

**A peer-reviewed version of this preprint was published in PeerJ on 14 November 2016.**

[View the peer-reviewed version](https://peerj.com/articles/cs-98) (peerj.com/articles/cs-98), which is the preferred citable publication unless you specifically need to cite this preprint.

Katsikas GP, Enguehard M, Kuźniar M, Maguire Jr GQ, Kostić D. 2016. SNF: synthesizing high performance NFV service chains. PeerJ Computer Science 2:e98 <https://doi.org/10.7717/peerj-cs.98>

# SNF: Synthesizing high performance NFV service chains

Georgios P Katsikas <sup>Corresp., 1</sup>, Marcel Enguehard <sup>2,3</sup>, Maciej Kuźniar <sup>1</sup>, Gerald Q Maguire Jr. <sup>1</sup>, Dejan Kostić <sup>1</sup>

<sup>1</sup> Department of Communication Systems (CoS), School of Information and Communication Technology (ICT), KTH Royal Institute of Technology, Kista, Stockholm, Sweden

<sup>2</sup> Network and Computer Science Department (INFRES), Telecom ParisTech, Paris, France

<sup>3</sup> Paris Innovation and Research Laboratory (PIRL), Cisco Systems, Paris, France

Corresponding Author: Georgios P Katsikas  
Email address: katsikas@kth.se

In this paper we introduce SNF, a framework that synthesizes (S) network function (NF) service chains by eliminating redundant I/O and repeated elements, while consolidating stateful cross layer packet operations across the chain. SNF uses graph composition and set theory to determine traffic classes handled by a service chain composed of multiple elements. It then synthesizes each traffic class using a minimal set of new elements that apply single-read-single-write and early-discard operations.

Our SNF prototype takes a baseline state-of-the-art network functions virtualization (NFV) framework to the level of performance required for practical NFV service deployments. Software-based SNF realizes long (up to 10 NFs) and stateful service chains that achieve line-rate 40 Gbps throughput (up to 8.5x greater than the baseline NFV framework). Hardware-assisted SNF, using a commodity OpenFlow switch, shows that our approach scales at 40 Gbps for Internet Service Provider-level NFV deployments.

# SNF: Synthesizing high performance NFV service chains

Georgios P. Katsikas<sup>1</sup>, Marcel Enguehard<sup>2,3</sup>, Maciej Kuźniar<sup>1</sup>,  
Gerald Q. Maguire Jr.<sup>1</sup>, and Dejan Kostić<sup>1</sup>

<sup>1</sup>KTH Royal Institute of Technology, Kista, Sweden

<sup>2</sup>Network and Computer Science Department (INFRES), Telecom ParisTech, Paris,  
France

<sup>3</sup>Paris Innovation and Research Laboratory (PIRL), Cisco Systems, Paris, France

Corresponding author:

Georgios P. Katsikas<sup>1</sup>

Email address: katsikas@kth.se

## ABSTRACT

In this paper we introduce SNF, a framework that synthesizes (S) network function (NF) service chains by eliminating redundant I/O and repeated elements, while consolidating stateful cross layer packet operations across the chain. SNF uses graph composition and set theory to determine traffic classes handled by a service chain composed of multiple elements. It then synthesizes each traffic class using a minimal set of new elements that apply single-read-single-write and early-discard operations.

Our SNF prototype takes a baseline state-of-the-art network functions virtualization (NFV) framework to the level of performance required for practical NFV service deployments. Software-based SNF realizes long (up to 10 NFs) and stateful service chains that achieve line-rate 40 Gbps throughput (up to 8.5x greater than the baseline NFV framework). Hardware-assisted SNF, using a commodity OpenFlow switch, shows that our approach scales at 40 Gbps for Internet Service Provider-level NFV deployments.

## INTRODUCTION

Middleboxes hold a prominent position in today's networks as they substantially enrich the dataplane's functionality (Sherry et al., 2012; Gember-Jacobson et al., 2014). However, to manage traditional middleboxes requires costly capital and operational expenditures; hence, network operators are adopting network functions virtualization (NFV) (European Telecommunications Standards Institute, 2012).

Among the first challenges in NFV was to scale software-based packet processing by exploiting the characteristics of modern hardware architectures. To do so, several works leveraged parallelism first across multiple servers and then across multiple cores, sockets, memory controllers, and graphical processing units (GPUs) (Han et al., 2010; Kim et al., 2015b) within a single server (Dobrescu et al., 2009, 2010).

Attaining hardware-based forwarding performance was difficult to achieve, even with highly-scalable software-based packet processing frameworks. The main reason was the poor I/O performance of these frameworks. Thus, the focus of both industry and academia shifted to customizing the operating systems (OSs) to achieve high-speed network I/O. For example, by using batch packet processing (Kim et al., 2012), static memory pre-allocation, and zero copy data transfers (Rizzo, 2012; DPDK, 2016).

