Fiber-to-Application: Optical Slicing to Enhance Application Performance over a Metro Transport Network

Cen Wang[1], Xiao Xue[2], Noboru Yoshikane[1], Filippos Balasis[1], Hongxiang Guo[2], Takehiro Tsuritani[1]

1 KDDI Research Inc., Saitama, Japan. {ce-wang, yoshikane, fi-balasis, tsuri}@kddi-research.jp
2 Beijing University of Posts and Telecommunications, Beijing, China. {xuexiao, hxguo}@bupt.edu.cn

Abstract: We demonstrate the fiber-to-application transport slicing architecture and mechanism. The experiment shows ultrahigh throughput ($> 5Gbps$ per application) and significant acceleration for 100 applications in 4 categories.

1. Introduction

Optical transport with multiple resources (i.e., spectral and spatial) is the cornerstone of emerging network scenarios, such as 5G enhanced mobile broadband (eMBB), internet-of-things (IoT), vehicle-to-everything (V2X) and edge computing. To achieve optical resource slicing for specific scenarios, a data model and slicing procedure in the transport domain were proposed [1–3]. In addition, recent works [4–6] moved a step forward to demonstrate end-to-end optical slicing use cases.

However, the aforementioned scenarios are application-centric, meaning that each application has specific demands. Moreover, it has witnessed rapid increments of such demands (e.g., several Gbps for the distributed machine learning applications), in a short future, the traffic from a single application will occupy full of the resources in an optical channel. End-to-end slicing usually leads to application-agnostic effects due to a focus on the aggregated traffic level. In addition, TCP/IP-based and packet-switched applications struggle with throughput bottlenecks, even when allocated with sufficient bandwidth, which makes high-throughput applications, such as holographic video and federated computing [7] driven by 5G hard to fulfill. Thus, two key issues remain to be discussed: (1) a method of mapping per-application demand to multidimensional optical resources; and (2) a flattened solution for applications to directly utilize the high bit-rate and high-quality transmission of the transport domain.

In this paper, we extend end-to-end optical slicing to a fiber-to-application (FA) style. More specifically, we use the application group concept [8] to perform application-specific bandwidth allocation, and then map the group demand onto the transport resources. For one of the groups, we deploy FPGA-based FA network interface cards (FANICs) to let the application stream bypass the TCP/IP protocol and seamlessly move onto the optical resources. Considering this, we modify the conventional *application group slicing* (AGS) and propose *FA + application group slicing* (FAAGS). We demonstrate the proposed architecture and mechanism, then compare the performance of the proposed FAAGS with the AGS and *traffic-level slicing* (TLS). Experimental results verify that FAAGS outperforms AGS and TLS by achieving a throughput of $> 5Gbps$ per application and significantly accelerating the application completion times (ACTs).

2. Architecture

Fig. 1 displays the holistic slicing architecture. In the data plane, FANICs (Xilinx KC705) are deployed via PCIe on partial computing resources. After bypassing the OS kernel, the application data are streamed and sent to a $10Gbps$

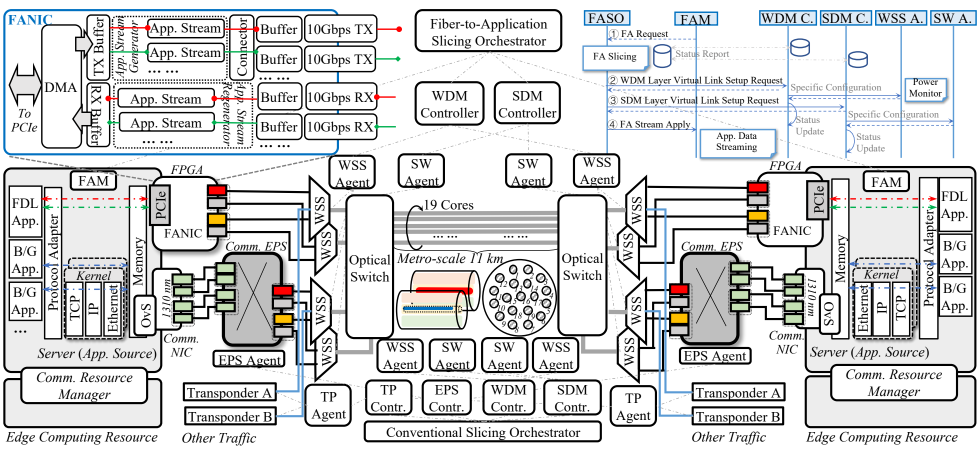

Fig. 1: Holistic slicing architecture.

