

# Reducing Latency in Virtual Machines

## Enabling Tactile Internet for Human Machine

### Co-working

Zuo Xiang\*, Frank Gabriel\*, Elena Urbano\*, Giang T. Nguyen\*, Martin Reisslein<sup>†</sup>, and Frank H.P. Fitzek\*

\*Deutsche Telekom Chair of Communication Networks  
Technische Universität Dresden, 01062 Dresden, Germany  
Email: {firstname}.{lastname}@tu-dresden.de

<sup>†</sup>School of Electrical, Computer, and Energy Engineering  
Arizona State University, Tempe, AZ 85287-5706, USA  
Email: reisslein@asu.edu

#### Abstract

Software Defined Networking (SDN) and Network Function Virtualization (NFV) processed in Multi-access Edge Computing (MEC) cloud systems have been proposed as critical paradigms for achieving the low latency requirements of the tactile Internet. While Virtual Network Functions (VNFs) allow greater flexibility compared to hardware based solutions, the VNF abstraction also introduces additional packet processing delays. In this paper, we investigate the practical feasibility of NFV with respect to the tactile Internet latency requirements. We develop, implement, and evaluate Chain bAsed Low latency VNF ImplemeNtation (CALVIN), a low-latency management framework for distributed Service Function Chains (SFCs). CALVIN classifies VNFs into elementary, basic, and advanced VNFs. CALVIN implements elementary and basic VNFs in the kernel space, while advanced VNFs are implemented in the user space. Throughout, CALVIN employs a distributed mapping with one VNF per Virtual Machine (VM) in a MEC system. Moreover, CALVIN avoids the metadata structure processing and batch processing of packets in the conventional Linux networking stack so as to achieve short per-packet latencies. Our rigorous measurements on off-the-shelf conventional networking and computing hardware demonstrate that CALVIN achieves round-trip times from a MEC ingress point via two elementary forwarding VNFs (one in kernel space and one in user space) and a MEC server to a MEC egress point on the order of 0.32 ms. Our measurements also indicate that MEC network coding and encryption are feasible for small 256 byte packets with an MEC latency budget of 0.35 ms; whereas, large 1400 byte packets can complete the network coding, but not the encryption within the 0.35 ms.

# Reducing Latency in Virtual Machines

## Enabling Tactile Internet for Human Machine Co-working

### I. INTRODUCTION AND MOTIVATION

Low latency communication is the central requirement for enabling the tactile Internet for human-machine co-working [1]–[6]. Both, humans and machines require latencies below one millisecond for a wide range of co-working scenarios. For instance, for humans operating in a virtual world and for interactions with robots and other machines, the latencies for visual, audio, or tactile multi-sensoric feedback should be below 15 ms, 3 ms, or 1 ms, respectively [7]. Every machine based on control loops also requires low latencies in order to work efficiently or to operate in a stable manner [8], [9]. As a concrete example, consider a classical inverted pendulum whose controller is placed in the cloud. Closing the control loop through a communication network will likely introduce some delays and packet losses. Fig. 1 shows the influence of the delay between the angle sensor and pendulum actuator (motor) on the pendulum stability. For long delays (50 ms in Fig. 1), the system becomes unstable, and the pendulum will never reach stability in the inverted position. For shorter delays (40 ms), the system takes some time to achieve stability. This time delay could imply lack of quality of service, and may affect other systems if the pendulum is part of a more complex environment with interconnected systems, or multiple pendulums coexisting in the same physical space.

The 5G communication standard for automation in vertical domains [10] defines the allowed latency requirement of one millisecond for an end-to-end communication. Based on this standard, we consider a typical allocation of the individual 5G communication network delay budget components, as illustrated in Figure 2. Figure 2 assumes that a total of 0.4 ms is consumed in the embedded systems and the wireless links on both ends. For instance, 0.1 ms may be allocated for the embedded sensor computing platform to evaluate the sensed information, 0.1 ms for (one-way) communication latency in uplink and downlink, and 0.1 ms for embedded computing at the actuator.

This leaves 0.6 ms for the wired domain. The wired domain has two main delay components. One delay component is the basic communication over fiber where we are bound to the speed of light ( $3.34 \mu\text{s}$  per kilometer) and the physical fiber characteristics. The second delay component is based on the communication nodes. In conventional “store and forward” communication networks, these communication nodes are routers or switches. But in upcoming future communication, there is a paradigm shift from “store and forward” to “compute and forward”. Now the communication nodes can process and manipulate the incoming data. The idea of “compute and forward” is realized by new technologies, such as Software Defined Networking (SDN) [11], Network Function Virtualization (NFV) [12]–[19], and Service Function Chaining (SFC) [20]–[34] standardized by the IETF/IRTF. These novel technologies enable the concept of Multi-access Edge Computing (MEC) [26], [35]–[40], which allows for local processing of data, which in turn will reduce the latencies on the pure communication path. Assuming a maximum

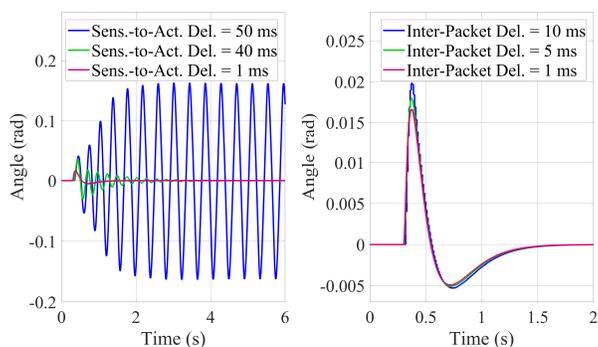


Fig. 1: Angle of an inverted pendulum trying to reach stability for different sensor-to-actuator delays (50 ms, 40 ms, and 1 ms) for 1 ms inter-packet delay as well as for 1 ms sensor-to-actuator delay for different inter-packet delays (10 ms, 5 ms, and 1 ms).

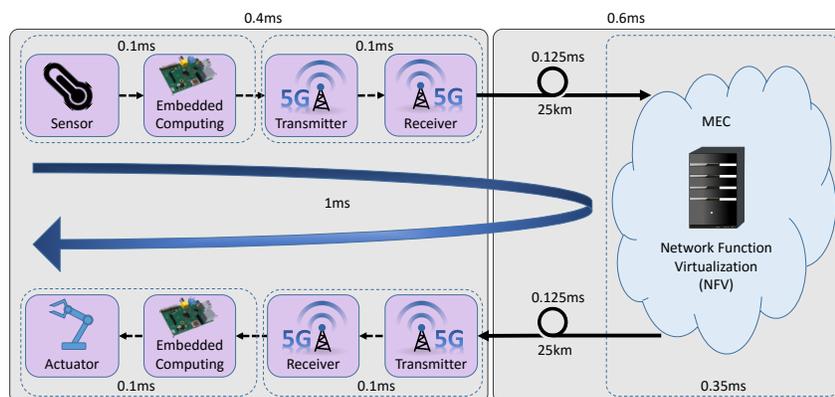


Fig. 2: Typical latency budget for sensor-to-actuator control loop that meets one millisecond round-trip latency requirement of 5G [10]: 0.4 ms are allocated for embedded sensor and actuator processing and wireless communication, leaving 0.6 ms for the wired link propagation as well as the virtualized communication environment (“compute and forward”) processing in the MEC.

62 distance of 25 km between the sensor/actuator and the MEC, the (round-trip) communication will require 0.25 ms  
 63 and will leave 0.35 ms for NFV processing in the MEC system.

64 However, achieving low latencies is a notoriously difficult problem in communication networks [41], [42]. While  
 65 delays that are proportional to the available transmission bitrate and data amounts can be addressed through scaling  
 66 up the transmission capacities and compression, processing delays with their various constant delay contributions  
 67 pose significant challenges [42]–[44]. Moreover, recent studies [45], [46] have demonstrated that NFV, which is  
 68 highly desirable for flexibility [47], imposes heavy data transfer and computation demands, incurring relatively long  
 69 latencies. Nowadays, virtual switches are quite fast [48]–[50]; however, virtual machines (VMs), specifically the  
 70 packet IO and processing operations inside the VM, are slow. The centralized approach proposed in [51] gives  
 71 more than 2 ms latency inside the virtual communication environment for a single VM running the elementary  
 72 forwarding function, which is clearly above the 0.35 ms delay budget.

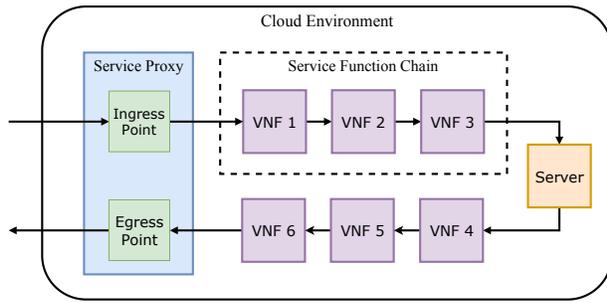


Fig. 3: Illustration of service loop in a cloud (MEC/virtualization) environment: Packets traverse an ingress service function chain (SFC) consisting of an ordered sequence of virtual network functions (VNFs) to reach the server for cloud processing and leave the cloud via the egress point.

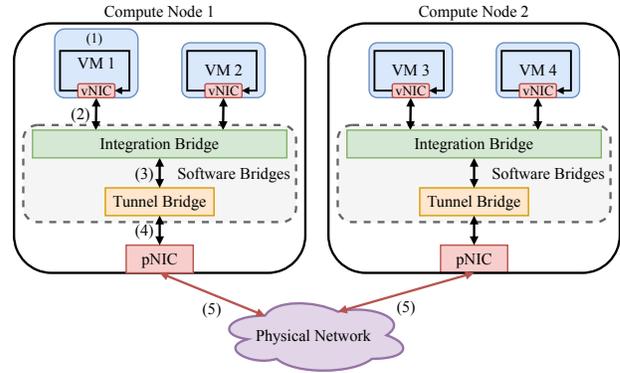


Fig. 4: Illustration of typical virtual network connection between two compute nodes in a cloud environment, whereby each compute node hosts multiple virtual machines (VMs). Each VM connects via a virtual network interface controller (vNIC) to the bridges and onwards via a physical NIC to the physical network.

73 In this empirical measurement study, we examine low-latency NFV in real general-purpose MEC systems built  
 74 with off-the-shelf hardware and software. We design, implement, and evaluate a low-latency service function  
 75 chain (SFC) management framework named Chain bAsed Low latency VNF ImplemeNtation (CALVIN). CALVIN  
 76 implements VNFs either in the kernel space or in the user space and distributes each VNF to its own VM. CALVIN  
 77 employs fast packet input/output (IO) mechanisms that avoid the metadata structures and the batch processing of data  
 78 packets in the conventional Linux networking stack. Our extensive measurements demonstrate that CALVIN achieves  
 79 MEC latencies on the order of 0.32 ms, which allows for additional processing of data, e.g., for network coding  
 80 and encryption. The proposed CALVIN approach makes it for the first time possible to process advanced network  
 81 functions, such as network coding and encryption, in a general-purpose virtualized MEC setting while meeting the  
 82 one millisecond delay target of the tactile Internet. Our CALVIN SFC management and VNF implementation codes  
 83 for OpenStack, which is the defacto industry standard for general-purpose cloud computing platforms, are openly  
 84 available at [52], [53].

## 85 II. BACKGROUND AND RELATED WORK

### 86 A. Background

87 The typical NFV service loop inside the MEC part of Fig. 2 is presented in Fig. 3. Each packet is received by the  
 88 cloud through the ingress point of a service proxy, processed by a chain of virtual network functions (VNFs), i.e.,  
 89 by an ingress SFC, transmitted to the requested server, processed by the server, and leaves the cloud via the egress  
 90 point of the service proxy (after optional egress SFC processing). The ingress and egress points are endpoints or  
 91 instances that are exposed to the external network (outside of the cloud network, which is commonly a dedicated  
 92 private network). For several reasons, including security considerations, not all internal components are exposed to  
 93 the external network. Therefore, ingress and egress points are required for clients in the external network to access

94 the cloud services. As illustrated in Fig. 3, an SFC consists of an ordered set of VNFs that operate typically as a  
95 pipeline.

96 The virtualized service loop of Fig. 3 is typically implemented on a cloud computing infrastructure platform. A  
97 cloud computing infrastructure platform provides flexible management control over underlying physical computing  
98 resources, such as servers, networking facilities, and storage. All components of the virtualized service loop are  
99 implemented and run in VMs that are orchestrated by the cloud computing platform.

100 Fig. 4 illustrates a typical cloud computing infrastructure scenario where multiple VMs are connected with a  
101 virtualized networking overlay. In order to enable multi-tenant networking with configurable networking resources  
102 and isolation, a virtualized networking overlay needs to be built on top of the physical networking infrastructure.  
103 For example, VM1 and VM3 can be allocated in the same broadcast domain in a virtual network (or named tenant  
104 network) even though they are hosted on different physical nodes that are connected by layer three routing entities.

105 In order to provide a virtual overlay network on top of the underlying heterogeneous physical network, two  
106 software bridges (or virtual switches) are used to connect the virtual network interfaces (vNICs) with the physical  
107 network interface (pNIC). In particular, the integration bridge connects all VMs running on the same physical node.  
108 These VMs can belong to different virtual tenant networks even if they are on the same physical node. The tunnel  
109 bridge is used to encapsulate and transfer tenant network data through a tunneling protocol, e.g., Generic Routing  
110 Encapsulation (GRE) or Virtual extensible LAN (VXLAN).

111 The networking components in Fig. 4 introduce different types of latency:

- 112 ● Between VM and integration bridge (1, 2): This latency component is the focus of this study and includes  
113 two main parts:
  - 114 ○ Transfer packets between the integration bridge and VM through vNIC (2): This latency depends on  
115 the vNIC technology and has two subparts:
    - 116 ■ Data transfer between virtual bridge data buffer and vNIC ring buffer: With the performance im-  
117 provements of virtual bridges, e.g., the DPDK fast path in the Open vSwitch (OVS-DPDK) [54], the  
118 latency of this subpart has been reduced to the order of microseconds [48].
    - 119 ■ Frame copying between the vNIC ring buffer and the VM memory: This subpart now becomes a  
120 bottleneck for low-latency data transfer in the virtual network overlay.
  - 121 ○ Process packets inside VM (1): This latency depends on the implementation of the VNF processing  
122 workflow. Optimizations from the implementation perspective should process the frame as fast as possible  
123 for the low-latency tactile Internet.
- 124 ● Between integration bridge and pNIC (3, 4): This latency component depends on the employed virtual bridges  
125 as well as the physical resource management and orchestration (MANO) platform [55]. Cloud platforms  
126 commonly employ OVS-DPDK and OpenStack, which we also employ in our testbed.
- 127 ● Between pNICs (5): This latency component depends on the physical network technologies.

128 To the best of our knowledge, the reduction of latencies (1) and (2) is an open research question. Our proposed  
129 CALVIN approach significantly reduces these latencies so that the overall end-to-end latency of the service loop is  
130 within the 5G one millisecond latency requirement.

## 131 *B. Related Work*

132 Several recent studies based on mathematical analysis and simulations have considered SFC latencies, see  
133 e.g., [56]–[68]. These studies have mainly considered moderate to high workloads that result in substantial queuing  
134 delays for VNF processing. In contrast, we conduct an experimental study with empirical latency measurements  
135 on lightly loaded real networking and computing systems. Our goal is to rigorously empirically investigate the  
136 baseline latencies of real SFC implementations. Our measurement study complements the existing mathematical  
137 analysis and simulation studies and provides baseline latency values, which can serve as reference points for  
138 future analysis and simulation studies on low-latency VNF processing. We also develop and evaluate the CALVIN  
139 low-latency SFC framework which achieves significant latency reductions compared to existing frameworks. We  
140 also note that several recent studies have examined VNF placement strategies, e.g., [69]–[77], mainly from the  
141 perspective of mathematical modeling and optimization of the placement. A few recent implementation oriented  
142 studies have examined scheduling and flexibility aspects [25], [78]. Complementarily, our implementation oriented  
143 study examines the low-latency aspects of VNF placement.