Modern applications require combinations of network functions (NFs), also known as service chains, to satisfy their services' quality requirements (Quinn and Nadeau, 2015). With all the above advancements in place, NFV instances achieved line-rate forwarding at tens of millions of packets per second (Mpps); however, performance issues remain when several NFs are chained together. State-of-the-art frameworks such as ClickOS (Martins et al., 2014) and NetVM (Hwang et al., 2014) have reported substantial throughput degradation when realizing chains of interconnected, monolithic NFs.

The first consolidation attempts targeted application layer (e.g., deep packet inspection (DPI)) (Bremner-Barr et al., 2014) and session layer (e.g., HTTP) (Sekar et al., 2012) consolidation. However, a lot of redundancy still resides lower in the network stack. Anderson et al. (2012) describe how xOMB allows them to build programmable and extensible open middleboxes specialized for request/response based communication. In addition, Slick (Anwer et al., 2015) introduced a programming language to deploy

48 network-wide service chains, driven by a controller. Slick avoids redundant operations and shares common  
49 elements; however, its decentralized consolidation still realizes a chain of NFs as distributed processes.  
50 Most recently, E2 (Palkar et al., 2015) showed how to schedule NFs across a cluster of machines for high  
51 throughput. Also, OpenBox (Bremner-Barr et al., 2016) introduced an algorithm that merges processing  
52 graphs from different NFs into a single processing graph. Contemporaneously with E2 and OpenBox, our  
53 work implements the mechanisms fully specified in (Enguehard, 2016) and represents the next logical  
54 step of high-performance NFV research\*.

55 In the case of network-wide deployments, chains suffer from the latency imposed by interconnecting  
56 different machines, processes, and switches, along with potential virtualization overheads. In the case  
57 of single-server deployments, where the NFs are pinned to a specific (set of) core(s), throughput is  
58 bounded by the increasing number of context switches as the length of the chain increases. Based on  
59 our measurements, context switches cause a domino effect on cache utilization because of continuous  
60 data invalidations and the number of CPU cycles spent forwarding packets along the chain. This leads to  
61 increased end-to-end packet latency and considerable variation in latency (jitter).

62 In this paper, we describe the design and implementation of the Synthesized Network Functions (SNF),  
63 our approach for dramatically increasing the performance of NFV service chains. The idea in SNF is  
64 simple: create spatial correlation to execute service chains as close as possible to the speed of CPU cores  
65 operating on the fastest, L1 cache of modern multi-core machines. SNF leverages the ever-continuing  
66 increases in core counts of modern machines and the recent advances in user-space networking.

67 SNF automatically derives traffic classes of packets that are traversing a provider-specified service  
68 chain of NFs. Packets in a traffic class are all processed the same way. Additionally, SNF handles stateful  
69 NFs. Using its understanding of each of the per-traffic class chains, SNF then *synthesizes equivalent,*  
70 *high-performance NFs* for each of the traffic classes. In a straightforward SNF deployment, one CPU core  
71 processes one traffic class. In realistic scenarios, SNF allocates multiple CPU cores to execute different  
72 sets of traffic classes in isolation (see § 2).

73 SNF’s optimization process performs the following tasks: (i) consolidates all the **read** operations of a  
74 traffic class into one element, (ii) early-discards those traffic classes that lead to packet drops, and (iii)  
75 associates each traffic class with a **write-once** element. Moreover, SNF shares elements among NFs to  
76 avoid unnecessary overhead, and compresses the number and length of the chain’s traffic classes. Finally,  
77 SNF scales with an increasing number of NFs and traffic classes.

78 This architecture shifts the challenge to packet classification, as one component of SNF has to  
79 classify an incoming packet into one of the pre-determined traffic classes, and pass it to the synthesized  
80 function. We extended popular, open-source software to improve the performance of software-only packet  
81 classification. In addition, we employed an OpenFlow (McKeown et al., 2008) switch as a packet classifier  
82 demonstrating the performance possible by a sufficiently powerful programmable network interface  
83 (commonly abbreviated as NIC). The benefits for network operators are multifold: (i) SNF dramatically  
84 increases the throughput of long NF chains, and achieves low latency, and (ii) it does so while preserving  
85 the functionality of the original service chains.

86 We implemented the SNF design principles into an appropriately modified version of the Click (Kohler  
87 et al., 2000) framework. To demonstrate SNF’s superior performance, we compare it against the fastest  
88 Click variant to date, called FastClick (Barbette et al., 2015). To show SNF’s generality we tested its  
89 performance in three use cases: (i) a chain of software routers, (ii) nested network address and port  
90 translators (NAPT) (Liu et al., 2014), and (iii) access control lists (ACLs) using actual NF configurations  
91 taken from Internet Service Providers (ISPs) (Taylor and Turner, 2007).

92 Our evaluation shows that software-based SNF achieves 40 Gbps, even with small Ethernet frames,  
93 across long (up to 10 NFs), stateful chains. In particular, it achieves up to 8.5x more throughput and 10x  
94 lower latency with 2-3.5x lower latency variance than the original NF chains implemented with FastClick-  
95 when running on the same hardware. Offloading traffic classification to a commodity OpenFlow switch  
96 allows SNF to realize realistic ISP-level chains at 40 Gbps (for most of the frame sizes), while bounding  
97 the median chain latency below 100  $\mu$ s (measured from separate sending and receiving machines).

