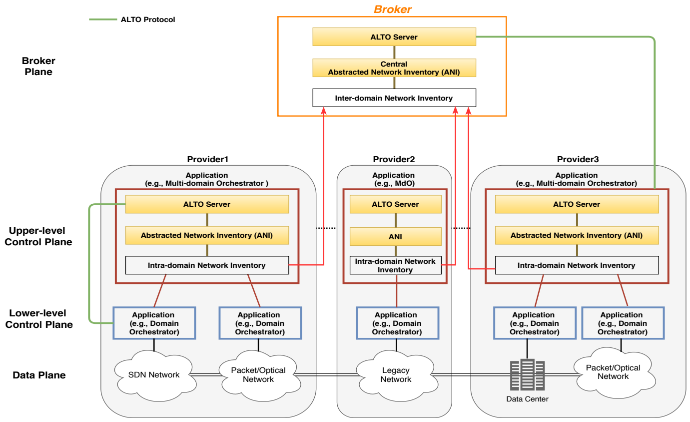

Fig. 4: MUDED: Integrating Networks with Applications through Multi-Domain Exposure and Discovery Mechanisms.

Current Key Issues	Envisioned mechanisms
Server-to-Client ALTO communication	**Server-to-Server ALTO communication**
Domain connectivity discovery	**Multi-domain connectivity discovery**
ALTO server discovery	**Multi-domain ALTO server discovery**
Single-domain composition	**Unified Resource Representation**
Simple resource query language	**Generic/Flexible query language**
Scalability	**Computation complexity optimization**
Security & Privacy	**Security/Privacy preserving**

TABLE I: ALTO in Multi-domain: Issues & solutions [14].

networking and network-aware applications.

2) Key challenges for systematically realizing deep NAI.
3) We present the design of MUDED (See Fig. 4), a NAI possibilities exposure and discovery system for network service placement in multi-domain scenarios.
4) A prototype implementation of MUDED[1], following the 5GEx project architectural design[2]

From the systematic analysis to detailed design and experimental validation, the topics above have been addressed in the contributions listed in [8, 11, 13, 18], including a best paper award [9] and a contribution in a European funded initiate[3].

B. Multi-domain Info. Exposure & Discovery using ALTO

The current ALTO base protocol is not designed for a multi-domain setting of discovering and exposing network information. Via an analytical investigation of the ALTO framework, we identify several key design issues and summarize on-going efforts and potential solutions that suggest an answer to the investigated research question #2. Table I shows the relationship between the key ALTO design issues and their corresponding mechanisms to consider in a multi-domain ALTO framework.

The main contributions can be summarized as follows:

1) We identify what network information the emerging multi-domain applications need and the benefit of using it.
2) We give a systematic review of the ALTO design issues for multi-domains settings of exposing network information.
3) We summarize envisioned solutions and on-going efforts to design a multi-domain ALTO framework.

These contributions lead to the related publications that are listed in [10, 14, 15, 16, 17, 23], which include several standardization proposals at the IETF ALTO working group.

C. ANI: Abstracted Network Inventory

Suggesting answers for the proposed research question #3, The ANI component proactively constructs multiple network views over the same network infrastructure, called LNIs. The ANI receives two inputs (See Fig. 5): (i) services requirements from a catalog, and (ii) a network inventory representation. Both inputs guide the right level of abstraction to generate a LNI. Each LNI is optimized to a service in terms of its requirements such, as CPU, memory, latency, etc. Every new service in a catalog triggers the creation of another LNI that will be part of the optimized network inventory.

The concrete contributions are as follows:

1) We propose the ANI as a novel component that allows the creation of service-optimized network inventory views in distributed environments. To the best of our knowledge, ANI is the first approach that uses service requirements as an input to generate an abstract network inventory.
2) We formally define a network model along with the development of three algorithms to generate an LNI given the capacity-related resources and requirements of a network inventory and network service, respectively.
3) We evaluate the proposed algorithms through extensive experiments using random and real-world topologies[4]. Results show significant benefits of using an LNI to simplify the service management and placement: (i) the relationship between compute nodes and links (*i.e.*, density) in an LNI is reduced between 1.8-2.7x compared to a full network inventory topology (See Fig. 6a); and (ii) up to 50% of time can be saved for service placement after abstracting around 20% of the compute nodes (See Fig. 6b).

The contributions above, including a work as per patent application, are listed in [12, 40].