144 Some recent studies have sought to reduce communications delays and increase communications reliability through  
145 specific communications strategies [6], such as short packet transmissions [79], [80]. Complementarily to these  
146 communications mechanism focused studies, we address the NFV latency in MEC systems for arbitrary packet  
147 sizes.

148 The remainder of this section surveys existing related SFC frameworks **both in kernel and user space** and  
149 distinguishes our approach from the existing frameworks. Our approach is not based on full kernel-bypassing;  
150 instead, we exploit the complementary strengths of both kernel and user space technologies to reduce the latency  
151 (while efficiently utilizing the CPU resources).

152 *1) Kernel Space:* Recent kernel space approaches have typically been based on the eXpress Datapath (XDP)  
153 framework (a new system in the Linux kernel) [81]–[84] which is built using custom extended Berkeley Packet  
154 Filter (eBPF) programs. Since the XDP framework is relative new (it became available with the Linux kernel  
155 versions after 4.1X), few studies have been conducted with XDP [84].

156 Miano et al. [85] have conducted quantitative characterizations of a range of eBPF aspects. Miano et al.  
157 documented several limitations when building complex network services with eBPF. Due to these limitations we  
158 do not implement all functions in the kernel space in CALVIN. Our own preliminary XDP evaluations (with Linux  
159 kernel 4.17.0-041700-generic), which are not included here due to space constraints, indicated that XDP achieves  
160 low latencies for elementary packet processing. However, due to the main XDP limitations (e.g., limited number of  
161 instructions, not Turing complete since loops are not allowed) we do not implement advanced packet processing in  
162 the kernel space. Nevertheless, XDP has several benefits over kernel-bypassing technologies. XDP does not require  
163 large pages (the smallest unit or block for memory management in the operating system (OS)), nor dedicated CPUs.  
164 Also, XDP does not replace the kernel TCP/IP stack; rather, XDP works in concert with the kernel TCP/IP stack  
165 along with all eBPF benefits. Moreover, XDP avoids the need to define a new security model by utilizing the Linux  
166 kernel security model.

167 In-kernel processing with XDP based on the eBPF with an NFV focus has recently been examined in the InKeV  
168 study [86]. InKeV employs XDP based on the eBPF, which can quickly execute simple functions; however, the  
169 implementation of advanced functions is very challenging. The InKeV latency measurements were performed on a

170 single machine and thus do not capture the latencies incurred when traversing a physical network to complete an SFC  
171 consisting of physically distributed VMs. In contrast, we consider physically distributed VMs in our evaluations.  
172 InKeV has been compared with the OpenStack neutron [87] networking infrastructure, which we also employ in our  
173 evaluations. Thus, InKeV could be a good competitor for the underlying virtual switches (OVSS) used by OpenStack  
174 (latency components (3)–(5) in Fig. 4). In contrast, we focus on the latency introduced by the VM, i.e., on the  
175 latency components (1) and (2) in Fig. 4.

176 For completeness, we note that relatively simple networking functions relating to security and virtual switching  
177 in the kernel space have recently been studied in [88]–[91].

178 2) *User Space*: User space frameworks, which are also referred to as kernel-bypassing frameworks, have been  
179 studied more frequently than kernel space frameworks [92]–[94]. However, most user space framework evaluations  
180 have focused on throughput and did not consider latency performance in detail. The performance of three widely  
181 known kernel-bypassing high-speed packet IO frameworks, including netmap, PF\_RING ZC, and Intel DPDK, has  
182 been measured in [94], revealing a general trade-off between throughput and latency. Recent studies have examined  
183 several aspects of user space frameworks, including CPU scheduling [95], flexible programmability [96], [97],  
184 and resilience [98], [99]. These recent studies are complementary to our CALVIN study and did not specifically  
185 focus on low latency. For instance, the NetStar study [96] has examined flexible asynchronous network function  
186 programming, which is mainly useful for network management traffic, e.g., flow table updates; whereas we focus on  
187 low-latency processing of data traffic. The HyperVDP study [97] has developed a hypervisor that flexibly offloads  
188 some CPU compute-intensive tasks to programmable network interface cards to reduce the resource usage (at the  
189 expense of slightly reduced throughput and increased latency).

190 3) *Combined Kernel and User Space*: Closer related to our approach are recent frameworks that combine kernel  
191 and user space techniques. General architectural principles for building hybrid kernel-user space VNFs have been  
192 explored in [100]. With the combined kernel and user space approaches, the VNF applications can be programmed  
193 with the common socket interfaces; thus, some legacy applications can be deployed without modification. The  
194 combined approaches can utilize the scheduler provided by the guest OS and can simultaneously utilize kernel and  
195 user space tools. Thus, the combined kernel and user space approaches allow for low complexity implementations.

196 Zhang et al. [51] have recently implemented network coding on typical VMs by employing DPDK [101] with the  
197 kernel network interface (KNI) [102]. We refer to this approach as the “centralized approach” as it strives to pack  
198 all VNFs into a single VM so as to avoid the latency introduced by transmissions between VMs, see Fig 5. In order  
199 to make multiple VNFs cooperate properly, the centralized approach employs both kernel and user space tools. As  
200 illustrated in Fig. 5, a packet needs to be transmitted between the kernel space and user space at least four times  
201 (red lines in Fig. 5 indicate slow path or bottleneck). Additional copies of packets between the kernel space and  
202 user space introduce non-negligible latencies, especially in a virtual guest OS. Moreover, the resources, especially  
203 virtual CPU (vCPU) resources, need to be shared between multiple processes both in kernel and user space; these  
204 context switches also incur latencies. Furthermore, the cache behavior of the CPU cannot be optimized because  
205 of this context switching; specifically, some processing related instructions cannot be pre-fetched and consistently  
206 stored in the CPU cache. The resulting relatively high latency is one main drawback of the centralized approach.  
207 The latency cannot be readily scaled down; for instance, utilizing two vCPU cores (one core to handle packet  
208 receptions and another core to handle packet transmissions) does not reduce the latency (as we have verified with

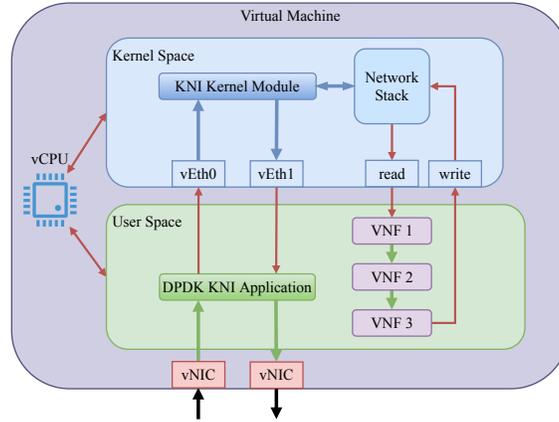


Fig. 5: Graphical illustration of operational principles of centralized combined kernel and user space approach described in [51]. The combined approach requires four or more packet transmissions between kernel space and user space, illustrated by the vertical red arrows crossing over between kernel space and user space.

our own measurements of the centralized approach). To overcome the drawbacks of the centralized approach, our main CALVIN strategy is to distribute VNFs over a chain of VMs that run the network functions *either completely in the kernel space or completely in the user space*.

We also note for completeness that a specific combined kernel and user space approach, referred to as “Tuna”, for a 5G wireless access point has been developed in [103]. The Tuna approach places management frames in the user space for virtualization, while placing control and data frames in the kernel space to reduce packet processing delays. In contrast, our approach is suitable for general VM processing and not tied to a particular application.

### III. PROPOSED APPROACH: CHAIN BASED LOW LATENCY VNF IMPLEMENTATION (CALVIN)

The proposed CALVIN approach aims to take advantage of both the high-performance in-kernel network data path and user space (kernel bypass) techniques. At the same time, CALVIN avoids the overhead of context switching of the centralized approach [51]. This section presents the overview, architectural design, and workflow of CALVIN.

#### A. Overview

The underlying idea of CALVIN is to assess the nature of a VNF in terms of its complexity when processing packets. Accordingly, CALVIN decides to implement each VNF in either the kernel space or the user space. Each implemented VNF is then encapsulated inside a separate VM. These two CALVIN design choices completely eliminate the context switching overhead of the vCPU of each VM to process packets in different spaces and to schedule different VNF programs, thus reducing the overall end-to-end service latency.

The main advantages of this CALVIN strategy are:

- Reduce cost of context switching inside VM: As quantified in [104], context switching can produce direct and indirect costs. Direct costs are incurred mainly for storing the state of a process. The cache sharing between multiple processes creates an indirect cost that can exceed one millisecond for high system workloads. Running a VNF in a single space mitigates the negative effects of both costs on latency.

- 231 • Avoid copying packet data structures between spaces: For ultra-low-latency VNF implementation, the cost of  
232 data copying and format conversion at any location of the data path should be considered. As illustrated in  
233 Fig. 5, data exchanges between spaces are critical bottlenecks for the VNF processing. Conducting all packet  
234 operations in a single space reduces data copying.
- 235 • Increase scalability and flexibility: In a centralized approach, such as [51], multiple VNFs run on the same  
236 VM, sharing the same resources and configurations. Because of the virtual resource contention, the latency  
237 performance is not scalable without resizing resources when new VNFs are added into the processing pipeline.  
238 In contrast, the CALVIN performance can be scaled due to the flexible structure of the dynamic SFC. Each  
239 new VNF can be assigned a specific VM with appropriately allocated resources for the current functional  
240 requirements.

## 241 B. VNF Classification

242 Based on these design imperatives, we take a first step towards the development of the CALVIN SFC management  
243 framework by classifying VNFs. According to the introduced VNF classification, a given VNF is either implemented  
244 in the kernel space or in the user space. The network functions are divided into three main categories in CALVIN:

245 1) *Elementary (Skeleton) Functions*: This type covers the fundamental functionalities for all VNFs: *i*) Retrieve  
246 packets from the ingress virtual interface. *ii*) Create data structures to store packets for operations. *iii*) Transmit  
247 processed packets through egress virtual egress interfaces. These functions can be implemented in both kernel and  
248 user spaces.

249 2) *Basic Functions*: The main characteristics of basic functions are: *i*) Operations are mainly performed on the  
250 header (or metadata); not on the packet payload. Headers have typically small sizes; thus, header operations can  
251 be performed without iterative execution, which is not fully supported in the current eBPF and XDP versions.  
252 *ii*) The computational intensity and complexity are low so that the processing delay is limited to an acceptable  
253 range, even without acceleration mechanisms, such as Single Instruction, Multiple Data (SIMD), hugepages, or  
254 CPU cache prefetching, which are not fully available in most in-kernel frameworks. *iii*) The implementation has  
255 no strict dependencies on specific running environments or frameworks. Basic functions with these characteristics,  
256 such as router, load balancer, network address translation (NAT), and packet filter, are suitable for kernel space  
257 implementation.

258 3) *Advanced Functions*: Compared to basic functions, advanced functions involve complex and compute-intensive  
259 operations with the following main characteristics: *i*) Both packet header and payload need to be processed. *ii*)  
260 Acceleration mechanisms in the user space are required to keep processing delays within reasonable ranges. *iii*) The  
261 implementation of advanced functions typically requires specific runtime or execution environments. For example,  
262 the widely used open source network coding library Kodo [105] requires the C++ runtime environment, which is  
263 not available in the kernel space. According to these characteristics it is beneficial to implement advanced functions,  
264 such as data encryption, compression, and network coding, in the user space.

## 265 C. Selecting VNF Implementations for VNF Classes

266 Compared to the numerous kernel-bypassing approaches [94], [106], [107], such as Netmap, PF\_RING, DPDK,  
267 relative few in-kernel fast packet IO approaches have been developed. To the best of our knowledge, XDP is

268 currently the fastest in-kernel programmable network data path which provides bare metal packet processing at the  
269 lowest point in the software stack [84], [108]. Therefore, we select XDP for the in-kernel VNF implementation in  
270 CALVIN.

271 Based on a thorough literature review that covered [94], [106], [107] and related studies, we selected DPDK  
272 for the user space VNF implementation, mainly for the following reasons: i) High performance: According to the  
273 evaluation in [94], DPDK has demonstrated favorable bandwidth and latency performance. ii) Open source, low  
274 level with high configurability and very good documentation: This allows our VNF implementation to have full  
275 control of all processing functions invoked in the packet IO and to adopt design optimizations to reduce latency.  
276 For example, we have integrated the network coding library Kodo [105] with the native DPDK data structure to  
277 avoid data transfer overheads, thus achieving short latencies for advanced network coding VNFs. iii) Wide support:  
278 DPDK supports a wide range of physical, paravirtualized, and software NICs. The Open vSwitch also implements  
279 the DPDK fast path [54], simplifying the deployment of CALVIN on the OpenStack Cloud platform, which uses  
280 Open vSwitch as default software bridge.

281 However, we acknowledge that bypassing the kernel with DPDK (Version 18.02) has several disadvantages [108]:  
282 i) The driver can only run in the polling mode. ii) Network and upper layer protocols require third-party imple-  
283 mentations (whereas these protocols are already implemented in the mainline Linux kernel). iii) The Linux kernel  
284 has its own security model for managing the networking hardware. Bypassing the kernel requires a new security  
285 model in the user space, which is an important direction for future research.

#### 286 *D. Architecture Design*

287 The architectural design of CALVIN is illustrated in Fig. 6. CALVIN is built on top of the research-oriented  
288 SFC framework SFC-Ostack [46], [52]. We extend the SFC-Ostack control and data plane components to enable  
289 latency optimization strategies:

- 290 • Control plane: We add the VNF classifier module which translates between the VNF description and the SFC  
291 description, including the VNF classification into basic and advanced functions. The processing pipeline of  
292 these VNFs is then converted into a function chain of VMs with their networking configurations. The life-  
293 cycle management of all instances in the SFC description is handled by the SFC manager that communicates  
294 with OpenStack services.
- 295 • Data plane: In the data plane, multiple VMs are launched to run service functions over several physical  
296 compute nodes. In each VM, the service function is positioned either in the kernel space or in the user space.  
297 Packets are received from the ingress interface (vNIC), processed by the function program, and sent out  
298 through the egress interface (vNIC). VMs in the service chain are connected in a prescribed order by Open  
299 vSwitch with DPDK (OVS-DPDK).

300 Notice in Fig. 6 that CALVIN uses the distributed mapping with one VM for one VNF (and not the compact  
301 mapping of multiple VNFs in the same VM). This CALVIN design choice is mainly based on the reasoning in  
302 Section II-B3 about the flexibility of VNF resource allocation and on the results in [25], which indicate that for  
303 non-trivial computational VNF tasks, i.e., in particular for our advanced functions, the distributed mapping achieves  
304 lower latencies.