98 In the rest of this paper, we provide an overview of SNF in § 2. We introduce our synthesis approach  
99 in § 3 and a motivating example in § 4. Implementation details and performance evaluation are presented  
100 in § 5 and § 6 respectively. We discuss verification aspects in § 7. § 8 discusses the limitations of this  
101 work and § 9 positions our work with respect to the state of the art. Finally, § 10 concludes this paper.

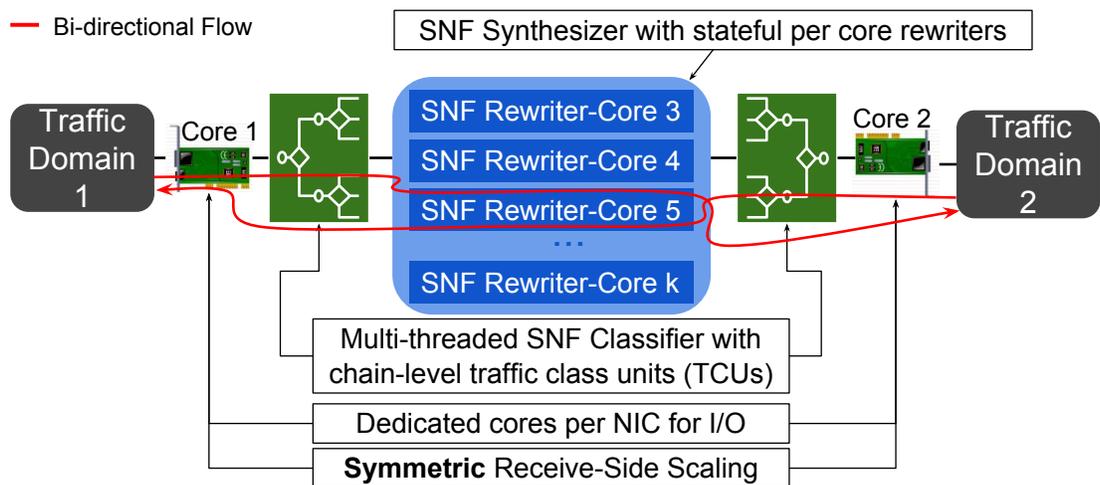
---

\*We provide a detailed comparison of our work with both E2 and OpenBox in § 9.

## SNF OVERVIEW

The idea of synthesizing network service components consorts with a powerful property: *the data correlation of network traffic*. In a network system, this property is mapped to a *spatial locality with respect to the caches*. SNF aggregates parts of the flow space into traffic class units (TCUs) (the detailed definition is in § 3.1), which are then mapped to sets of (re)write operations. By carefully setting the CPU affinity of each TCU, this aggregation enforces a large degree of correlation in the traffic requests (seen as logical units of data) resulting in high cache hit rates.

Our overarching goal is to design a system that efficiently utilizes per core and across cores cache hierarchies. With this in mind, we design SNF based on Figure 1. Let us assume that a network operator wants to deploy a service chain between network domains 1 and 2. For simplicity let us also assume that there is one NIC per domain. A set of dedicated cores (i.e., Core 1 and 2 for the NICs facing domains 1 and 2, respectively) undertakes to read and write frames at line-rate. Once a set of frames is received, say by core 1, it is transferred to the available processing cores (i.e., Cores 3 to k). Frame transfers can occur at high speed via a shared cache, which has substantial capacity in modern hardware architectures.



**Figure 1.** An overview of SNF running on a machine with  $k$  CPU cores and 2 NICs. Dedicated CPU cores per NIC deliver bi-directional flows to packet processing CPU cores via Symmetric RSS. Processing cores concurrently classify traffic and access individual, stateful SNF rewriters to modify the traffic.

Once a processing core acquires a frame, it executes SNF as shown in Figure 1. First the core classifies the frame (green rectangles in Figure 1) in one of the chain’s TCUs and then applies the required synthesized modifications (blue rounded-rectangle in Figure 1) that correspond to this TCU. Both classification and modification processes are highly parallelized as different cores can simultaneously drive frames that belong to different TCUs out of the chain. We detail both processes in § 3.2.

However, the key point of Figure 1 lies in the fact that a core’s pipeline shares nothing with any other pipeline. We employed the symmetric Receive Side Scaling (RSS) (Intel, 2016) scheme by Woo and Park (2012) to hash input traffic in a way that a flows’ bi-directional packets are always served by the same SNF rewriter, hence the same processor. This scheme allows a processing core to drive a TCU at the maximum processing speed of the machine.

### Main Objectives

The primary goal of SNF is to eliminate redundancy along the chain. The sources of redundancy in current NF chains and the solutions that our approach offers are:

**A. Multiple network I/O** interactions between the chain and the backend dataplane occur because each NF is an individual process. We solve this by placing NF chains in a single logical entity. Once a packet enters this entity, it does not exit until all the chain operations are applied.

**B. Late packet drops** appear in NF chain implementations when packets unnecessarily pass through several elements before getting dropped. SNF discards these packets as early as possible.