V. FINAL REMARKS

In evolving networking scenarios (*e.g.*, 5G) where applications provide softwarized services requiring resource orchestration across multiple network domains, application and network integration is fundamental. The networking community (academy/industry) has pointed at the importance of revisiting this topic regularly to identify its possibilities and

---

[1] https://github.com/intrig-unicamp/alto-based-broker-assisted-mdo
[2] https://doi.org/10.1002/ett.3085
[3] https://tinyurl.com/y3w348bj
[4] https://github.com/intrig-unicamp/ani

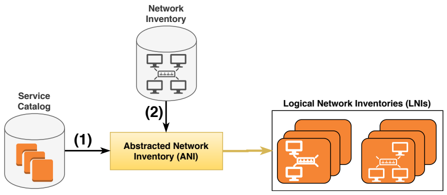

Fig. 5: ANI: Logical Network Inventory (LNI) generation

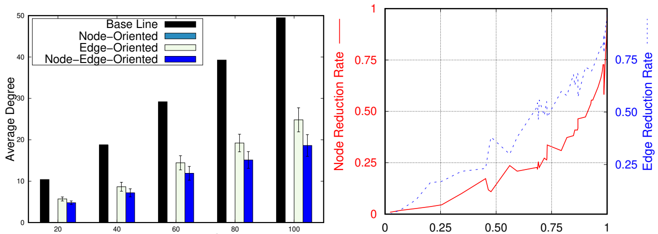

(a) Degree

(b) Time saving rate

Fig. 6: ANI impact on service Mngmt (a) & Provisioning (b).

challenges (SIGCOMM'20: Workshop on NAI[5]). This work contributes with a systematic analysis of huge possibilities in designing and implementing NAI by application-aware networking and network-aware applications. We design and evaluate MUDED, a multi-domain NAI possibilities discovery and exposure framework to address the key barriers of systematically realizing NAI. By achieving objectives **O1-O3**, MUDED demonstrates, on the one hand, how barriers **B1** and **B2** can be removed taking advantage of maturing NAI protocols such as ALTO (and extensions); and while the proposed solutions are not mature enough to have an immediate impact on industry standards, they are in an excellent position to be adopted in relevant bodies like IETF. On the other hand, towards addressing the scalability concerns and remove barrier **(B3)**, the ANI component is implemented. ANI emerged as a per patent application which, when filed, it was incorporated into the MUDED system, demonstrating the two-way transfer of benefits between academy and industry.