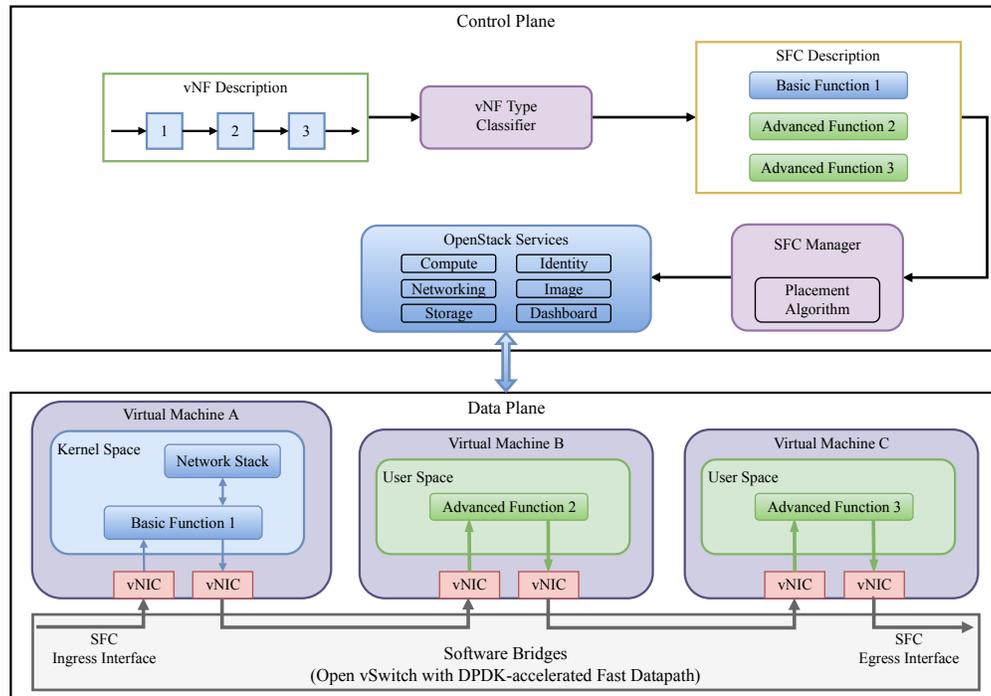


Fig. 6: Illustration of CALVIN architectural design, including fundamental components in control and data plane: The SFC manager utilizes OpenStack services to deploy VNFs (which are classified into basic and advanced functions) into a chain of VMs. The functions are implemented either in the kernel or user space. VMs in the SFC are connected by Open vSwitch with DPDK datapath (OVS-DPDK).

### 305 E. Setup and Workflow

306 The CALVIN operation with the distinct processing of basic and advanced functions requires various settings for  
 307 both hardware and software used by the OpenStack cloud platform.

#### 308 1) Configurations for Accelerating Virtual Networking Infrastructure:

- 309 • The physical NIC of each node must support DPDK [101], [109] to deploy the official OVS-DPDK plugin  
 310 of OpenStack.
- 311 • Each compute node should have dedicated CPUs assigned only for OVS-DPDK: The current version of OVS-  
 312 DPDK works in polling mode and any interruptions from other processes can cause performance degradations  
 313 and even lead to packet losses. Consequently, dedicated CPU cores must be configured for OVS-DPDK by  
 314 using the LINUX CPU isolation tool.
- 315 • The Input/Output Memory Management Unit (IOMMU) should be enabled to allow guest VMs to directly  
 316 use physical NICs through Direct Memory Access (DMA).
- 317 • Sufficient memory should be reserved to allocate hugepages for OVS-DPDK on all nodes. For compute nodes,  
 318 additional memory needs to allocate hugepages for the guest OS of each VM instance.

319 2) VNF Processing Configurations: Figs. 7 and 8 present the workflow of running basic functions in the kernel  
 320 space and running advanced functions in the user space; the full source code is available from [53]. Since the  
 321 Kernel-based Virtual Machine (KVM) is the default hypervisor of the OpenStack computing service (until latest

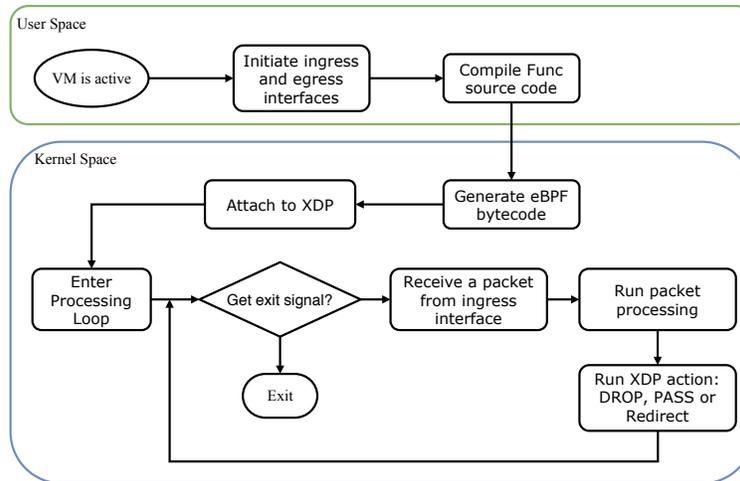


Fig. 7: CALVIN workflow for basic functions running in the kernel space

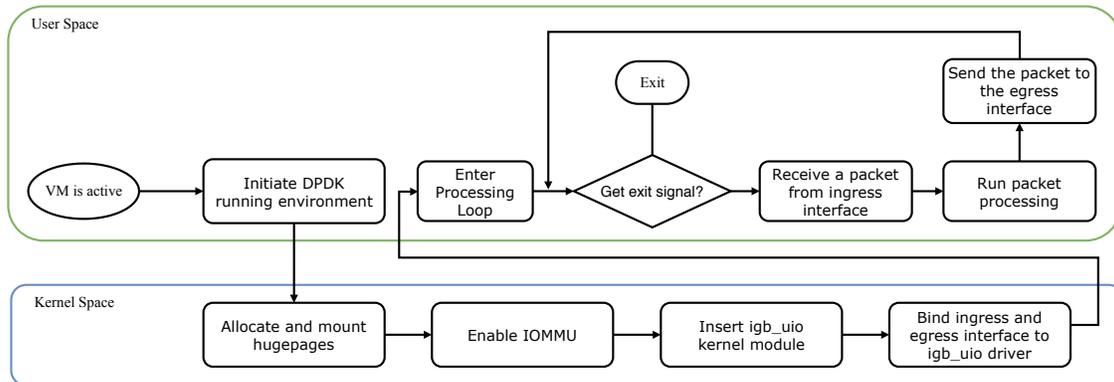


Fig. 8: CALVIN workflow for advanced functions running in the user space

322 version: rocky) [110], the following configurations are deployed for VMs launched by the KVM hypervisor:

323 a) *Kernel Space:*

- 324 • Due to the XDP requirements [111], the virtual interface of each VM should support the allocation of
- 325 a dedicated transmit queue (TX queue), e.g., through the virtio\_net patch [111] to OpenStack or through
- 326 extending the OpenStack Nova project.
- 327 • The Linux kernel of the guest OS running inside the VM should be updated to at least version 4.8 to run the
- 328 XDP program [112].
- 329 • The eBPF compiler framework BCC [112] should be installed to compile the XDP program and to attach
- 330 the compiled program to the virtual interfaces.

331 b) *User Space:*

- 332 • Virtual interfaces should be assigned with IOMMU for minimal latency overhead.
- 333 • Sufficient memory should be allocated for hugepages. These memories are used not only to initialize the
- 334 running environment of each DPDK application, but also to store packets that must be queued before
- 335 transmission.

- The DPDK kernel module `igb_uio` needs to be loaded to enable the kernel Poll Mode Drivers (PMD) driver.

This driver should be bound to the VM ingress and egress interfaces.

We note that exchanging data between VNFs purely in user space in a given VM is not straightforward. Normally, Inter Process Communication (IPC) mechanisms e.g., the Unix Domain Socket, are used to exchange data between processes running on the same VM. However, the IPC mechanisms are provided by the OS, which inject packets back into the kernel space. In order to avoid being forced back into the kernel space for inter-VNF data exchanges, we adopt the distributed mapping of one VNF per VM for CALVIN. Both the centralized approach and CALVIN use fast DPDK frame IO (which is substantially faster than the conventional raw socket frame handling by the OS) for the ingress and egress vNICs. The critical difference between CALVIN and the centralized approach is that CALVIN avoids injecting packets back into the kernel space of a given VM.

We also note that legacy advanced VNF, such as network coding function programs, have been written for normal layer 3 or raw sockets, which are typically not directly compatible with DPDK. For CALVIN, we implemented the examined advanced VNFs that run in user space based on the DPDK architecture and data structures. Building advanced VNFs on the native DPDK data structures avoids data transfers, aiding in achieving short latencies.

*3) Future Research Directions for CALVIN Workflow:* The current CALVIN version developed and evaluated in this study employs the distributed mapping of one VM per one VNF. The data is exchanged between VMs with low-latency OVS-DPDK software bridging, see Fig. 6, and is processed with the respective workflows illustrated in Figs. 7 and 8. The software bridging avoids the latencies due to the conventional IPC mechanisms provided by the OS. Ongoing research and development of kernel space and user space processing may enable low-latency data exchanges between distinct VNFs purely in the kernel space or purely in the user space. For instance, eBPF supports tail calls and maps to chain and to exchange data between multiple eBPF programs purely in the kernel space. Future research could design and evaluate low-latency XDP-based in-kernel SFC mechanisms that could process the elementary and basic VNFs. In the user space, the principle of shared memory regions [113], [114] could enable low-latency data exchanges between VNFs and form the foundation for low-latency DPDK-based user space SFC mechanisms for processing advanced VNFs.

Another (albeit less important) motivation for the distributed mapping in the current CALVIN version is the flexibility of dynamic resource scaling of distributed VMs. Dynamic VM resource scaling without service downtime can be achieved through live resizing, which has recently become feasible in the commercial VMware integrated OpenStack [115], but is not yet fully supported in OpenStack [116]. Alternatively, a new VM with more vCPUs can be provisioned when VNFs become bottlenecks. The compact mapping with all VNFs operating on a single VM requires that the old VM and the new VM operate simultaneously during the transition phase, which can be demanding for the typically limited memory resources of cloud compute systems. The distributed mapping allows the transition to be performed sequentially, with only the old VM and new VM supporting one VNF operating simultaneously. Nevertheless, the ongoing advances may make dynamic resource scaling (live resizing) of VM instances common in the future. It would then be interesting to examine live resizing in the context of low-latency VNF processing.

As low-latency inter-VNF data exchanges purely in the kernel or user space in a given VM and live resizing of VMs mature, it would be interesting to develop a compact mapping version of CALVIN that processes all basic VNFs in a single kernel-space-focused VM and all advanced VNFs in a single user-space-focused VM. It would

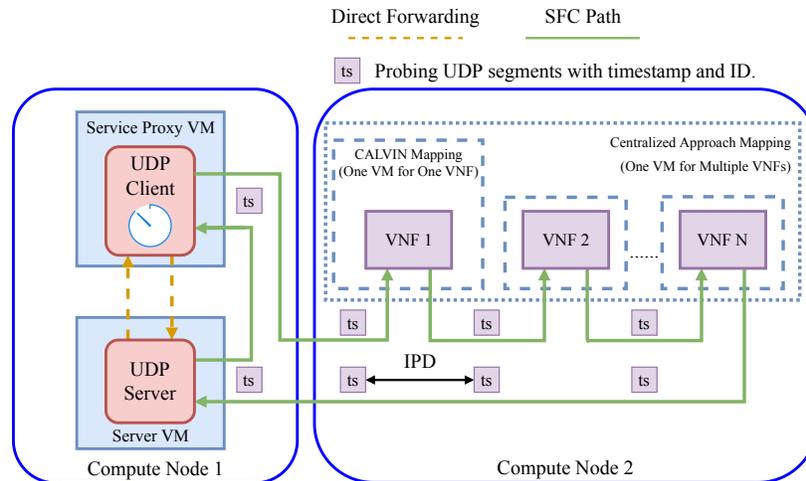


Fig. 9: Illustration of RTT measurement set-up: The service proxy and the server are allocated on compute node 1. Multiple VNFs are launched on compute node 2 to implement the SFC for the probing UDP traffic generated by the service proxy. The RTT is measured at the UDP client.

375 then be interesting to compare this compact mapping CALVIN version with the distributed mapping CALVIN that  
 376 is examined in this paper.

#### 377 IV. PERFORMANCE EVALUATION OF ELEMENTARY AND BASIC FUNCTIONS

378 In order to evaluate the feasibility of CALVIN for the 1 ms control loop of the tactile Internet, we first conduct  
 379 measurements for elementary and basic functions with a minuscule computational workload in this section. The  
 380 measurement results for elementary and basic functions indicate the baseline latencies. The evaluation methodology  
 381 and testbed are extended from the testbed introduced in our conference paper [46].

##### 382 A. Measurement Set-up

383 Fig. 9 illustrates the measurement set-up. VMs are launched on two compute nodes to measure the latency of the  
 384 elementary service loop introduced in Section II-A. We use an active measurement strategy to measure the latency  
 385 for UDP traffic. We consider the UDP protocol for the probing traffic since UDP is typically used for low-latency  
 386 applications. Latency measurements for TCP are difficult due to the complex TCP flow control and congestion  
 387 control. UDP traffic provides us with full control of both the sending and receiving processes.

388 1) *Architecture*: The measurement setup is based on the service loop in Fig. 3. The service proxy runs the UDP  
 389 client, which sends probing packets to the server located on the same compute node 1. The server simply bounces  
 390 all received packets back to the client as fast as possible. Alternatively, probing packets can be forwarded directly  
 391 to the server to measure the direct forwarding latency (without any VNFs) of the underlying network infrastructure.

392 In order to reflect authentic practical networking scenarios, the service proxy VM and the server VM do not use  
 393 fast IO technologies, such as DPDK or XDP. In particular, the service proxy and the server are normally working  
 394 in the network layer (in contrast to VNFs running in the data link layer). Therefore, we employ the conventional  
 395 Linux networking stack in the service proxy VM and server VM, so as to capture the latencies introduced by the  
 396 networking layer.

397 Compared to the centralized measurement setup in [86], the components of our measurement architecture are  
398 distributed over two different physical nodes, which mimics practical cloud environments.

399 A time stamp (ts) and an identification number (ID) are added before the payload part of each probing packet.  
400 The time stamps are used for the latency calculation. The IDs are used to identify out-of-order packets which  
401 would indicate the accumulation of packets in queues [117]. We set the packet traffic rates sufficiently low, see  
402 Section IV-A5 to avoid out-of-order packets.

403 2) *Testbed*: All measurements were performed on our NFV testbed consisting of off-the-shelf computers that are  
404 connected by two separate Gigabit Ethernet networks. Each computer had 4 CPU cores (Intel 4th Generation Core  
405 i5), 16 GB RAM, 128GB SSD, and two Gigabit NICs (Intel 9301CT Gigabit CT). OpenStack (Pike version) [118]  
406 was deployed on these computers, installed with the Ubuntu Server 16.04 TLS operation system. For each node,  
407 one NIC is used for management and public traffic, and the other NIC is used to build a separate internal data  
408 network for the virtual instance. The purpose of using a separate network is to reduce the impact of management  
409 and public traffic on latency measurements. Besides the compute, networking, identification, and storage services  
410 of OpenStack, the official SFC and OVS-DPDK plugins were installed. In particular, we used Virtio [119] for the  
411 vNIC and OVS-DPDK as the virtual bridge. For the management of the virtual instances, KVM was used as the  
412 hypervisor and the customized Ubuntu Cloud image was used to implement different VNFs.

413 3) *Elementary and Basic VNFs*: We implemented and measured elementary forwarding and two basic VNFs.

414 FWD Elementary Forwarding (FWD): Packets are retrieved from the VNF ingress interface and directly forwarded  
415 to the egress interface without any operations. FWD is an elementary function of each VNF and other VNF  
416 functions are built on top of the FWD function.