**C. Multiple read operations** on the same field occur because each NF contains its own decision elements. A typical example is an Internet protocol (IP) lookup in a chain of routers. While SNF is parsing the

136 initial chain, it marks the read operations and constructs traffic classes encoded as paths of elements in a  
 137 directed acyclic graph (DAG). Then, SNF synthesizes these elements into a *single* classifier to realize  
 138 both routing and filtering.

139 **D. Multiple write operations** on the same field overwrite previous values. For example, the IP checksum  
 140 is modified twice when a decrement time to live (TTL) operation follows a destination IP address  
 141 modification. SNF associates a set of (stateful) write operations with a traffic class, hence it can modify  
 142 each field of a traffic class all at once.

143 Next, we describe in detail how SNF *automatically* synthesizes the equivalent of a service chain.

## 144 SNF ARCHITECTURE

145 Taking into account the main objectives listed above, this section presents the design of SNF. § 3.1  
 146 defines the synthesis abstraction, § 3.2 presents the formal synthesis steps, and § 3.3 describes how  
 147 stateful functions are realized.

### 148 Abstract Service Chain Representation

149 The crux of SNF's design is an abstract service chain representation. We begin by describing a  
 150 mathematical model to represent packet units in § 3.1.1. Next, we model an NF's behavior in an abstract  
 151 way in § 3.1.2. Finally, we define our target service-level network function in § 3.1.3.

### 152 Packet Unit Representation

Inspired by the approach of Kazemian et al. (2012), we represent each packet as a vector in a multi-  
 dimensional space. However, we follow a protocol-aware approach by dividing a packet according to the  
 unsigned integer value of the different header fields. Thus, if  $p$  is an IPv4/TCP packet, we represent it as:

$$p = (p_{ip\_version}, p_{ip\_ihl}, \dots, p_{tcp\_sport}, p_{tcp\_dport}, \dots)$$

From now on, we call  $P$  the space of all possible packets. For a given header field  $f$  of length  $l$  bits, we  
 define a field filter  $F_f$  as a union of disjoint intervals  $(0, 2^l - 1)$ :

$$F_f = \bigcup_{s_i \subset (0, 2^l - 1)} s_i \text{ where } \begin{cases} \forall i, & s_i \text{ is an interval} \\ \forall i \neq j, & s_i \cap s_j = \emptyset \end{cases}$$

This allows grouping packets into a data structure that we call a *packet filter*, defined as a logical  
 expression of the form:

$$\phi = \{(p_1, \dots, p_n) \in P \mid (p_1 \in F_1) \wedge \dots \wedge (p_n \in F_n)\}$$

where  $(F_1, \dots, F_n)$  are field filters. The space of all possible packet filters is  $\Phi$ . Then:

$$u : \begin{cases} \phi & \mapsto (F_1, \dots, F_n) \\ \Phi & \mapsto \{(F_1, \dots, F_n) \mid \forall i, F_i\}_{(F_1, \dots, F_n)} \end{cases}$$

153 is a bijection and we can assimilate  $\phi$  to  $(F_1, \dots, F_n)$ .

154 If  $\phi_1$  and  $\phi_2$  are two packet filters defined by their field filters  $(F_{1,1}, \dots, F_{1,n})$  and  $(F_{2,1}, \dots, F_{2,n})$ , then  
 155  $\phi_1 \cap \phi_2$  is also a packet filter and is defined as  $(F_{1,1} \cap F_{2,1}, \dots, F_{1,n} \cap F_{2,n})$ .

### 156 Network Function Representation

Network functions typically apply read and write operations to traffic. While our packet unit  
 representation allows us to compose complex read operations across the entire header space, we still need  
 the means to modify traffic. For this, we define an operation as a function  $\omega : P \mapsto \Phi$  that associates a set  
 of possible outputs to a packet. We add the additional constraint that for any given operation  $\omega$ , there is  
 $\omega_1, \dots, \omega_n \in \mathbb{N}^{\mathbb{N}}$  such as:

$$\forall p = (p_1, \dots, p_n) \in P, \omega(p) = (\omega_1(p_1), \dots, \omega_n(p_n))$$

157 Note that we use sets of possible values (instead of fixed values) to model cases where the actual value is  
 158 chosen at run-time (e.g., source port in an S-NAT). *Therefore, SNF does support both deterministic and*  
 159 *conditional operations.*

If we define  $\Omega$  as the space of all possible operations, we can express a **processing unit**  $PU$  as a conditional function that maps packet filters to operations:

$$PU : p \mapsto \begin{cases} \omega_1(p) & \text{if } p \in \phi_1 \\ \dots & \\ \omega_m(p) & \text{if } p \in \phi_m \end{cases}$$

160 where  $(\omega_1, \dots, \omega_m) \in \Omega^m$  are operations and  $(\phi_1, \dots, \phi_m) \in \Phi^m$  are mutually distinct packet filters.