REFERENCES

[1] M. Campanellaet al. Bandwidth on demand services for european research and education networks. In *2006 IEEE First International Workshop on Bandwidth on Demand*, pages 65–72. IEEE, 2006.
[2] NGMN Alliance. 5g white paper. *Next generation mobile networks, white paper*, pages 1–125, 2015.
[3] The CMS Collaboration. The CMS experiment at the CERN LHC. *Journal of Instrumentation*, 3(08), 2008.
[4] A. Gupta et al. Sdx: A software defined internet exchange. *ACM SIGCOMM Computer Communication Review*, 2015.
[5] A. Verma et al. Large-scale cluster management at google with borg. In *Proceedings of X European Conference on Computer Systems*, 2015.
[6] B. Hindman et al. Mesos: A platform for fine-grained resource sharing in the data center. In *NSDI*, 2011.
[7] B. Sonkoly et al. Multi-domain service orchestration over networks and clouds: a unified approach. In *ACM SIGCOMM Computer Communication Review*. ACM, 2015.
[8] D. Lachos et al. Broker-assisted multi-domain network service orchestration. In *IEEE Wireless Comms. and Networking Conf. (WCNC)*, 2018.
[9] D. Lachos et al. Multi-domain orchestration leveraging the application-layer traffic optimization protocol. In *V Workshop Pré-IETF*. SBC, 2018.
[10] D. Lachos et al. Supporting multi-domain use cases with alto. In *Proceedings of the Applied Networking Research Workshop*, 2019.
[11] D. Lachos et al. ALTO-based Broker-assisted Multi-domain Orchestration. Draft draft-lachosrothenberg-alto-brokermdo-04, IETF, July 2020.
[12] D. Lachos et al. Ani: Abstracted network inventory for streamlined service placement in distributed clouds. In *2020 6th IEEE Conference on Network Softwarization (NetSoft)*, pages 319–325. IEEE, 2020.
[13] D. Lachos et al. Multi-domain E2E Network Services. Internet-Draft draft-lachosrothenberg-alto-md-e2e-ns-02, IETF, July 2020.
[14] D. Lachos et al. Multi-domain information exposure using alto: The good, the bad and the solution. In *Proceedings of the Applied Networking Research Workshop*, pages 52–54, 2020.
[15] D. Lachos et al. Multi-domain Service Function Chaining with ALTO. Internet-Draft draft-lachos-sfc-multi-domain-alto-01, IETF, July 2020.
[16] D. Lachos et al. Multi-domainn Information Exposure using ALTO. Internet-Draft draft-lachos-alto-md-info-exposure-00, IETF, July 2020.
[17] D. Lachos et al. Supporting Multi-domain Use Cases with ALTO. Draft draft-lachos-alto-multi-domain-use-cases-01, IETF, July 2020.
[18] D. Lachos et al. Towards deep network & application integration: Possibilities, challenges, and research directions. In *Proceedings of the Workshop on Network Application Integration/CoDesign*, 2020.
[19] E. Pujol et al. Steering hyper-giants' traffic at scale. In *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies*, ACM CoNEXT '19, New York, NY, USA, 2019.
[20] H. Madhyastha et al. iplane: An information plane for distributed services. In *Proceedings of the 7th symposium on Operating systems design and implementation*, 2006.
[21] H. Xie et al. P4p: Provider portal for applications. *ACM SIGCOMM Computer Communication Review*, 38(4), 2008.
[22] J. Lee et al. Application-driven bandwidth guarantees in datacenters. In *Proceedings of the 2014 ACM conference on SIGCOMM*, 2014.
[23] L. Contretas et al. Use of ALTO for Determining Service Edge. Internet-Draft draft-contreras-alto-service-edge-01, IETF, July 2020.
[24] M. Fiorani et al. Transport abstraction models for an sdn-controlled centralized ran. *IEEE Communications Letters*, 19(8):1406–1409, 2015.
[25] M. Fiorani et al. Abstraction models for optical 5g transport networks. *IEEE/OSA Journal of Optical Communications and Networking*, 2016.
[26] Marian Babik et al. HEPiX Network Functions Virtualisation Working Group Report, April 2020.
[27] P. Schmidt et al. Socket intents: Leveraging application awareness for multi-access connectivity. In *Proceedings of IX ACM Conf. on Emerging Networking Experiments & Technologies (CoNEXT)*, USA, 2013.
[28] Q. Xiang et al. Fine-grained, multi-domain network resource abstraction as a fundamental primitive to enable high-performance, collaborative data sciences. In *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 58–70, 2018.
[29] Q. Xiang et al. Sfp: Toward interdomain routing for sdn networks. In *ACM SIGCOMM 2018 Conference on Posters & Demos*, 2018.
[30] R. Soulé et al. Merlin: A language for provisioning network resources. In *Proceedings of the 10th ACM International on Conference on emerging Networking Experiments and Technologies*, pages 213–226, 2014.
[31] S. Kiesel et al. Application-Layer Traffic Optimization (ALTO) Protocol. RFC 7285, September 2014.
[32] T. Soenen et al. A model to select the right infrastructure abstraction for service function chaining. In *IEEE Conference on NFV-SDN*, 2016.
[33] X. Zheng et al. Cheetah: Circuit-switched high-speed end-to-end transport architecture testbed. *IEEE Communications Magazine*, 2005.
[34] ETSI. Zero touch network & service management. https://www.etsi.org/technologies/zero-touch-network-service-management, 2020.
[35] G. G. Sun, Y. Li, D. Liao, and V. Chang. Service function chain orchestration across multiple domains: A full mesh aggregation approach. *IEEE Transactions on Network and Service Management*, 15(3), 2018.
[36] A. Kuzmanovic and E. Knightly. Tcp-lp: low-priority service via end-point congestion control. *IEEE/ACM Transactions on Networking*, 2006.
[37] The Linac Coherent Light Source. https://lcls.slac.stanford.edu/.
[38] Michael Ott. Intelligent network based application recognition, November 1 2005. US Patent 6,961,770.
[39] Pacific Wave. R&E Networks in CA, OR, and WA Announce Collaborative Support for COVID-19 Wester States Pact, April 2020.
[40] M Santos, D Lachos, P Gomes, C Rothenberg, and A Vidal. Technique for simplifying management of a service in a cloud computing environment, 2019. Patent No. PCT/EP2019/058274, Filed April 02, 2019.

---

[5] https://conferences.sigcomm.org/sigcomm/2020/workshop-nai.html