417 ATS Appending Time Stamps (ATS): The timestamps of the reception and transmission of a given packet by the  
418 current VNF are appended to the end of the UDP payload of the packet just before the packet is sent out.  
419 The ATS function modifies the payload size; therefore, the layer 3 and layer 4 header checksums must be  
420 recalculated. In order to ensure fair latency comparisons, we recalculate the checksum with standard methods  
421 in software (and do not employ checksum offloading). Thus, the ATS function can be used to estimate the  
422 latency introduced by a trivial operation on the packet payload.

423 XOR XORing UDP payload (XOR): This function performs an XOR operation with the same static key on all  
424 bytes of the UDP payload. Compared to ATS, the XOR function can be used to estimate the additional  
425 latency introduced by a non-trivial computational operation on the packet payload.

426 We make the source code of all VNF implementations openly available at [53].

427 4) *Metrics*:

428 a) *Latency*: We adopt the Round-Trip-Time (RTT) (i.e., the per-packet delay) of each UDP packet sent by  
429 the UDP client program running on the service proxy VM as the latency metric for the following reasons: *i*) The  
430 RTT includes the delay introduced by the forward and backward paths. *ii*) The RTT measurement does not require  
431 time synchronization between the VMs. (According to the experiments in [120], VM-level time synchronization is  
432 error prone due to the interference between the VMs.) Note that the measured RTT includes all latency components  
433 (1)–(5) illustrated in Fig. 4 across the entire SFC.

434 b) *Bandwidth*: We measure the maximum achieved bandwidth (packet throughput in bit/second) with iPerf  
435 (version 2) in the UDP mode [121] (We did not use iPerf for RTT measurements because iPerf currently does not

436 support per-packet delay measurements; we wrote our own Python tools [53] for the per-packet RTT measurements.)  
437 We run an iPerf UDP bounce server on the server VM and an iPerf UDP client on the service proxy VM (all on  
438 compute node 1 in Fig. 9). Through trials that manually adjusted the iPerf client target bandwidth in steps of  
439 10 kbit/s, each lasting for five minutes (with ten independent repetitions), we determined the maximum bandwidth  
440 such that there are no lost or out-of-order packets detected at the iPerf client side.

441 5) *Active Measurement Parameters:* Active measurements require the setting of two UDP traffic parameters,  
442 namely the Inter-Packet-Delay (IPD) and the payload size. Based on our preliminary measurements (which are not  
443 included in detail due to space constraints), relative small IPDs on the order of a few milliseconds give consistent RTT  
444 values. Much shorter IPDs and correspondingly high traffic rates would lead to queueing for the VNF processing,  
445 as analyzed in [57], [58], [61], [63], [64], [66]–[68]; in contrast, our measurements focus on the latencies of lightly  
446 loaded systems without significant queueing. The RTT values slightly increase with increasing IPD, which is mainly  
447 due to OS batching mechanisms and the queueing mechanisms of the underlying virtual bridges that are activated for  
448 low network traffic loads. To avoid the additional latency introduced by batching (queueing) and to ensure consistent  
449 measurements, we set the IPD to 5 ms in all measurements.

450 We select 256 and 1400 bytes as the lower and upper limits of the payload size. According to our preliminary  
451 evaluations, UDP segments with less than 256 bytes of payload cannot be properly processed by XDP. At the  
452 same time, the official SFC plugin of OpenStack [122] currently does not support jumbo frames. The maximum  
453 UDP payload size depends on the Maximum Transmission Unit (MTU) of the underlying physical network. For  
454 the default MTU of Ethernet of 1500 bytes, the maximum UDP payload size is limited to 1472 bytes (1500 – 20  
455 (minimum IPv4 header) – 8 (UDP header)). A payload size of 1400 bytes is chosen to provide enough free spaces  
456 reserved for IP header options or additional service function related headers.

457 For each scenario, the measurements were repeated 50 times; for each measurement, 500 UDP segments were  
458 sent by the probing client.

459 6) *Measurement Scenarios:*

460 a) *Comparison of Different VNF Technologies:* We selected XDP [84] to implement VNFs in kernel space  
461 and DPDK [101] to implement VNF in user space for CALVIN (see Section III-C). To benchmark the latencies  
462 of these two selected technologies, we also consider two other frequently considered VNF technologies, namely  
463 Linux Kernel Forwarding (LKF) as a kernel space technology and the Click modular router [123] as a user space  
464 technology.

465 XDP: The XDP workflow flow is outlined in Section III-E. Due to following XDP restrictions, only the FWD  
466 and XOR functions with the 256 bytes UDP payload size are implemented: *i*) The maximum number of  
467 instructions per XDP program is currently 4096 BPF instructions [124], limiting the complexity of the  
468 computational operations as well as the data amounts that can be manipulated. *ii*) The range of the XDP  
469 operational memory for each packet is limited by the size of the original received packet. Operations outside  
470 this range are prohibited. Thus, the ATS function cannot be implemented with XDP.

471 DPDK: The high flexibility and programmability of DPDK allow all three functions to be implemented as DPDK  
472 applications. By default, DPDK applications run in polled mode and can consume 100% of the CPU resources.  
473 The CPU and corresponding power consumption could be reduced with a sleeping mechanism. To minimize  
474 the processing delay, we set the number of packets in a processed batch (burst) to one.

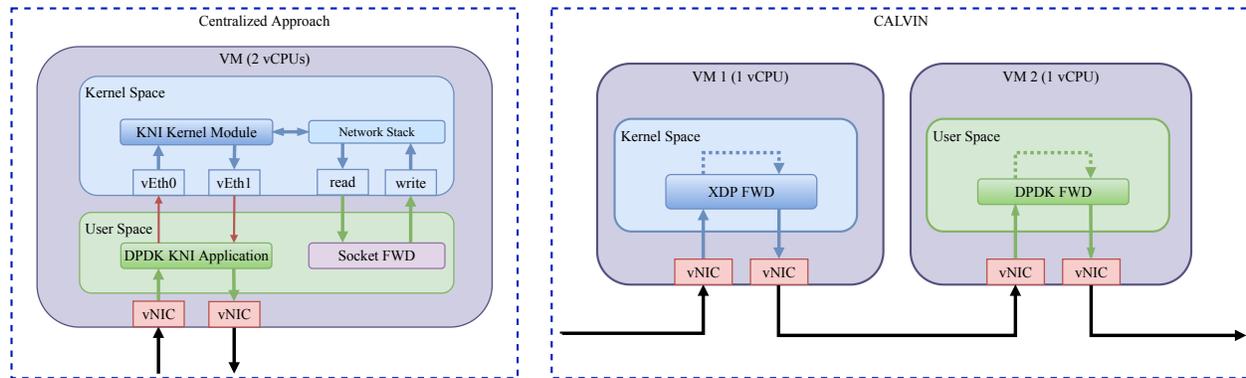


Fig. 10: Detailed illustration of measurement set-up for the RTT comparison between centralized approach and CALVIN on compute node 2 of Fig. 9.

475 LKF: Linux Kernel Forwarding (LKF) is a built-in Linux kernel feature for network-layer packet forwarding [125].  
 476 Due to its trivial operations, LKF is one of the fastest functions running in the kernel space. Since LKF does  
 477 not provide any programmability, LKF cannot be used to implement different VNFs in the kernel space.

478 Click: The Click modular switch is a software framework for building flexible and configurable routers [123].  
 479 Compared to the low-level IO operations offered by DPDK, Click provides multiple encapsulated packet  
 480 processing modules called elements to build processing pipelines. Elements can be connected with a directed  
 481 graph described in a router configuration file. The evaluated VNFs are implemented by combining built-in  
 482 and customized elements.

483 For the VNF technologies comparison, we launch a single VM on compute node 2 in Fig. 9 to run the network  
 484 function.

485 *b) Comparison of Centralized Approach [51] vs. CALVIN:* We benchmark the proposed CALVIN approach  
 486 against the state-of-the-art centralized approach [51], which was one of the first well-studied approaches for  
 487 implementing advanced functions, such as network coding, as VNFs. (We re-implemented the centralized approach as  
 488 its source code was not publicly available.) In order to examine the distributed VNF aspect of CALVIN, we consider  
 489 two FWD VNFs on two separate VMs for CALVIN, while for clarity of comparison we implement only one FWD  
 490 VNF in the centralized approach. We give the centralized approach a VM with twice the computational resources  
 491 (vCPU and memory) compared to each individual VM used by CALVIN. Thus, the comparison is effectively  
 492 between two FWD VNFs in CALVIN and one FWD VNF in the centralized approach, whereby both approaches  
 493 having the same computational resources. More specifically, for the centralized approach, which is illustrated in  
 494 the left part of Fig. 10, packets enter the VM through the left vNIC, traverse the FWD VNF and exit through right  
 495 vNIC. For CALVIN, which is illustrated in the right part of Fig. 10, packets enter the left vNIC of VM1, traverse  
 496 the XDP FWD, exit through the right vNIC of VM1, enter the left vNIC of VM2, traverse the DPDK FWD, and  
 497 exit through the right vNIC of VM2.

## 498 B. Results

499 *1) RTT Measurements of Elementary and Basic VNFs for Different VNF Technologies:* The average values and  
 500 95% confidence intervals of the measured RTTs of the elementary and basic VNFs implemented (one VNF on

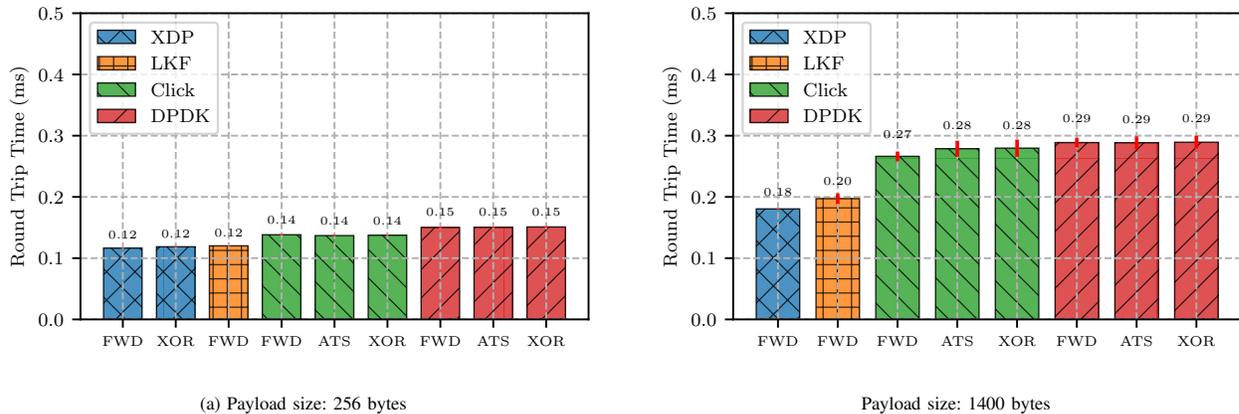


Fig. 11: Means and 95% confidence intervals for round-trip time (RTT) of different VNF technologies in kernel space (XDP and LKF) and user space (Click and DPDK). The 95% confidence intervals for the 256 byte payload size are very tight and barely visible in the plot.

compute node 2) with different VNF technologies are plotted in Fig. 11. We observe from Fig. 11 that the RTTs for the two basic VNFs, namely ATS and XOR, are equivalent to the respective RTTs for the elementary FWD VNF for all VNF technologies. This result implies that the additional latencies for the payload processing in the basic ATS and XOR VNFs are negligible compared to the elementary FWD latencies.

Examining closely the in-kernel VNF technologies, we observe from Fig. 11 that although XDP and LKF have the same FWD RTTs for small packets, XDP FWD is around 10% faster than LKF FWD for large packets. Moreover, we observe from Fig. 11 that the RTTs for in-kernel processing (XDP and LKF) and user space processing (Click and DPDK) are very similar for the small 256 bytes packet size. On the other hand, for the large 1400 bytes packet size, user space processing incurs substantially longer (about 50% longer) RTTs than in-kernel processing.

These observed latency differences appear to be primarily due to the two main types of overhead introduced by basic functions: (i) The overhead of copying frames [126] from the virtual NIC ring buffer to the VM memory, and (ii) the overhead of additional metadata pre-processing of the frames. In the Linux networking stack (LKF), a data structure `sk_buff` containing some metadata is allocated for each received frame. The latency overheads for extracting the metadata and allocating `sk_buff` are non-trivial [24], [84], [127]. In contrast, XDP operates at the lowest point in the Linux software networking stack, without allocating a metadata data structure, nor any parsing and pre-processing of packets. XDP supports the packet forwarding to another interface (XDP\_REDIRECT Action) without the metadata processing and without checking the routing table of the kernel (MAC addresses are provided and written by the eBPF program). For the conventional Linux kernel forwarding (LKF), the metadata structures must be created and the kernel requires checking the routing table to write the proper MAC addresses. Thus, LKF is slightly slower than XDP for large packets (which require more time to extract the metadata). In summary, both XDP and LKF copy the frames into kernel space memory; XDP does not allocate metadata structures, whereas LKF does allocate metadata structures.

Turning to the user space technologies, Click and DPDK copy the frame data into user space memory and allocate different metadata structures. In particular, in DPDK (kernel bypassing), the Poll Mode Driver (PMD) [128] copies

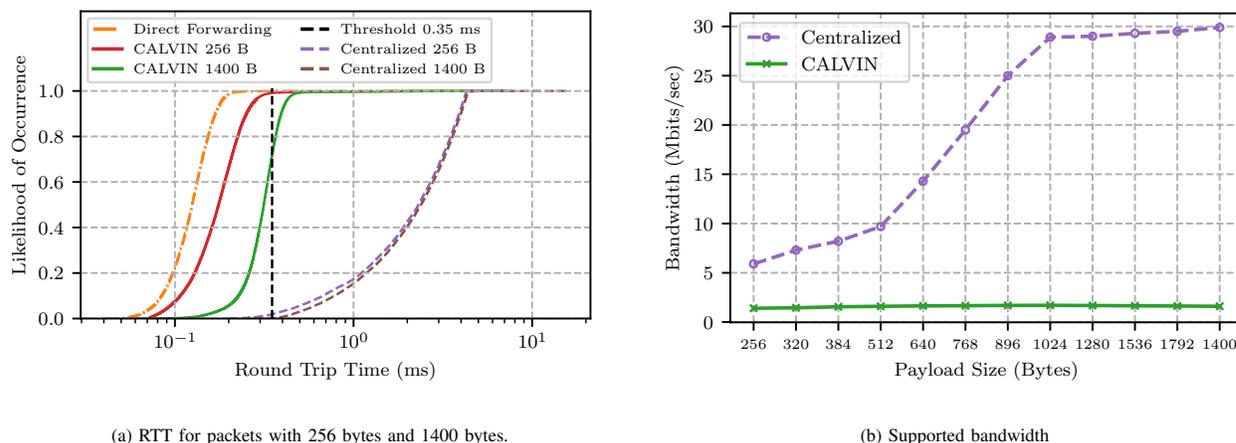


Fig. 12: FWD VNF performance comparison: Proposed CALVIN (two FWD VNFs, namely XDP FWD and DPDK FWD, see right side of Fig. 10) vs. the state-of-the-art centralized approach [51] (one FWD VNF, see left side of Fig. 10), whereby both CALVIN and centralized approach utilize same computational resources (two equivalent vCPUs).

frames from the NIC ring buffer to the DPDK user space memory pool (a pre-allocated fixed-length memory using hugepages make it resident in the physical memory [101], [129]). During this copying, an additional `mbuf` data structure [130] containing metadata is created for each frame. These metadata are required by DPDK for further processing and many advanced features. For large packets, this copying into the user space memory pool and metadata structure creation can incur significantly longer latencies compared to the in-kernel approaches. Click involves similar latency increasing copy and metadata operations.