An NF is simply a DAG of PUs. For instance, SNF can express a simplified router's NF as follows:

$$NF_{ROUTER} : PU\{Lookup\} \rightarrow PU\{DecIPTTL\} \rightarrow PU\{IPChecksum\} \rightarrow PU\{MAC\}$$

161 where, 4 PUs take place. An IP lookup PU is followed by decrement IP TTL, IP checksum update, and  
162 source and destination MAC address modification PUs.

### 163 **The Synthesized Network Function**

In the previous section we laid the foundation to construct NFs as graphs of PUs. Now, at the service level where multiple NFs can be chained, we define a TCU as a set of packets/flows, represented by disjoint unions of packet filters, that are processed in the same fashion (i.e., undergo the same set of synthesized operations). This definition allows us to construct the service chain's *SynthesizedNF* function (in short SNF) as a DAG of PUs, or equivalently, as a map of TCUs that associates operations to their packet filters:

$$SynthesizedNF : \Phi \mapsto \Omega$$

164 Formally, the complexity of the *SynthesizedNF* is upper-bounded by the function  $O(n \cdot m)$ , where  $n$  is the  
165 number of TCUs and  $m$  is the number of packet filters (or conditions) per TCU. Each TCU turns a textual  
166 packet filter specification (such as "proto tcp && dst net 10.0/16 && src port 80") into a binary decision  
167 tree traversed by each packet. Therefore, in the absolute worst case, an input packet might traverse a  
168 skewed binary tree of the last TCU, yielding the above complexity bound. The average case occurs  
169 in a relatively balanced tree ( $O(\log m)$ ), in which case the average complexity of the *SynthesizedNF* is  
170 bounded by the function  $O(n \cdot \log m)$ .

### 171 **Synthesis Steps**

172 Leveraging the abstractions introduced in § 3.1, we detail the steps that translate a set of NFs into an  
173 equivalent SNF. The SNF architecture is comprised of three modules (shown in Figure 2). We describe  
174 each module in the following sections.

#### 175 **Service Chain Configurator**

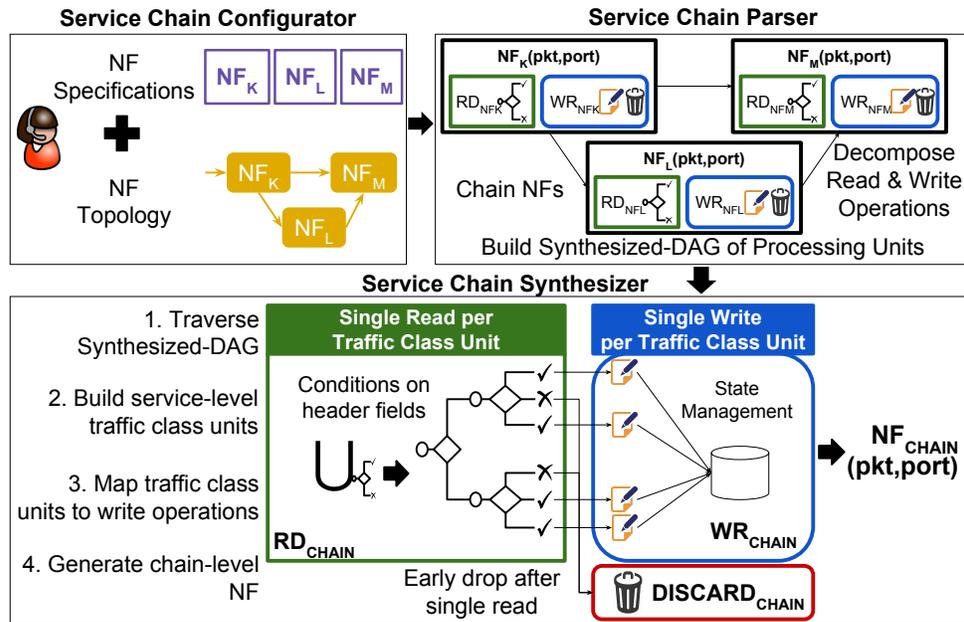
176 The top left box in Figure 2 is the Service Chain Configurator; the interface that a network operator  
177 uses to specify a service chain to be synthesized by SNF. Two inputs are required: a set of service  
178 components (i.e., NFs), along with their topology. SNF abstracts packet processing by using graph theory.  
179 That said, a chain is described as a DAG of interconnected NFs (i.e., chain-level DAG), where each NF is  
180 a DAG of abstract packet processing elements (i.e., NF DAG). The NF DAG is implementation-agnostic,  
181 similar to the approaches of [Bremner-Barr et al. \(2016\)](#); [Anwer et al. \(2015\)](#); [Kohler et al. \(2000\)](#). The  
182 network operator enters these inputs in a configuration file using the following notation:

183 **Vertices (NFs):** Each service component (i.e., an NF) of a chain is a vertex in the chain-level DAG  
184 for which, the Service Chain Configurator expects a name and an NF DAG specification (see Figure 2).  
185 Each NF can have any number of input and output ports as specified by its DAG. An NF with one input  
186 and one output interface is denoted as:  $[interface_0]NF_1[interface_1]$ .