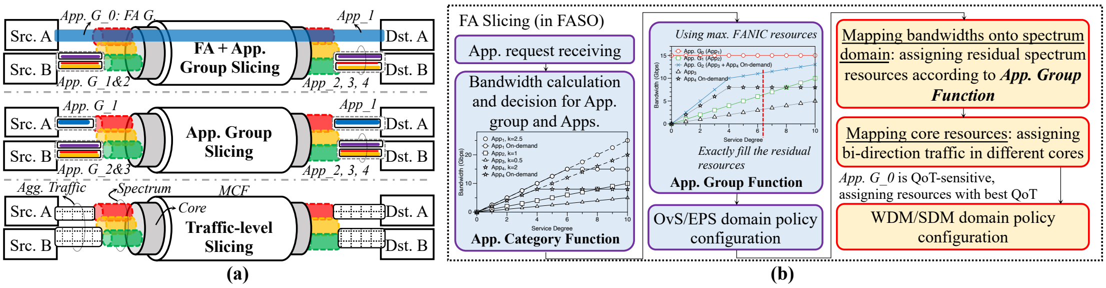

Fig. 2: Slicing mechanism.

interface. The interface is bound to a DWDM SFP+ for the direct reach of spectrum resources. Legacy NICs are also applied for common usage. OpenvSwitchs (OvSs) are linked to the NICs for application identification and per-application bandwidth allocation. The LR SFP+ on the NICs are connected to a packet switch (Dell S5224). The packet switches perform group-level bandwidth allocation. For simplicity, we plug DWDM SFP+ on a packet switch for an application to utilize spectrum resources. The transponders (ADVA), on the other hand, only generate signals with more spectrum to influence the slicing policy. All the DWDM SFP+ and transponders access WSSs, and the WSSs are connected to the optical switch (Polatis) to occupy different spatial resources (cores) of a 19-core MCF with $11km$.

In the control plane, we develop a fiber-to-application slicing orchestrator (FASO). FASO conducts a slicing mechanism (see details in next section) and sends policies to fiber-to-application managers (FAMs), OvSs, packet switches, WDM and SDM controllers. The latter two use *NETCONF* to control the corresponding devices. The FAM configures the application to choose a legacy NIC or a FANIC. Fig. 1 additionally highlights the procedure of an application to use FANIC and to obtain related optical resources.

3. Slicing Mechanism

Fig. 2 (a) shows the concepts of TLS, AGS and FAAGS. The TLS allocates optical resources according to the traffic amount between all source-destination pairs (usually IP-level). The AGS is inspired from B4 [8]. This reduces the burden of application identification on the core site. In the FAAGS, we define $App.G_0$ as the FA group, and high-throughput application is preferred to be divided in $App.G_0$. Fig. (b) details the slicing mechanism of FAAGS implemented in FASO.

The *App. Category Function* (ACF) and the *App. Group Function* (AGF) is predefined in the mechanism. ACF denotes the bandwidth value to be allocated for all of the existing applications belonging to an application category under a given service degree (dimensionless and depends on the residual bandwidth), where k is the slope. The ACF is flattened according to the real demands. Then, the applications are grouped. Applications in $App.G_0$ occupy maximum resources for FANICs by default, and since resources are not shared with other groups, the AGF of $App.G_0$ becomes flat. For other groups, the AGF is the summation of ACFs of all group members. The bandwidth for the groups is allocated according to the service degree. For example, in AGF of Fig. 2 (b), the red dashed line indicates the service degree when allocating, and the intersections between the red dashed line and the bandwidth curves are the allocated bandwidth values (the summation of which exactly equals the residual resources).