## 2) Comparison Between Centralized Approach and CALVIN:

*a) RTT:* Fig. 12(a) presents the RTT measurement results for the elementary FWD VNF. As presented in Fig. 9, the direct forwarding is the baseline of the transmission delay introduced by the underlying virtual networking infrastructure. We observe from Fig. 12(a) that the centralized approach exceeds the 0.35 ms delay threshold in the best case, while the CALVIN RTT is below the 0.35 ms threshold with about 70% probability for 1400 bytes packets, and with nearly 100% for 256 bytes packets. The measured average (mean) RTTs (accounting for all latency components (1)–(5) in Fig. 4) for the 256 bytes and 1400 bytes packets are as follows: CALVIN: 0.19 ms and 0.32 ms, respectively; centralized approach: 2.30 ms and 2.39 ms, respectively. Thus, we conclude that CALVIN can meet the low latency requirement introduced in Section I and allow additional data processing inside the virtualized environment.

We remark that the measurements without fast packet IO in [25] indicated that for VNFs with very low computation demands, the compact mapping gave lower latencies than the distributed mapping, while for computationally intensive VNFs, the distributed mapping gave lower latencies. Our latency measurements with our fast packet IO based CALVIN demonstrate that the distributed mapping achieves low latencies within the latency constraints of the tactile Internet for elementary and basic VNFs with low computational demands. Intuitively, with the accelerated packet IO and the correspondingly supported high packet rates and high CPU resource consumption for fast packet IO (compared to the conventional slow networking stack), even VNFs with low computational processing demands

do no longer present a trivially low computation load on the CPU. Thus, we believe that the distributed mapping is a good design choice for implementing all classes of low-latency VNFs in CALVIN. The distributed mapping allows each VNF to focus on conducting its processing tasks quickly in its own space and leaves the communication to the underlying software integration bridge, which has mature low-latency performance. The allocation of one VM per VNF with the distributed mapping may be considered wasteful; however, each VM could be optimized for its usage and the distributed mapping is consistent with the emerging unikernel concept [131], [132]. For example, if the packet processing occurs mainly in the userspace, then the kernel components could be slimmed down to make the kernel space as small as possible.

*b) Bandwidth:* Fig. 12(b) presents the bandwidth measurement results for payload sizes ranging from 256 to 1400 bytes. We observe from Fig. 12(b) that the bandwidth supported by the centralized approach is substantially higher than the CALVIN bandwidth, especially for large packets. For a payload size of 1400 bytes, the centralized approach bandwidth is over 15 times higher than the CALVIN bandwidth. Compared to the centralized approach with a minimal supported bandwidth around 6 Mbits/s, the maximal CALVIN bandwidth is in the range from 1.4–1.7 Mbits/s.

*c) Latency-Bandwidth Trade-off:* Taken together, the RTT results in Fig. 12(a) and the bandwidth results in Fig. 12(b) demonstrate the tradeoff between per-packet latency and throughput in real system implementations. The centralized approach uses the standard Linux socket implementation which was designed primarily for high-throughput best-effort service applications, e.g., file transfers, which have typically bursty traffic. Accordingly, the kernel networking stack includes a range of throughput enhancing mechanisms, such as batch processing. Batch processing collects multiple packets and then processes the batch of packets with acceleration methods, e.g., with single instruction multiple data (SIMD) instructions and enhanced CPU caching. Batch processing increases the per-packet delay as the first packet in a batch must wait for subsequent packets to fill up a batch before processing commences.

The batch processing in the Linux networking stack slows down the centralized approach, even if no batch processing is employed in the DPDK user space application. This is because the KNI injects packets back into the normal Linux kernel networking stack, where batching is employed. The batching in the kernel stack cannot be avoided without modifying the kernel source code.

CALVIN avoids this batch processing in the kernel space by employing fast path technologies, e.g., XDP, for in-kernel VNF implementation. Since our goal is to reduce the per-packet delay, CALVIN avoids batch processing in all VNF implementations. The VNF implementations in CALVIN operate in a run-to-completion mode, i.e., they retrieve one packet, process it (batch size of one packet), and send it out as quickly as possible. CALVIN thus reduces the per-packet latency, as observed from Fig. 12(a), at the expense of supporting a lower packet throughput, as observed in Fig. 12(b).

For the tactile Internet for human-machine co-working, low per-packet delay is typically much more important than support for high bandwidth. The human-machine co-working packet traffic, e.g., the control messages of a robot arm are typically small so that support for low bandwidths is sufficient. Also, the 5 ms IPD is sufficient as a sample period of an angle sensor for a typical pendulum application, as illustrated by the right side of Fig. 1. On the other hand, control messages that are delayed by batch processing can profoundly disrupt human-machine co-working. Therefore, CALVIN trades support for only low bandwidth for reduced per-packet latency. The bandwidth supported

TABLE I: CPU usage of the physical compute node

Approach	User (%)	Sys (%)	Guest (%)	IDLE (%)
Centralized	25.2	47.2	2.9	24.7
CALVIN	25.2	23.0	2.2	49.6

by CALVIN can be improved in future work with bandwidth management mechanisms, e.g., a load balancer that distributes a packet flow over a set of duplicated VNFs to enable parallel processing. Parallel processing requires a fast per-flow load balancer and a traffic merger to handle out-of-order packets arising from the parallelism.

*d) CPU Resource Usage:* We measured the usage of the four cores of the physical CPU on compute node 2 in Fig. 9 with `mpstat` [133]. Since the scheduling of CPU resources for all running VMs is managed by the OpenStack compute service (Nova), the global average usage of all cores has been measured. We measured with a sample period of 1 second for a duration of 10 minutes. Table I shows the CPU usage levels of the centralized approach and CALVIN at the user level (User), kernel level (Sys), and for a niced guest (Guest).

We observe Table I that compared to CALVIN, the centralized approach doubles the CPU resources consumed at the kernel (Sys) level. KVM uses the Linux kernel of the host OS as the hypervisor and uses POSIX threads to implement vCPUs of the guest OS [134]; thus, the Sys CPU usage reflects the usage of the vCPU of the VM running the VNF. By avoiding the overhead of context switching and metadata processing of each vCPU, CALVIN significantly reduces the kernel (Sys) CPU usage of the host OS. Thus, CALVIN leaves a significantly higher proportion of time of the physical CPU idle. This higher CPU idle time achieved by CALVIN can be utilized to run cloud management services, software bridges, and other essential background processes to stabilize the latency performance.

## V. EVALUATION OF COMPUTATION-INTENSIVE ADVANCED VNFs

The evaluations in Section IV-B2 indicated that CALVIN can complete elementary VNFs with an RTT (within the MEC) on the order of 0.32 ms. Thus, considering the 0.35 ms MEC latency budget from Fig. 2 there is still a remaining latency budget of about 0.02 ms for some advanced VNFs. This section evaluates network coding and encryption as two examples of advanced VNFs that have relatively high computational demands. We first describe the practical applications and relevance of these two VNFs. Then, we use the test setup from Section IV-A to evaluate the processing delay incurred by these advanced VNFs. The purpose of this evaluation is to assess whether the RTT reduction for elementary VNFs achieved by CALVIN is sufficient to permit practical advanced VNFs within the latency requirements of the tactile Internet.

### A. Network Coding

Network coding linearly combines several original packets with coding coefficients to form encoded packets that are transferred through the network [135]. In Random Linear Network Coding (RLNC), the coding coefficients are randomly generated. The key benefits of RLNC include: i.) the ability to recode with partially received data at all nodes in the network without requiring coordination, thus being suitable for distributed environments [136], ii.) versatile coding matrix, permitting sparsity (judiciously added zeros) to reduce computation complexity [137],

618 [138], iii.) low latency support due to on-the-fly coding capabilities [139], [140], iv.) support of heterogeneous field  
619 sizes for communication entities, increasing flexibility in heterogeneous contexts [141], and v.) reduced overhead  
620 between the storage and transmission layers, as the same code can also be used for distributed storage [142].

621 Network coding research has proposed many different RLNC variants. We focus on the two main RLNC types,  
622 namely systematic block codes and convolutional sliding window codes. Block-based RLNC was introduced to  
623 reduce the computational requirements and control for network coding [143]. To further improve the performance,  
624 a systematic code does not code every packet, but sends original packets as “coded” packets [144]. The packets  
625 built from linear combinations are then sent in between original packets or at the end of the block [145].

626 Sliding window network coding has been introduced to reduce the in-order delay of coded transmissions [146].  
627 In the form of a systematic code with a limited coding window, sliding window network coding has shorter in-order  
628 delay compared to block codes, while generally requiring comparable computational resources [140].

629 Although network coding has been extensively studied in recent years, the deployment of RLNC in real-world  
630 networks is still rare. The main obstacle for the deployment of network coding is the limited availability of pro-  
631 grammable computing resources at network nodes, which are currently only used for switching and routing decisions.  
632 However, NFV and SDN provide new flexibilities for deploying innovative functions within a network [16]. With  
633 NFV, network coding can be implemented for abstract VMs or containers, which can be instantiated at arbitrary  
634 NFV-capable network nodes. In addition, SDN can direct the data flows towards the network coding VNFs and  
635 orchestrate them in an SFC [147], [148]. However, the latency of network coding as a VNF in a general-purpose  
636 MEC system has to the best of our knowledge not been previously examined in detail.

637 Our per-packet processing delay measurements consider the encoding (which is computationally equivalent to  
638 recoding in a network node) with a Galois field size of  $GF(2^8)$  and 25% redundancy. We consider block coding  
639 with a block size of 32 packets and sliding window coding with a window size of 8 packets.

## 640 *B. Encryption*

641 Encryption and decryption are critical security components in communication, ensuring the confidentiality and  
642 integrity of the transferred data. More than 40% of the web traffic is transported in encrypted form over HTTPS,  
643 with an increasing trend [149]. As a result, decryption is required for a multitude of network functions, e.g., caching  
644 and deep packet inspection. As with network coding, encryption requires the entire payload to be processed which is  
645 a considerable computational effort. We focus on the Advanced Encryption Standard (AES), which is a commonly  
646 used encryption standard for data transfers and storage.

647 A previous study showed a prohibitive end-to-end latency of at least 30 ms for encryption as a VNF [150].  
648 However, this previous study did not take advantage of fast packet processing mechanisms, such as DPDK. Other  
649 VNF implementations of AES encryption have used Graphics Processing Units (GPUs) to increase the throughput  
650 and scalability [151], [152], while incurring latencies of over 150  $\mu$ s [151]. In contrast, our CALVIN approach  
651 enables encryption of small packets within a 20  $\mu$ s delay budget on a general-purpose MEC system.

## 652 *C. Measurement Set-up*

653 Compared to the measurement set-up for elementary and basic functions described in Section IV-A, the following  
654 additional considerations are required for evaluating advanced functions:

655 *1) VNF Implementation:* Due to the limitations of in-kernel technologies, CALVIN uses DPDK to implement all  
656 advanced VNFs. Both network coding and AES encryption functions are implemented on top of the elementary  
657 DPDK FWD application. We implemented network coding with the Network Coding Kernel Library (NCKernel),  
658 which is built on top of the Kodo library [105], to support the different variants of network coded communications.  
659 We used the portable AES Implementation Tiny-AES-C [153] to build the encryption application.

660 We implement multiple VNFs in parallel (each VNF on its own VM according to the CALVIN architecture  
661 principles, see Section III-D) so as to evaluate the scalability of our VNF implementation. Scalability is a key  
662 performance indicator for virtualization systems since a key aspect of virtualization is to run multiple virtual  
663 instances on limited hardware resources.

664 *2) Metric:* Since the RTT of elementary forwarding has been evaluated in Section IV, the measurements for the  
665 advanced VNFs focus on the packet processing delay. We define processing delay as the time duration required  
666 for the complete processing of a packet by a VNF, i.e., as the latency component (1) in Fig. 4. For VNFs that can  
667 generate redundant packets, such as network coding, this processing delay also includes the time duration required  
668 to create redundant packets.

669 *3) Methodology:* In order to evaluate the impact of VNF processing demands on the processing delay, the  
670 computational operations should be performed in parallel. This requirement is very challenging if the probing  
671 traffic is generated by a remote VM. Accurate synchronization mechanisms would need to be deployed on the  
672 virtualized networking infrastructure to ensure that probing packets arrive at each VNF at the same time. Therefore,  
673 instead of using an additional client to generate probing UDP traffic, for the evaluation of the advanced VNFs, the  
674 UDP segments are generated locally by each allocated VM. The locally generated traffic ensures that the VMs are  
675 continuously backlogged so that we obtain the worst-case processing delay: Every VM is always busy working and  
676 the OpenStack scheduler needs to handle the resource allocation among them. The delay values of warm-up and  
677 tail probing packets are not included in the measurement results. For each number of VNFs, 50000 valid probing  
678 packets are generated for processing.

#### 679 *D. Evaluation*

680 The evaluation of elementary functions in Section IV-B2, indicated a mean delay of 0.32 ms for the elementary  
681 FWD VNF of 1400 bytes packets in CALVIN. Considering the MEC latency budget of 0.35 ms from Fig. 2 and  
682 a safety margin around 0.01 ms, we consider a latency budget of 20  $\mu$ s for the advanced function processing (for  
683 smaller packets this latency budget could be larger as 256 bytes packets had only 0.19 ms mean RTT and 0.25 ms  
684 90%ile RTT in Section IV-B2).

685 Figure 13 shows the measured processing times for small 256 bytes packets and large 1400 bytes packets. For  
686 small packets, the processing time is within the 20  $\mu$ s requirement for all evaluated functions. With increasing  
687 number of VNFs, the load on the CPU increases and the processing time increases linearly as soon as the number  
688 of VNFs exceeds the number of CPU cores that are available exclusively for VM processing (one of the available  
689 four CPU cores is heavily utilized by the OVS-DPDK software bridge, which runs in polling mode with default  
690 DPDK functionalities). The latency increase is due to the contention for CPU resources. For a prescribed maximum  
691 latency requirement, e.g., 5  $\mu$ s, we can read off the number of permitted parallel running VNFs, e.g., three VNFs.

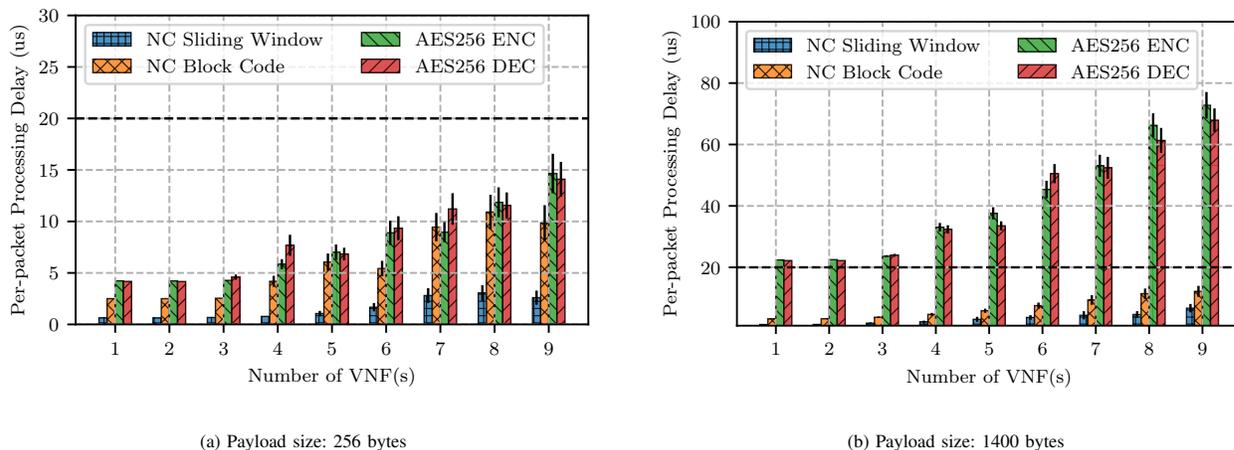


Fig. 13: Means and 95% confidence intervals for processing times in microseconds for computationally intensive advanced VNFs: Per-packet processing latency in given VNF (latency component (1) in Fig. 4) on a given VM as a function of the number of VNFs running in parallel (with one VNF per VM) on one compute node with four CPU cores.