187 **Edges (NF inter-connections):** The connections between NFs are the edges of the chain-level DAG.  
188 We interconnect two NFs as follows:  $NF_1[interface_1] \rightarrow [interface_0]NF_2$ .

189 **No loops:** Since the chain-level DAG is acyclic by construction, SNF must prevent loops (e.g., two  
190 interfaces of the same NF cannot be connected to each other).

191 **Entry points:** In addition to the internal connections within a chain (i.e., connections between NFs),  
192 the Service Chain Configurator also requires the entry points of the chain. These points are the interfaces  
193 of the chain with the outside world and indicate the existence of traffic sources. An interface that is neither  
194 internal nor an entry point can only be an end-point; these interfaces are discovered by the Service Chain  
195 Parser as described below.



**Figure 2.** The SNF framework. The network operator inputs a service chain and its topology (top left part). SNF parses the chained NFs, decomposes their read and write parts, and composes a Synthesized-DAG (top right part). While traversing the Synthesized-DAG, SNF builds the TCUs of the chain, associates them with write/discard operations, leading to a synthesized chain-level NF.

### 196 **Service Chain Parser**

197 The Service Chain Configurator outputs a chain-level DAG that describes the chain to the Service  
 198 Chain Parser. As shown in the top right box of Figure 2, the parser iterates through all of the input NF  
 199 DAGs (i.e., one per NF); while parsing each NF DAG, the parser marks each element according to its  
 200 type. We categorize NF elements in four types: I/O, parsing, read, and write elements. As an example  
 201 NF, consider a router that consists of interconnected elements, such as *ReadFrame*, *StripEthernetHeader*,  
 202 *IPLookUp*, and *DecrementIPTTL*. *ReadFrame* is an I/O element, *StripEthernetHeader* is a parsing  
 203 element (moves a frame’s pointer), *IPLookUp* is a read element, while *DecrementIPTTL* is a write  
 204 element.

205 The parser stitches together all the NF DAGs based on the topology graph and builds a Synthesized-  
 206 DAG (see Figure 2) that represents the entire chain. This process begins from an entry point and searches  
 207 recursively until an output element is found. If the output element leads to another NF, the parser keeps a  
 208 jump pointer and cross checks that the encountered interfaces match the interfaces declared in the Service  
 209 Chain Configurator. After collecting this information, the parser omits the I/O elements because one of  
 210 SNF’s objectives is to eliminate inter-NF I/O interactions. The process continues until an output element  
 211 that is not in the topology is found; such an element can only be an **end-point**. Along the path to an  
 212 output element the parser separates the read from the write elements and transforms NF elements into  
 213 PUs, according to § 3.1.2. Next, the parser considers the next entry point until all are exhausted.

214 The final output of the Service Chain Parser is a large Synthesized-DAG of PUs that models the  
 215 behavior of the entire input service chain.

### 216 **Service Chain Synthesizer**

217 After building the Synthesized-DAG, our next target is to create the *SynthesizedNF* introduced  
 218 in § 3.1.3. To do so, we need to derive the SNF’s TCUs. To build a TCU we execute the following steps:  
 219 from each entry port of the Synthesized-DAG, we start from the identity TCU  $tcu_0 \in \Phi \times \Omega$  defined  
 220 as:  $tcu_0 = (P, id_P)$ , where  $id_P$  is the identity function of  $P$ , i.e.,  $\forall x \in P, id_P(x) = x$ . Conceptually,  $tcu_0$   
 221 represents an empty packet filter and no operations, which is equivalent to a transparent NF. Then, we  
 222 search the Synthesized-DAG, while updating our TCU as we encounter conditional (read) or modification  
 223 (write) elements. Algorithms 1 and 2 build the TCUs using an adapted depth-first search (DFS) of the  
 224 Synthesized-DAG.

225 Now let us consider a TCU  $t$ , defined by its packet filter  $\phi$  and its operation  $\omega$ , that traverses a PU  $U$   
 226 using the adapted DFS. The TRAVERSE function in Algorithm 1 creates a new TCU for each possible

227 pair of  $(\omega_i, \phi_i)$ . In particular, it creates a new packet filter  $\phi'$  returned by the INTERSECT function (line 3).  
 228 This function is described in Algorithm 2 and considers previous write operations while updating a packet  
 229 filter. For each field filter  $\phi_i$  of a packet filter, the function checks whether the value has been modified by  
 230 the corresponding  $\omega_i$  operation (condition in line 8) and whether the written value is in the intersecting  
 231 field filter  $\phi_i^0$  (line 10). It then updates the TCU by intersecting it with the new filter, if the value has not  
 232 been modified (action in line 8). After the INTERSECT function returns in Algorithm 1, TRAVERSE creates  
 233 a new operation by composing  $\omega$  and  $\omega_i$  (line 4).