After obtaining the allocated bandwidth, the FASO starts to configure the OvS and packet switch. Subsequently, FASO maps the bandwidth into spectral resources group by group. Such mapping denotes generating control policies for packet switching (the outputs) and WSSs (adding/modifying connections). FASO then allocates cores. For convenience of management, we use the same spectrum for the bidirectional traffic of an application group; thus, the bidirectional traffic is assigned to different cores. More essentially, as the FANIC offloads a higher layer error correction and retransmission, it becomes much more sensitive to the QoT. Assigning optical resources for $App.G_0$, OSNR and a crosstalk (for the high crosstalk MCFs) needs to be considered. FASO finally generates the control policies for the transport domain.

4. Experiment

In the experiment, we generate 100 applications with 4 categories divided into 3 application groups, as shown in Table 1. We operate three slicing mechanisms in Fig. 3 (a) and compare the ACT and average throughput of the applications. Fig. 3 (a) shows the control message related to the FANIC setup highlighted in Fig. 1. Note that this is an initial establishment, and thus, the control delay is high. The subsequential applications in $App.G_0$ reuse these resources to reduce the delay.

Fig. 3 (b) shows the CDF of ACTs under three slicing mechanisms. This shows that FAAGS avoids ultralong ACTs due to the high throughput of FANIC. In addition, using FANIC for $App.G_0$ increases the service degree for the other groups, which becomes the reason FAAGS outperforms AGS. Under TLS, although the bandwidth is sufficient for aggregated traffic, the applications compete for such bandwidth, and any congestion reduces the throughput (TCP behavior) of all of these applications. As a result, the ACTs are extremely long.

Table 1: Applications in the experiment.

App.	Description	App. G.	No.	App.	Description	App. G.	No.
App_1	Federated machine leaning (FML)	$App.G_0$	10	App_3	Control msg. and IoT data	$App.G_2$	40
App_2	Data synchronization and backup	$App.G_1$	20	App_4	File download/upload	$App.G_2$	30

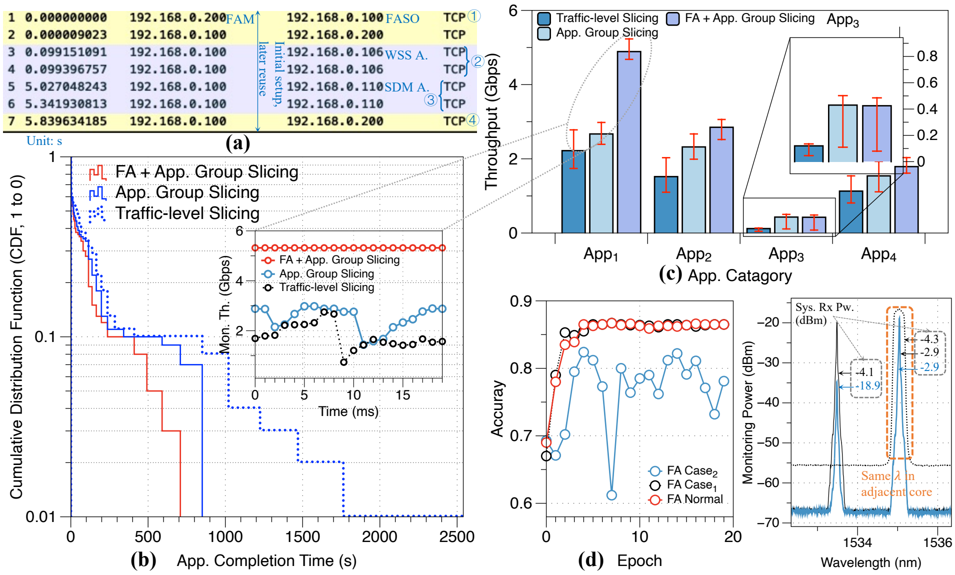

Fig. 3: Experimental results.