692 With a load balancer redirecting a given flow to multiple VNFs (and VMs), the larger numbers of supported VNFs  
 693 would correspond to increased supported throughput.

694 For the large 1400 byte packets, we observe from Figure 13 increased processing times compared to the small  
 695 256 bytes packets. While the processing time for network coding remains relatively low and well within the 20  $\mu$ s  
 696 budget, encryption is not feasible even when the VNFs have exclusive access to a CPU core, i.e., for three or less  
 697 VNFs. For network coding, the sliding window code has substantially shorter processing times than the block code.  
 698 Even for large packets and high contention for the CPU resources, e.g., for nine parallel VNFs, the sliding window  
 699 network coding delays remain below 7  $\mu$ s.

## 700 VI. CONCLUSION

701 We have designed, implemented, and evaluated Chain bAsed Low latency VNF ImplemeNtation (CALVIN),  
 702 an approach for managing distributed service function chains (SFCs) for low-latency tactile Internet applications.  
 703 CALVIN implements virtual network functions (VNFs) either in the kernel space (if VNFs require only simple  
 704 processing) or in the user space (if VNFs require advanced processing) so as to avoid transmissions between kernel  
 705 space and user space for processing a given VNF. CALVIN further implements VNFs in a distributed manner with  
 706 one VNF per VM and employs fast packet input/output (IO) to avoid the metadata and batch processing of the  
 707 conventional Linux network stack.

708 We initially measured the elementary forwarding latencies of various current VNF implementations. We found  
 709 that the eXpress Data Path (XDP) achieved latencies of 120  $\mu$ s for small payloads and 180  $\mu$ s for large payloads,  
 710 while the native Linux kernel incurred about 10% higher forwarding latencies. The Data Plane Development Kit  
 711 (DPDK) approach and the Click router approach have up to 50% higher latencies than XDP. Based on these  
 712 measurements, we adopted XDP for implementing computationally simple VNF in CALVIN, while we adopted  
 713 DPDK for implementing computationally complex VNFs in CALVIN.

714 We extensively benchmarked CALVIN against the state-of-the-art centralized SFC management approach [51],  
715 which processes a given VNF with both the kernel space and the user space. Our measurements demonstrated  
716 that CALVIN achieves significantly shorter latencies (0.32 ms mean latency for 1400 byte packets) for an SFC  
717 consisting of two distributed elementary forwarding VNFs (XDP forwarding and DPDK forwarding) compared to a  
718 single elementary forwarding VNF in the centralized approach (2.39 ms for 1400 byte packets). On the downside,  
719 CALVIN supports only a lower packet throughput (bandwidth) of around 1.5 Mbit/s than the centralized approach  
720 (between 6 and close to 30 Mbit/s depending on the packet size). CALVIN thus trades in reduced packet throughput  
721 in order to achieve shorter per-packet latency, which is required for typical tactile Internet applications with a 1 ms  
722 round-trip delay budget.

723 There are many important future research directions for SFC management in the tactile Internet. The implemen-  
724 tation and measurements reported in this article have focused on the network function virtualization (NFV) in the  
725 MEC, i.e., the rightmost dashed box (the MEC cloud) in Fig. 2. Future research could integrate the MEC into a  
726 holistic 5G testbed that encompasses the entire end-to-end sensor-to-actuator loop in Fig. 2. Another direction is  
727 to examine novel “compute and forward” functions, such as video frame preprocessing for object detection [154]  
728 or transcoding [155], that bring more intelligence to the network edge [156]. The video preprocessing with limited  
729 edge cloud computing could extract key information to reduce the amount of data that needs to be transmitted  
730 through the network to the remote computationally powerful cloud.

## 731 REFERENCES

- 732 [1] K. Antonakoglou, X. Xu, E. Steinbach, T. Mahmoodi, and M. Dohler, “Toward haptic communications over the 5G tactile internet,”  
733 *IEEE Commun. Surv. & Tut.*, vol. 20, no. 4, pp. 3034–3059, Fourth Qu. 2018.
- 734 [2] H. Chen, R. Abbas, P. Cheng, M. Shirvanimoghaddam, W. Hardjawana, W. Bao, Y. Li, and B. Vucetic, “Ultra-reliable low latency  
735 cellular networks: Use cases, challenges and approaches,” *IEEE Commun. Mag.*, vol. 56, no. 12, pp. 119–125, Dec. 2018.
- 736 [3] K.-C. Chen, T. Zhang, R. D. Gitlin, and G. Fettweis, “Ultra-low latency mobile networking,” *IEEE Network*, in print, 2019.
- 737 [4] O. Holland, E. Steinbach, R. V. Prasad, Q. Liu, Z. Dawy, A. Aijaz, N. Pappas, K. Chandra, V. S. Rao, S. Oteafy *et al.*, “The IEEE  
738 1918.1 “Tactile Internet” standards working group and its standards,” *Proc. IEEE*, in print, 2019.
- 739 [5] Z. Hou, C. She, Y. Li, T. Q. Quek, and B. Vucetic, “Burstiness-aware bandwidth reservation for ultra-reliable and low-latency  
740 communications in tactile internet,” *IEEE J. on Selected Areas in Commun.*, vol. 36, no. 11, pp. 2401–2410, Nov. 2018.
- 741 [6] A. Nasrallah, A. Thyagaturu, Z. Alharbi, C. Wang, X. Shao, M. Reisslein, and H. El Bakoury, “Ultra-Low Latency (ULL) networks:  
742 The IEEE TSN and IETF DetNet standards and related 5G ULL research,” *IEEE Commun. Surv. & Tut.*, vol. 21, no. 1, First Qu. 2019.
- 743 [7] Y. Yang and A. M. Zador, “Differences in sensitivity to neural timing among cortical areas,” *Journal of Neuroscience*, vol. 32, no. 43,  
744 pp. 15 142–15 147, Oct. 2012.
- 745 [8] L. Zhang, H. Gao, and O. Kaynak, “Network-induced constraints in networked control systems—a survey,” *IEEE Transactions on*  
746 *Industrial Informatics*, vol. 9, no. 1, pp. 403–416, Feb. 2013.
- 747 [9] X.-M. Zhang, Q.-L. Han, and X. Yu, “Survey on recent advances in networked control systems,” *IEEE Transactions on Industrial*  
748 *Informatics*, vol. 12, no. 5, pp. 1740–1752, Oct. 2016.
- 749 [10] *Technical Specification Group Services and System Aspects; Study on Communication for Automation in Vertical Domains (Release 16)*,  
750 3GPP, 05 2018, 22.804 TR, V2.0.0.
- 751 [11] Q.-Y. Zhang, X.-W. Wang, M. Huang, K.-Q. Li, and S. K. Das, “Software defined networking meets information centric networking: A  
752 survey,” *IEEE Access*, vol. 6, pp. 39 547–39 563, 2018.

- 753 [12] I. Afolabi, T. Taleb, K. Samdanis, A. Ksentini, and H. Flinck, "Network slicing & softwarization: A survey on principles, enabling  
754 technologies & solutions," *IEEE Commun. Surv. & Tut.*, vol. 20, no. 3, pp. 2429 – 2453, Third Qu. 2018.
- 755 [13] M. Bagaa, T. Taleb, A. Laghrissi, A. Ksentini, and H. Flinck, "Coalitional game for the creation of efficient virtual core network slices  
756 in 5G mobile systems," *IEEE J. on Selected Areas in Commun.*, vol. 36, no. 3, pp. 469–484, Mar. 2018.
- 757 [14] S. Fu, J. Liu, and W. Zhu, "Multimedia content delivery with network function virtualization: The energy perspective," *IEEE MultiMedia*,  
758 no. 3, pp. 38–47, Jul.–Sep. 2017.
- 759 [15] Y. Harchol, D. Hay, and T. Orenstein, "FTvNF: Fault tolerant virtual network functions," in *Proc. ACM Symp. on Architectures for Netw.  
760 and Commun. Sys.*, 2018, pp. 141–147.
- 761 [16] W. Kellerer, P. Kalmbach, A. Blenk, A. Basta, M. Reisslein, and S. Schmid, "Adaptable and data-driven softwarized networks: Review,  
762 opportunities, and challenges," *Proc. IEEE, in print*, 2019.
- 763 [17] L. Linguaglossa, D. Rossi, S. Pontarelli, D. Barach, D. Marjon, and P. Pfister, "High-speed data plane and network functions virtualization  
764 by vectorizing packet processing," *Computer Networks*, vol. 149, pp. 187–199, Feb. 2019.
- 765 [18] I. Trajkovska, M.-A. Kourtis, C. Sakkas, D. Baudinot, J. Silva, P. Harsh, G. Xylouris, T. M. Bohnert, and H. Koumaras, "SDN-based  
766 service function chaining mechanism and service prototype implementation in NFV scenario," *Computer Standards & Interfaces*, vol. 54,  
767 pp. 247–265, Nov. 2017.
- 768 [19] B. Yi, X. Wang, K. Li, and M. Huang, "A comprehensive survey of network function virtualization," *Computer Networks*, vol. 133, pp.  
769 212–262, Mar. 2018.
- 770 [20] A. AbdelSalam, F. Clad, C. Filsfils, S. Salsano, G. Siracusano, and L. Veltri, "Implementation of virtual network function chaining  
771 through segment routing in a Linux-based NFV infrastructure," in *Proc. IEEE Conf. on Network Softwarization (NetSoft)*, 2017, pp. 1–5.
- 772 [21] L. Askari, A. Hmaity, F. Musumeci, and M. Tornatore, "Virtual-network-function placement for dynamic service chaining in metro-area  
773 networks," in *Proc. IEEE Int. Conf. on Optical Network Design and Modeling (ONDM)*, 2018, pp. 136–141.
- 774 [22] N. F. S. de Sousa, D. A. L. Perez, R. V. Rosa, M. A. Santos, and C. E. Rothenberg, "Network service orchestration: A survey," *arXiv  
775 preprint arXiv:1803.06596*, 2018.
- 776 [23] A. Gupta, M. F. Habib, U. Mandal, P. Chowdhury, M. Tornatore, and B. Mukherjee, "On service-chaining strategies using virtual network  
777 functions in operator networks," *Computer Networks*, vol. 133, pp. 1–16, Mar. 2018.
- 778 [24] H. Hantouti, N. Benamar, T. Taleb, and A. Laghrissi, "Traffic steering for service function chaining," *IEEE Commun. Surv. & Tut., in  
779 print*, 2019.
- 780 [25] W. Hahn, B. Gajic, F. Wohlfart, D. Raumer, P. Emmerich, S. Gallenmueller, and G. Carle, "Feasibility of compound chained network  
781 functions for flexible packet processing," in *Proc. European Wireless Conf.*, May 2017, pp. 1–6.
- 782 [26] K. Han, S. Li, S. Tang, H. Huang, S. Zhao, G. Fu, and Z. Zhu, "Application-driven end-to-end slicing: When wireless network  
783 virtualization orchestrates with NFV-based mobile edge computing," *IEEE Access*, vol. 6, pp. 26 567 – 26 577, 2018.
- 784 [27] A. Hmaity, M. Savi, F. Musumeci, M. Tornatore, and A. Pattavina, "Protection strategies for virtual network functions placement and  
785 service chains provisioning," *Networks*, vol. 70, no. 4, pp. 373–387, Dec. 2017.
- 786 [28] T.-W. Kuo, B.-H. Liou, K. C.-J. Lin, and M.-J. Tsai, "Deploying chains of virtual network functions: On the relation between link and  
787 server usage," *IEEE/ACM Transactions on Networking*, vol. 26, no. 4, pp. 1562–1576, Aug. 2018.
- 788 [29] A. M. Medhat, T. Taleb, A. Elmangoush, G. A. Carella, S. Covaci, and T. Magedanz, "Service function chaining in next generation  
789 networks: State of the art and research challenges," *IEEE Commun. Mag.*, vol. 55, no. 2, pp. 216–223, Feb. 2017.
- 790 [30] F. Rath, J. Krude, J. R uth, D. Schemmel, O. Hohlfeld, J.  . Bitsch, and K. Wehrle, "SymPerf: Predicting network function performance,"  
791 in *Proc. of the SIGCOMM Posters and Demos*, 2017, pp. 34–36.
- 792 [31] D. Raumer, S. Bauer, P. Emmerich, and G. Carle, "Performance implications for intra-node placement of network function chains," in  
793 *Proc. IEEE Int. Conf. on Cloud Networking (CloudNet)*, 2017, pp. 1–6.
- 794 [32] A. Morton, "Considerations for benchmarking virtual network functions and their infrastructure," RFC 8172, Jul. 2007. [Online].  
795 Available: <https://rfc-editor.org/rfc/rfc8172.txt>