234 The recursive algorithm terminates in two cases: (i) when the packet filter of the current TCU is the  
 235 empty set, in which case the function does not return anything, (ii) when the PU  $U$  does not have any  
 236 successors, in which case it returns the current TCUs. In the latter case, the returned TCUs comprise the  
 237 final *SynthesizedNF* function.  
 238

#### Algorithm 1 Building the SNF TCUs

```

1: function TRAVERSE( $t = (\phi, \omega), U = \{(\phi_i, \omega_i)_{i \leq m}\}$ )
2:   for  $i \in (1, m)$  do 0
3:      $\phi' \leftarrow \text{INTERSECT}(t, \phi_i)$ 
4:      $\omega' \leftarrow \omega_i \circ \omega$ 
5:      $t' = (\phi', \omega')$ 
6:     TRAVERSE( $t', U.successors[i]$ )
  
```

#### Algorithm 2 Intersecting a TCU with a filter

```

1: function INTERSECT( $t = (\phi, \omega), \phi^0$ )
2:    $\phi' \leftarrow P$ 
3:    $(\omega_1, \dots, \omega_n) \leftarrow \omega.COORDINATES$ 
4:    $(\phi_1, \dots, \phi_n) \leftarrow \phi.COORDINATES$ 
5:    $(\phi_1^0, \dots, \phi_n^0) \leftarrow \phi^0.COORDINATES$ 
6:    $(\phi'_1, \dots, \phi'_n) \leftarrow \phi'.COORDINATES$ 
7:   for  $i \in (1, n)$  do
8:     if  $\omega_i = id_{\mathbb{N}}$  then  $\phi'_i \leftarrow \phi_i \cap \phi_i^0$ 
9:     else
10:      if  $\omega_i(\phi_i) \subset \phi_i^0$  then  $\phi'_i \leftarrow \phi_i$ 
11:      else  $\phi'_i \leftarrow \emptyset$ 
12:   return  $\phi'$ 
  
```

### 239 Managing Stateful Functions

240 A difficulty when synthesizing NF chains is managing successive stateful functions. It is crucial to  
 241 ensure that the states are properly located in a synthesized NF and that every packet is matched against  
 242 the correct state table. At the same time, SNF should hold the promise that NFV service chains must be  
 243 realized without redundancy, hence single-read and single-write operations must be applied per packet.

244 To highlight the challenges of maintaining the state in a chain of NFs, consider the example topology  
 245 shown in Figure 3. In this example, a large network operator has run out of private IPv4 addresses in the  
 246 10.0/8 prefix and has been forced to share the same network prefix between two distinct zones (i.e., zones  
 247 1 and 2), using a chain of NAPT's. This is not unlikely to happen, as an 8-byte network prefix contains less  
 248 than 17 million addresses and recent surveys have predicted that 50 billion addresses will be connected to  
 249 the Internet by 2020 (Evans, D., 2011).

250 Consolidating this chain of NFs into a single SNF instance poses a problem. That is, traffic originating  
 251 from zones 1 and 2 shares the same source IP address and port range, but to ensure that all the traffic is  
 252 translated properly, the corresponding synthesized chain must share their NAPT table. However, since  
 253 traffic also shares the same destination prefix (i.e., towards the same Internet gateway), a host from the  
 254 outside world cannot possibly distinguish the zone where the traffic is originating from.

255 Obviously, the question that SNF has to address in general, and particularly in this example is: "How  
 256 can we synthesize a chain of NFs, ensuring that (i) traffic mappings are unique and (ii) no redundant  
 257 operations will be applied?" To solve this conundrum, the SNF design respects the following properties:

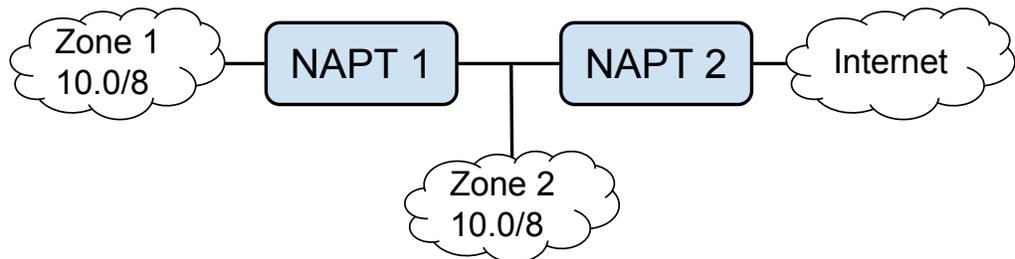
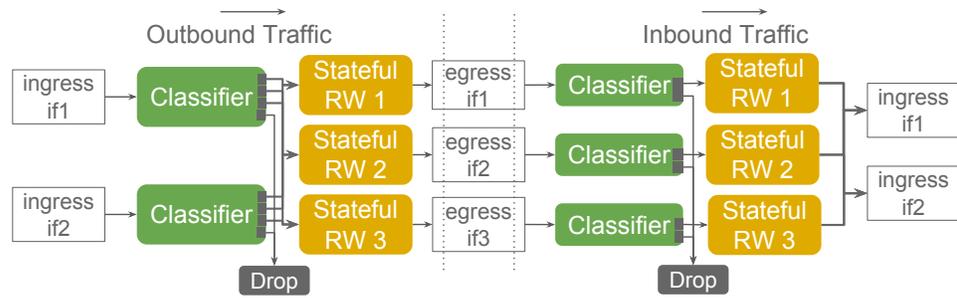


Figure 3. Example of stateful NAPT chains, where two zones share the same IPv4 prefix.