We analyze the average throughput of each application category in Fig. 3 (c). App_1 achieves the highest and most stable throughput when using FANIC. For those applications (except for App_3) without FANIC under FAAGS, the throughputs are higher than those of AGS. App_3 shows no improvement, which indicates that their bandwidth is already sufficient under AGS. The TLS shows the lowest throughput. The reason of the throughput performance is almost the same with that of ACT performance.

We also demonstrate the QoT influence for the FML (using FANICs), as shown in Fig. 3 (d). In $case_1$, we schedule the FANIC to use the same spectrum as the adjacent core. The accuracy under $case_1$ is almost equal to a normal case. That is to say, crosstalk does not impact the QoT required by FANIC. In contrast, as for $case_2$, we set an attenuation on the channel to emulate the unexpectedly low OSNR. The system reported RX power is lower than the receiver sensitivity. The accuracy under $case_2$ suffers misconvergence, which is resulted from the received error gradients.

Conclusion

In this paper, we contribute FANIC to bridge applications and optical resources by virtue of a high bit-rate and a good quality of optical transmission. We combine FANICs and legacy NICs, while proposing FAAGS to improve the performance of the diverse applications. Compared to traditional AGS and TLS, our architecture and mechanism show both a high throughput and a high acceleration for specific applications. These advantages meet the application-centric requirements driven by 5G, IoT, V2X and federated computing.

References

1. R. Casellas, A. Giorgetti, R. Morro, R. Martinez, R. Vilalta, and R. Munoz, "Enabling network slicing across a disaggregated optical transport network", in *2019 Optical Fiber Communications Conference and Exhibition (OFC)*, 2019, pp. 1–3.
2. R. Casellas, A. Giorgetti, R. Morro, R. Martinez, R. Vilalta, and R. Munoz, "Virtualization of disaggregated optical networks with open data models in support of network slicing", *IEEE/OSA Journal of Optical Communications and Networking*, vol. 12, no. 2, A144–A154, 2020.
3. R. Munoz, N. Yoshikane, R. Vilalta, J. M. Fabrega, L. Rodriguez, R. Casellas, D. M. S. Moreolo, R. Martinez, L. Nadal, D. Soma, Y. Wakayama, S. Beppu, S. Sumita, T. Tsuritani, and I. Morita, "Sdn control of sliceable multidimensional (spectral and spatial) transceivers with yang/netconf", *IEEE/OSA Journal of Optical Communications and Networking*, vol. 11, no. 2, A123–A133, 2019.
4. S. Nakamura, K. Shiomoto, H. Sugiyama, Y. Hirota, N. Yoshikane, K. Sugawara, M. Miyabe, T. Eguchi, S. Okamoto, M. Murakami, T. Hirayama, I. Sato, and T. Roux, "First demonstration of end-to-end network slicing with transport network coordination and edge cloud applications in 5g era", in *2019 24th OptoElectronics and Communications Conference (OECC) and 2019 International Conference on Photonics in Switching and Computing (PSC)*, 2019, pp. 1–3.
5. J. Li, X. Shen, L. Chen, J. Ou, L. Wosinska, and J. Chen, "Delay-aware bandwidth slicing for service migration in mobile backhaul networks", *IEEE/OSA Journal of Optical Communications and Networking*, vol. 11, no. 4, B1–B9, 2019.
6. J. Brenes, T. D. Lagkas, D. Klonidis, R. Munoz, S. Rommel, G. Landi, I. Tafur Monroy, E. Grivas, E. Pikasis, G. Bernini, J. M. Fabrega, and R. Vilalta, "Network slicing architecture for sdm and analog-radio-over-fiber-based 5g fronthaul networks", *IEEE/OSA Journal of Optical Communications and Networking*, vol. 12, no. 4, B33–B43, 2020.
7. Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications", *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 10, no. 2, pp. 1–19, 2019.
8. S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, et al., "B4: Experience with a globally-deployed software defined wan", *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4, pp. 3–14, 2013.