- 796 [33] B. Yang, Z. Xu, W. K. Chai, W. Liang, D. Tuncer, A. Galis, and G. Pavlou, "Algorithms for fault-tolerant placement of stateful virtualized  
797 network functions," in *Proc. IEEE Int. Conf. on Commun. (ICC)*, 2018, pp. 20–24.
- 798 [34] B. Yi, X. Wang, and M. Huang, "A dynamic heuristic for the recomposition of service function chain," *IET Communications*, vol. 12,  
799 no. 16, pp. 1984–1990, Oct. 2018.
- 800 [35] E. Ahmed and M. H. Rehmani, "Mobile edge computing: Opportunities, solutions, and challenges," *Future Generation Computer Systems*,  
801 vol. 70, pp. 59–63, May 2017.
- 802 [36] M. Chen, Y. Zhang, L. Hu, T. Taleb, and Z. Sheng, "Cloud-based wireless network: Virtualized, reconfigurable, smart wireless network  
803 to enable 5G technologies," *Mobile Networks and Applications*, vol. 20, no. 6, pp. 704–712, Dec. 2015.
- 804 [37] T. Taleb and A. Ksentini, "Follow me cloud: Interworking federated clouds and distributed mobile networks," *IEEE Network*, vol. 27,  
805 no. 5, pp. 12–19, Sep.-Oct. 2013.
- 806 [38] T. Taleb, M. Corici, C. Parada, A. Jamakovic, S. Ruffino, G. Karagiannis, and T. Magedanz, "EASE: EPC as a service to ease mobile  
807 core network deployment over cloud," *IEEE Network*, vol. 29, no. 2, pp. 78–88, Mar.-Apr. 2015.
- 808 [39] T. Taleb, K. Samdanis, B. Mada, H. Flinck, S. Dutta, and D. Sabella, "On multi-access edge computing: A survey of the emerging 5G  
809 network edge cloud architecture and orchestration," *IEEE Commun. Surv. & Tut.*, vol. 19, no. 3, pp. 1657–1681, Third Qu. 2017.
- 810 [40] H. Zhang, N. Liu, X. Chu, K. Long, A.-H. Aghvami, and V. C. Leung, "Network slicing based 5G and future mobile networks: Mobility,  
811 resource management, and challenges," *IEEE Communications Magazine*, vol. 55, no. 8, pp. 138–145, Aug. 2017.
- 812 [41] S. Cheshire, "It is the latency, stupid," 1996, available from <http://www.stuartcheshire.org/rants/latency.html>, Last accessed Jan. 16, 2019.
- 813 [42] —, "Latency and the quest for interactivity," in *White paper commissioned by Volpe Welty Asset Management,*  
814 *LLC, for the Synchronous Person-to-Person Interactive Computing Environments Meeting*, 1996, pp. 1–8, available from  
815 <http://www.stuartcheshire.org/papers/LatencyQuest.pdf>, Last accessed Jan. 16, 2019.
- 816 [43] O. S. Sella, A. W. Moore, and N. Zilberman, "FEC killed the cut-through switch," in *Proc. ACM Workshop on Networking for Emerging*  
817 *Appl. and Techn.*, 2018, pp. 15–20.
- 818 [44] N. Zilberman, M. Grosvenor, D. A. Popescu, N. Manihatty-Bojan, G. Antichi, M. Wójcik, and A. W. Moore, "Where has my time  
819 gone?" in *Proc. Int. Conf. on Passive and Active Network Measurement*, 2017, pp. 201–214.
- 820 [45] N. Hanford, V. Ahuja, M. K. Farrens, B. Tierney, and D. Ghosal, "A survey of end-system optimizations for high-speed networks," *ACM*  
821 *Computing Surveys (CSUR)*, vol. 51, no. 3, pp. 54:1–54:36, Jul. 2018.
- 822 [46] Z. Xiang, F. Gabriel, G. T. Nguyen, and F. H. P. Fitzek, "Latency measurement of service function chaining on OpenStack platform,"  
823 in *Proc. IEEE Conf. on Local Computer Networks (LCN)*, Chicago, IL, 2018.
- 824 [47] M. He, A. M. Alba, A. Basta, A. Blenk, and W. Kellerer, "Flexibility in softwarized networks: Classifications and research challenges,"  
825 *IEEE Commun. Surv. & Tut.*, in print, 2019.
- 826 [48] P. Emmerich, D. Raumer, S. Gallenmüller, F. Wohlfart, and G. Carle, "Throughput and latency of virtual switching with OpenvSwitch:  
827 A quantitative analysis," *Journal of Network and Systems Management*, vol. 26, no. 2, pp. 314–338, Apr. 2018.
- 828 [49] C.-L. Hsieh and N. Weng, "NF-switch: VNFs-enabled SDN switches for high performance service function chaining," in *Proc. IEEE*  
829 *Int. Conf. on Network Protocols (ICNP)*, 2017, pp. 1–6.
- 830 [50] G. Lettieri, V. Maffione, and L. Rizzo, "A survey of fast packet I/O technologies for network function virtualization," in *Proc. Int. Conf.*  
831 *on High Performance Computing*. Springer, Cham, Switzerland, 2017, pp. 579–590.
- 832 [51] L. Zhang, S. Lai, C. Wu, Z. Li, and C. Guo, "Virtualized network coding functions on the internet," in *Proc. IEEE Int. Conf. on Distr.*  
833 *Computing Systems*, Jun. 2017, pp. 129–139.
- 834 [52] SFC-Ostack: A Simple Research Framework for SFC on OpenStack. [Accessed 2019-1-15]. [Online]. Available:  
835 <https://github.com/stevetorenz/sfc-ostack>
- 836 [53] Build-VNF project repository. [Accessed 2018-10-5]. [Online]. Available: <https://github.com/stevetorenz/build-vnf/tree/master/CALVIN>
- 837 [54] Open vSwitch with DPDK. [Accessed 2018-09-15]. [Online]. Available: <http://docs.openvswitch.org/en/latest/intro/install/dpdk/>
- 838 [55] M.-A. Kourtis, M. J. McGrath, G. Gardikis, G. Xilouris, V. Riccobene, P. Papadimitriou, E. Trouva, F. Liberati, M. Trubian, J. Batallé

- 839 *et al.*, “T-NOVA: An open-source MANO stack for NFV infrastructures,” *IEEE Trans. on Network and Service Management*, vol. 14,  
840 no. 3, pp. 586–602, Sep. 2017.
- 841 [56] H. A. Alameddine, S. Sharafeddine, S. Sebbah, S. Ayoubi, and C. Assi, “Dynamic task offloading and scheduling for low-latency IoT  
842 services in multi-access edge computing,” *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 3, pp. 668–682, Mar. 2019.
- 843 [57] D. Cho, J. Taheri, A. Y. Zomaya, and L. Wang, “Virtual network function placement: Towards minimizing network latency and lead  
844 time,” in *Proc. IEEE Int. Conf. on Cloud Computing Techn. and Science (CloudCom)*, 2017, pp. 90–97.
- 845 [58] R. Gouareb, V. Friderikos, and A. H. Aghvami, “Delay sensitive virtual network function placement and routing,” in *Proc. IEEE Int.  
846 Conf. on Telecommun. (ICT)*, 2018, pp. 394–398.
- 847 [59] H. Halabian, “Distributed resource allocation optimization in 5G virtualized networks,” *IEEE Journal on Selected Areas in Communi-  
848 cations*, vol. 37, no. 3, pp. 627–642, Mar. 2019.
- 849 [60] H. Hawilo, M. Jammal, and A. Shami, “Network function virtualization-aware orchestrator for service function chaining placement in  
850 the cloud,” *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 3, pp. 643–655, Mar. 2019.
- 851 [61] D. Liao, Y. Wu, Z. Wu, Z. Zhu, W. Zhang, G. Sun, and V. Chang, “AI-based software-defined virtual network function scheduling with  
852 delay optimization,” *Cluster Computing*, in print, pp. 1–13, 2019.
- 853 [62] W. Miao, G. Min, Y. Wu, H. Huang, Z. Zhao, H. Wang, and C. Luo, “Stochastic performance analysis of network function virtualization  
854 in future internet,” *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 3, pp. 613–626, Mar. 2019.
- 855 [63] L. Qu, C. Assi, and K. Shaban, “Delay-aware scheduling and resource optimization with network function virtualization,” *IEEE Trans.  
856 on Commun.*, vol. 64, no. 9, pp. 3746–3758, Sep. 2016.
- 857 [64] L. Qu, C. Assi, K. Shaban, and M. Khabbaz, “A reliability-aware network service chain provisioning with delay guarantees in NFV-  
858 enabled enterprise datacenter networks,” *IEEE Trans. on Netw. and Service Managm.*, vol. 14, no. 3, pp. 554–568, Sep. 2017.
- 859 [65] M. Savi, M. Tornatore, and G. Verticale, “Impact of processing-resource sharing on the placement of chained virtual network functions,”  
860 *arXiv preprint arXiv:1710.08262*, 2017.
- 861 [66] L. Tang, H. Yang, R. Ma, L. Hu, W. Wang, and Q. Chen, “Queue-aware dynamic placement of virtual network functions in 5G access  
862 network,” *IEEE Access*, vol. 6, pp. 44 291–44 305, 2018.
- 863 [67] Q. Xu, J. Wang, and K. Wu, “Resource capacity analysis in network slicing with ensured end-to-end performance bound,” in *Proc. IEEE  
864 Int. Conf. on Commun. (ICC)*, 2018, pp. 1–6.
- 865 [68] Q. Ye, W. Zhuang, X. Li, and J. Rao, “End-to-end delay modeling for embedded VNF chains in 5G core networks,” *IEEE Internet of  
866 Things Journal*, in print, 2019.
- 867 [69] S. Agarwal, F. Malandrino, C. F. Chiasserini, and S. De, “VNF placement and resource allocation for the support of vertical services in  
868 5G networks,” *IEEE/ACM Trans. on Netw.*, in print, 2019.
- 869 [70] I. Benkacem, T. Taleb, M. Bagaa, and H. Flinck, “Optimal VNFs placement in CDN slicing over multi-cloud environment,” *IEEE J. on  
870 Sel. Areas in Commun.*, vol. 36, no. 3, pp. 616–627, Mar. 2018.
- 871 [71] J. Cao, Y. Zhang, W. An, X. Chen, Y. Han, and J. Sun, “VNF placement in hybrid NFV environment: Modeling and genetic algorithms,”  
872 in *Proc. IEEE Int. Conf. on Parallel and Distr. Sys. (ICPADS)*, 2016, pp. 769–777.
- 873 [72] X. Chen, W. Ni, I. B. Collings, X. Wang, and S. Xu, “Distributed placement and online optimization of virtual machines for network  
874 service chains,” in *Proc. IEEE Int. Conf. on Commun. (ICC)*, 2018, pp. 1–6.
- 875 [73] C. Galdamez, R. Pamula, and Z. Ye, “On efficient virtual network function chaining in NFV-based telecommunications networks,”  
876 *Cluster Computing*, in print, pp. 1–11, 2019.
- 877 [74] A. Laghrissi, T. Taleb, M. Bagaa, and H. Flinck, “Towards edge slicing: VNF placement algorithms for a dynamic & realistic edge  
878 cloud environment,” in *Proc. IEEE GLOBECOM*, 2017, pp. 1–6.
- 879 [75] A. Laghrissi, T. Taleb, and M. Bagaa, “Conformal mapping for optimal network slice planning based on canonical domains,” *IEEE J.  
880 on Selected Areas in Commun.*, vol. 36, no. 3, pp. 519–528, 2018.

- 881 [76] A. Laghrissi and T. Taleb, "A survey on the placement of virtual resources and virtual network functions," *IEEE Commun. Surv. & Tut.*,  
882 *in print*, 2019.
- 883 [77] M. A. T. Nejad, S. Parsaeefard, M. A. Maddah-Ali, T. Mahmoodi, and B. H. Khalaj, "vSPACE: VNF simultaneous placement, admission  
884 control and embedding," *IEEE J. on Sel. Areas in Commun.*, vol. 36, no. 3, pp. 542–557, Mar. 2018.
- 885 [78] Y. Hu and T. Li, "Towards efficient server architecture for virtualized network function deployment: Implications and implementations,"  
886 in *Proc. IEEE/ACM Int. Symp. on Microarch.*, 2016, pp. 1–8.
- 887 [79] G. Durisi, T. Koch, J. Östman, Y. Polyanskiy, and W. Yang, "Short-packet communications over multiple-antenna rayleigh-fading  
888 channels," *IEEE Trans. on Commun.*, vol. 64, no. 2, pp. 618–629, Feb. 2016.
- 889 [80] G. Durisi, T. Koch, and P. Popovski, "Toward massive, ultrareliable, and low-latency wireless communication with short packets," *Proc.*  
890 *IEEE*, vol. 104, no. 9, pp. 1711–1726, Sep. 2016.
- 891 [81] T. Herbert and A. Starovoirov, "eXpress Data Path (XDP): Programmable and high performance networking data path," Mar. 2016,  
892 available from [https://github.com/iovisor/bpf-docs/blob/master/Express\\_Data\\_Path.pdf](https://github.com/iovisor/bpf-docs/blob/master/Express_Data_Path.pdf), Last accessed Oct. 2, 2018.
- 893 [82] J. D. Brouer, "Linux Networking Subsystem, XDP – eXpress Data Path," 2016, available from [https://prototype-](https://prototype-kernel.readthedocs.io/en/latest/networking/index.html)  
894 [kernel.readthedocs.io/en/latest/networking/index.html](https://prototype-kernel.readthedocs.io/en/latest/networking/index.html), Last accessed Oct. 2, 2018.
- 895 [83] "BPF and XDP Reference Guide," 2018, available from <https://cilium.readthedocs.io/en/v1.2/bpf/>, Last accessed Oct. 2, 2018.
- 896 [84] T. Høiland-Jørgensen, J. D. Brouer, D. Borkmann, J. Fastabend, T. Herbert, D. Ahern, and D. Miller, "The eXpress Data Path: Fast  
897 programmable packet processing in the operating system kernel," in *Proc. ACM Int. Conf. on Emerging Netw. Experiments and Techn.*  
898 *(CoNEXT)*, 2018, pp. 54–66.
- 899 [85] S. Miano, M. Bertrone, F. Risso, and M. Tumolo, "Creating complex network services with eBPF: Experience and lessons learned," in  
900 *Proc. IEEE High Performance Switching and Routing (HPSR)*, 2018, pp. 1–8.
- 901 [86] Z. Ahmed, M. H. Alizai, and A. A. Syed, "InKeV: In-kernel distributed network virtualization for DCN," *ACM SIGCOMM Computer*  
902 *Commun. Rev.*, vol. 46, no. 3, p. 4, Jul. 2018.
- 903 [87] "Neutron documentation," Sep. 2018, available from <https://docs.openstack.org/neutron>, Last accessed Oct. 3, 2018.
- 904 [88] M. Bertrone, S. Miano, F. Risso, and M. Tumolo, "Accelerating Linux security with eBPF iptables," in *Proc. ACM SIGCOMM Conf.*  
905 *on Posters and Demos*, 2018, pp. 108–110.
- 906 [89] X. Li, Y. Wu, J. Ge, H. Zheng, E. Yuepeng, C. Han, and H. Lv, "A kernel-space POF virtual switch," *Computers & Electrical Engineering*,  
907 vol. 61, pp. 339–350, Jul. 2017.
- 908 [90] D. Scholz, D. Raumer, P. Emmerich, A. Kurtz, K. Lesiak, and G. Carle, "Performance implications of packet filtering with Linux eBPF,"  
909 in *Proc. Int. Teletraffic Congress (ITC)*, Sep. 2018, pp. 1–9.
- 910 [91] Y. Zhang, B. Anwer, V. Gopalakrishnan, B. Han, J. Reich, A. Shaikh, and Z.-L. Zhang, "Parabox: Exploiting parallelism for virtual  
911 network functions in service chaining," in *Proc. ACM Symp. on SDN Research*, 2017, pp. 143–149.
- 912 [92] J. L. García-Dorado, F. Mata, J. Ramos, P. M. S. del Río, V. Moreno, and J. Aracil, "High-performance network traffic processing  
913 systems using commodity hardware," in *Data Traffic Monitoring and Analysis, LNCS 7754*. Springer, Berlin, 2013, pp. 3–27.
- 914 [93] S. Gallenmüller, D. Scholz, F. Wohlfart, Q. Scheitle, P. Emmerich, and G. Carle, "High-performance packet processing and measurements,"  
915 in *Proc. IEEE Int. Conf. on Commun. Systems & Networks (COMSNETS)*, 2018, pp. 1–8.
- 916 [94] S. Gallenmüller, P. Emmerich, F. Wohlfart, D. Raumer, and G. Carle, "Comparison of frameworks for high-performance packet IO," in  
917 *Proc. ACM/IEEE Symp. on Architectures for Netw. and Commun. Systems*, 2015, pp. 29–38.
- 918 [95] A. Anthony, S. R. Chowdhury, T. Bai, R. Boutaba, and J. François, "UNiS: A user-space non-intrusive workflow-aware virtual network  
919 function scheduler," in *Proc. IEEE Int. Conf. on Netw. and Service Management (CNSM)*, Nov. 2018, pp. 152–160.
- 920 [96] J. Duan, X. Yi, J. Wang, C. Wu, and F. Le, "NetStar: A future/promise framework for asynchronous network functions," *IEEE Journal*  
921 *on Selected Areas in Communications*, vol. 37, no. 3, pp. 600–612, Mar. 2019.
- 922 [97] C. Zhang, J. Bi, Y. Zhou, and J. Wu, "HyperVDP: High-performance virtualization of the programmable data plane," *IEEE Journal on*  
923 *Selected Areas in Communications*, vol. 37, no. 3, pp. 556–569, Mar. 2019.

- 924 [98] J. Duan, X. Yi, S. Zhao, C. Wu, H. Cui, and F. Le, "NFVactor: A resilient NFV system using the distributed actor model," *IEEE Journal*  
925 *on Selected Areas in Communications*, vol. 37, no. 3, pp. 586–599, Mar. 2019.
- 926 [99] M. Zhang, J. Bi, K. Gao, Y. Qiao, G. Li, X. Kong, Z. Li, and H. Hu, "Tripod: Towards a scalable, efficient and resilient cloud gateway,"  
927 *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 3, pp. 570–585, Mar. 2019.
- 928 [100] N. Van Tu, K. Ko, and J. W.-K. Hong, "Architecture for building hybrid kernel-user space virtual network functions," in *Proc. Int. Conf.*  
929 *on Network and Service Management (CNSM)*, 2017, pp. 1–6.
- 930 [101] P. Emmerich, M. Pudelko, S. Bauer, and G. Carle, "User space network drivers," in *Proc. ACM Applied Netw. Res. Workshop*, 2018,  
931 pp. 91–93.
- 932 [102] "Kernel NIC Interface," 2018, available from [https://doc.dpdk.org/guides/prog\\_guide/kernel\\_nic\\_interface.html](https://doc.dpdk.org/guides/prog_guide/kernel_nic_interface.html), Last accessed Oct. 3,  
933 2018.
- 934 [103] X. Wang, C. Xu, G. Zhao, and S. Yu, "Tuna: an efficient and practical scheme for wireless access point in 5G networks virtualization,"  
935 *IEEE Commun. Letters*, vol. 22, no. 4, pp. 748–751, Apr. 2018.
- 936 [104] C. Li, C. Ding, and K. Shen, "Quantifying the cost of context switch," in *Proc. USENIX Workshop on Experimental Computer Science*,  
937 2007, pp. 1–4.
- 938 [105] M. V. Pedersen, J. Heide, and F. H. Fitzek, "Kodo: An open and research oriented network coding library," in *Proc. Int. Conf. on*  
939 *Research in Networking, LNCS, Vol. 6827*. Springer, Berlin, Heidelberg, 2011, pp. 145–152.
- 940 [106] D. Cerovic, V. Del Piccolo, A. Amamou, K. Haddadou, and G. Pujolle, "Fast packet processing: A survey," *IEEE Commun. Surv. and*  
941 *Tut.*, vol. 20, no. 4, pp. 3645–3676, Fourth Qu. 2018.
- 942 [107] D. Barach, L. Linguaglossa, D. Marion, P. Pfister, S. Pontarelli, and D. Rossi, "High-speed software data plane via vectorized packet  
943 processing," *IEEE Communications Magazine*, vol. 56, no. 12, pp. 97–103, Dec. 2018.
- 944 [108] IOVisor Homepage. [Accessed 2019-1-24]. [Online]. Available: <https://www.iovisor.org/>
- 945 [109] M.-A. Kourtis, G. Xilouris, V. Riccobene, M. J. McGrath, G. Petralia, H. Koumaras, G. Gardikis, and F. Liberal, "Enhancing VNF  
946 performance by exploiting SR-IOV and DPDK packet processing acceleration," in *Proc. IEEE Conf. Netw. Function Virt. and Software*  
947 *Defined Network (NFV-SDN)*, 2015, pp. 74–78.
- 948 [110] OpenStack Nova computing documentation. [Accessed 2018-10-3]. [Online]. Available:  
949 <https://www.spinics.net/lists/netdev/msg405175.html>
- 950 [111] Linux netdev mailing list: Patch 4/5. [Accessed 2018-10-3]. [Online]. Available: <https://www.spinics.net/lists/netdev/msg405175.html>
- 951 [112] BCC project homepage. [Accessed 2018-09-20]. [Online]. Available: <https://github.com/iovisor/bcc>
- 952 [113] J. Hwang, K. K. Ramakrishnan, and T. Wood, "NetVM: high performance and flexible networking using virtualization on commodity  
953 platforms," *IEEE Trans. on Netw. and Service Manag.*, vol. 12, no. 1, pp. 34–47, Mar. 2015.
- 954 [114] W. Zhang, G. Liu, W. Zhang, N. Shah, P. Lopreiato, G. Todeschi, K. Ramakrishnan, and T. Wood, "OpenNetVM: A platform for high  
955 performance network service chains," in *Proc. ACM Workshop on Hot Topics in Middleboxes and Network Function Virtualization*, 2016,  
956 pp. 26–31.
- 957 [115] "VMware Docs, Enable Live Resize," Nov. 2018, [https://docs.vmware.com/en/VMware-Integrated-](https://docs.vmware.com/en/VMware-Integrated-OpenStack/4.1/com.vmware.openstack.admin.doc/GUID-FEB48287-04AD-4BAB-9B42-99543DCC9733.html)  
958 [OpenStack/4.1/com.vmware.openstack.admin.doc/GUID-FEB48287-04AD-4BAB-9B42-99543DCC9733.html](https://docs.vmware.com/en/VMware-Integrated-OpenStack/4.1/com.vmware.openstack.admin.doc/GUID-FEB48287-04AD-4BAB-9B42-99543DCC9733.html), Last accessed Mar.  
959 4, 2019.
- 960 [116] "OpenStack Compute (nova), instance live resize," 2018, <https://blueprints.launchpad.net/nova/+spec/instance-live-resize>, Last accessed  
961 Jan. 5, 2019.
- 962 [117] P. Veitch and T. Long, "A Low-Latency NFV Infrastructure for Performance-Critical Applications," Intel Corp., Tech.  
963 Rep., 2017, [Accessed 2018-9-30]. [Online]. Available: [https://software.intel.com/en-us/articles/low-latency-nfv-infrastructure-for-](https://software.intel.com/en-us/articles/low-latency-nfv-infrastructure-for-performance-critical-applications)  
964 [performance-critical-applications](https://software.intel.com/en-us/articles/low-latency-nfv-infrastructure-for-performance-critical-applications)
- 965 [118] OpenStack Pike documentation. [Accessed 2018-10-4]. [Online]. Available: <https://docs.openstack.org/pike/>
- 966 [119] Virtio Documentation. [Accessed 2019-1-20]. [Online]. Available: <https://www.linux-kvm.org/page/Virtio>

- 967 [120] J. Chauhan, D. Makaroff, and A. Arkles, "Is doing clock synchronization in a VM a good idea?" in *Proc. IEEE Int. Performance*  
968 *Computing and Commun. Conf.*, 2010, pp. 1–2.
- 969 [121] Iperf homepage. [Accessed 2018-9-10]. [Online]. Available: <https://iperf.fr/>
- 970 [122] Service Function Chain Extension for OpenStack Networking. [Accessed 2018-10-1]. [Online]. Available:  
971 <https://docs.openstack.org/networking-sfc/latest/>
- 972 [123] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, "The Click modular router," *ACM Transactions on Computer Systems*  
973 *(TOCS)*, vol. 18, no. 3, pp. 263–297, Aug. 2000.
- 974 [124] BPF and XDP reference guide. [Accessed 2018-10-4]. [Online]. Available: <https://cilium.readthedocs.io/en/v1.2/bpf/>
- 975 [125] Linux Kernel networking documentation. [Accessed 2018-10-4]. [Online]. Available:  
976 <https://www.kernel.org/doc/Documentation/networking/ip-sysctl.txt>
- 977 [126] C. Sieber, R. Durner, M. Ehm, W. Kellerer, and P. Sharma, "Towards optimal adaptation of NFV packet processing to modern CPU  
978 memory architectures," in *Proc. ACM Workshop on Cloud-Assisted Netw.*, 2017, pp. 7–12.
- 979 [127] Y. Hu, M. Song, and T. Li, "Towards full containerization in containerized network function virtualization," *ACM SIGARCH Computer*  
980 *Architecture News*, vol. 45, no. 1, pp. 467–481, Mar. 2017.
- 981 [128] DPDK Poll Mode Driver. [Accessed 2019-1-20]. [Online]. Available: [https://doc.dpdk.org/guides/prog\\_guide/poll\\_mode\\_drv.html](https://doc.dpdk.org/guides/prog_guide/poll_mode_drv.html)
- 982 [129] D. Vladislavić, D. Huljenić, and J. Ožegović, "Enhancing VNF's performance using DPDK driven OVS user-space forwarding," in *Proc.*  
983 *IEEE Int. Conf. on Softw., Telecommun. and Computer Netw. (SoftCOM)*, 2017, pp. 1–5.
- 984 [130] DPDK Mbuf Library. [Accessed 2019-1-20]. [Online]. Available: [https://doc.dpdk.org/guides/prog\\_guide/mbuf\\_lib.html](https://doc.dpdk.org/guides/prog_guide/mbuf_lib.html)
- 985 [131] L. A. D. Knob, B. G. Xavier, and T. Ferreto, "An unikernels provisioning architecture for OpenStack," in *Proc. IEEE Symp. on Computers*  
986 *and Commun. (ISCC)*, 2018, pp. 903–908.
- 987 [132] P. L. Ventre, P. Lungaroni, G. Siracusano, C. Pisa, F. Schmidt, F. Lombardo, and S. Salsano, "On the fly orchestration of unikernels:  
988 Tuning and performance evaluation of virtual infrastructure managers," *IEEE Transactions on Cloud Computing, in print*, 2019.
- 989 [133] Mpstat Manpage. [Accessed 2019-1-24]. [Online]. Available: <http://man7.org/linux/man-pages/man1/mpstat.1.html>
- 990 [134] H. D. Chiramal, P. Mukhedkar, and A. Vettathu, *Mastering KVM Virtualization*. Packt Publishing Ltd, 2016.
- 991 [135] R. Ahlswede, N. Cai, S.-Y. Li, and R. W. Yeung, "Network information flow," *IEEE Transactions on Information Theory*, vol. 46, no. 4,  
992 pp. 1204–1216, Jul. 2000.
- 993 [136] T. Ho, M. Médard, R. Koetter, D. R. Karger, M. Effros, J. Shi, and B. Leong, "A random linear network coding approach to multicast,"  
994 *IEEE Transactions on Information Theory*, vol. 52, no. 10, pp. 4413–4430, Oct. 2006.
- 995 [137] S. Feizi, D. E. Lucani, and M. Médard, "Tunable sparse network coding," in *Proc. Int. Zurich Seminar on Commun. (IZS)*. Eidgenössische  
996 Technische Hochschule Zürich, 2012.
- 997 [138] V. Nguyen, G. T. Nguyen, F. Gabriel, D. E. Lucani, and F. H. Fitzek, "Integrating sparsity into Fulcrum codes: Investigating throughput,  
998 complexity and overhead," in *Proc. IEEE Int. Conf. on Commun. Workshops (ICC Workshops)*, 2018, pp. 1–6.
- 999 [139] F. Gabriel, S. Wunderlich, S. Pandi, F. H. Fitzek, and M. Reisslein, "Caterpillar RLNC with feedback (CRLNC-FB): Reducing delay in  
1000 selective repeat ARQ through coding," *IEEE Access*, vol. 6, pp. 44 787–44 802, 2018.
- 1001 [140] S. Wunderlich, F. Gabriel, S. Pandi, F. H. Fitzek, and M. Reisslein, "Caterpillar RLNC (CRLNC): A practical finite sliding window  
1002 RLNC approach," *IEEE Access*, vol. 5, pp. 20 183–20 197, 2017.
- 1003 [141] D. E. Lucani, M. V. Pedersen, D. Ruano, C. W. Sørensen, F. H. Fitzek, J. Heide, O. Geil, V. Nguyen, and M. Reisslein, "Fulcrum:  
1004 Flexible network coding for heterogeneous devices," *IEEE Access*, vol. 6, pp. 77 890–77 910, 2018.
- 1005 [142] A. G. Dimakis, P. B. Godfrey, Y. Wu, M. J. Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," *IEEE*  
1006 *Transactions on Information Theory*, vol. 56, no. 9, pp. 4539–4551, Sep. 2010.
- 1007 [143] P. A. Chou, Y. Wu, and K. Jain, "Practical network coding," in *Proc. Allerton Conference on Communication, Control, and Computing*,  
1008 vol. 41, 2003, pp. 40–49.

- 1009 [144] J. Barros, R. A. Costa, D. Munaretto, and J. Widmer, "Effective delay control in online network coding," in *Proc. IEEE INFOCOM*,  
1010 2009, pp. 208–216.
- 1011 [145] S. Pandi, F. Gabriel, J. A. Cabrera, S. Wunderlich, M. Reisslein, and F. H. Fitzek, "PACE: Redundancy engineering in RLNC for  
1012 low-latency communication," *IEEE Access*, vol. 5, pp. 20477–20493, 2017.
- 1013 [146] M. Karzand and D. J. Leith, "Low delay random linear coding over a stream," in *Proc. Allerton Conf. on Communication, Control, and  
1014 Computing*, 2014, pp. 521–528.
- 1015 [147] D. Szabo, A. Gulyas, F. H. Fitzek, and D. E. Lucani, "Towards the tactile internet: Decreasing communication latency with network  
1016 coding and software defined networking," in *Proc. European Wireless*, 2015, pp. 1–6.
- 1017 [148] F. Gabriel, G. T. Nguyen, R.-S. Schmoll, J. A. Cabrera, M. Muehleisen, and F. H. Fitzek, "Practical deployment of network coding for  
1018 real-time applications in 5G networks," in *Proc. IEEE Consumer Commun. & Netw. Conf. (CCNC)*, 2018, pp. 1–2.
- 1019 [149] A. P. Felt, R. Barnes, A. King, C. Palmer, C. Bentzel, and P. Tabriz, "Measuring HTTPS adoption on the web," in *Proc. USENIX  
1020 Security Symposium*, 2017, pp. 1323–1338.
- 1021 [150] V.-C. Nguyen, A.-V. Vu, K. Sun, and Y. Kim, "An experimental study of security for service function chaining," in *Proc. IEEE Int.  
1022 Conf. on Ubiquitous and Future Networks (ICUFN)*, 2017, pp. 797–799.
- 1023 [151] Z. Zheng, J. Bi, C. Sun, H. Yu, H. Hu, Z. Meng, S. Wang, K. Gao, and J. Wu, "GEN: A GPU-accelerated elastic framework for NFV,"  
1024 in *Proc. ACM Asia-Pacific Workshop on Networking*, 2018, pp. 57–64.
- 1025 [152] M. M. Rovnyagin and A. A. Kuznetsov, "Application of hybrid computing technologies for high-performance distributed NFV systems,"  
1026 in *Proc. IEEE Conf. of Russian Young Researchers in Electrical and Electronic Engineering (EIConRus)*, 2017, pp. 540–543.
- 1027 [153] Small portable AES implementation in C. [Accessed 2018-9-10]. [Online]. Available: <https://github.com/kokke/tiny-AES-c>
- 1028 [154] J. Ren, Y. Guo, D. Zhang, Q. Liu, and Y. Zhang, "Distributed and efficient object detection in edge computing: Challenges and solutions,"  
1029 *IEEE Network*, vol. 32, no. 6, pp. 137–143, Nov./Dec. 2018.
- 1030 [155] I. Benkacem, T. Taleb, M. Bagaa, and H. Flinck, "Performance benchmark of transcoding as a virtual network function in CDN as a  
1031 service slicing," in *Proc. Wireless Commun. and Netw. Conf. (WCNC)*, 2018, pp. 1–6.
- 1032 [156] D. Sanvito, G. Siracusano, and R. Bifulco, "Can the network be the AI accelerator?" in *Proc. ACM Workshop on In-Network Computing  
1033 (NetCompute)*, Aug. 2018, pp. 20–25.