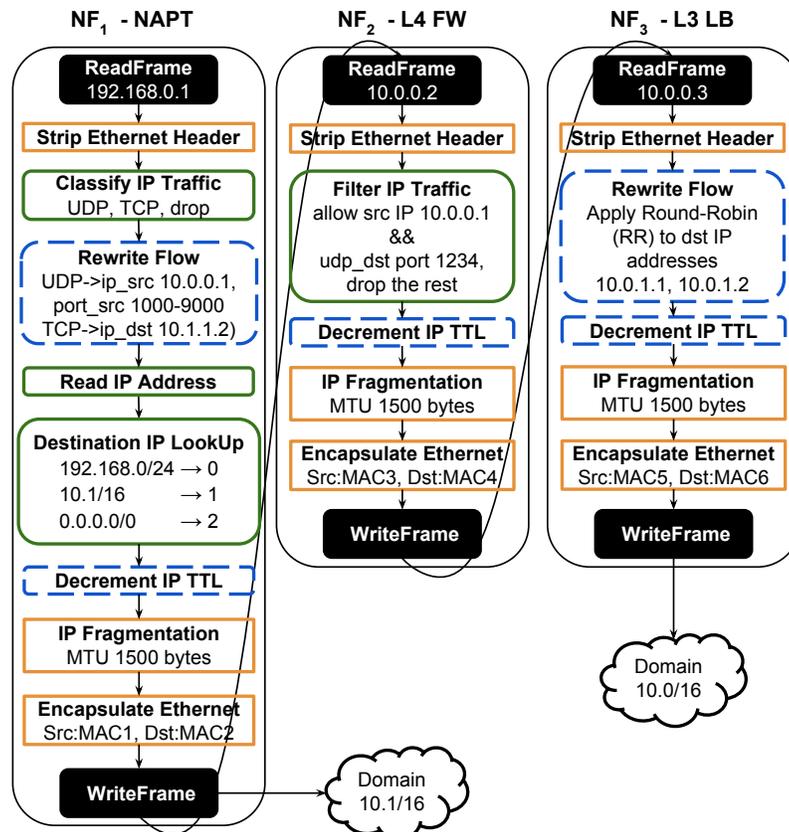
258 **Property 1** We enforce the uniqueness of flow mappings by ensuring that all egress traffic that shares  
 259 the same last stateful (re)write operation also shares the same state table.  
 260 **Property 2** The state table of SNF must be origin-aware. To redirect ingress traffic towards the  
 261 correct interface, while respecting the single-read principle of SNF, the SNF state table  
 262 must collocate flow information and the origin interface for each flow.  
 263 To generalize the state management problem, Figure 4 shows how SNF handles stateful configurations  
 264 with e.g., three egress interfaces. We apply “Property 1” by having exactly one stateful (re)write element  
 265 (denoted as Stateful RW) per egress interface. We apply “Property 2” by having one input port in each of  
 266 these (re)write elements, associated with an ingress interface. Therefore, a state table in SNF not only  
 contains flow-related information, but also keeps a linking of a flow entry with its origin interface.



267 **Figure 4.** State management in SNF.

## 268 A MOTIVATING USE CASE

269 To understand how SNF works and what benefits it can offer, we quantify the processing and I/O  
 270 redundancies in an example use case of an NF chain and then compare it to its synthesized counterpart.  
 271 We use Click to specify the NF DAGs of this example, but SNF is applicable to other frameworks.  
 272 The example chain consists of a NATP, a L4 firewall (FW), and a L3 load balancer (LB) that process  
 transmission control protocol (TCP) and user datagram protocol (UDP) traffic as shown in Figure 5.



273 **Figure 5.** The internal components of an example NATP - L4 FW - L3 LB chain.

274 The TCP traffic is NAPT'ed in the first NF and then leaves the chain, while UDP is filtered at the FW  
 275 (the second NF) and the UDP datagrams with destination port 1234 are load balanced across two servers  
 276 by the last NF. For simplicity, we discuss only the traffic going in the direction from the NAPT to the LB.

277 The rectangular operations in Figure 5 are interface-dependent, e.g., an “Encapsulate Ethernet”  
 278 operation encapsulates the IP packets in Ethernet frames before passing them to the next NF where a  
 279 “Strip Ethernet Header” operation turns them back into IP packets. Such operations occur 3 times because  
 280 there are 3 NFs, instead of only once (because the processing operates at the IP layer). Ideally, strip  
 281 should be applied before, and Ethernet encapsulation after all of the IP processing operations. Similarly,  
 282 the “IP Fragmentation” should only be applied before the final Ethernet encapsulation.

283 The remaining operations (illustrated as rounded rectangles) of the three processing stages are  
 284 those that (i) make decisions based upon the contents of specific packet fields (read operations with a  
 285 solid round outline, e.g., “Classify IP Traffic” and “Filter IP Traffic”) or (ii) modify the packet header  
 286 (rewrite operations with a blue dashed outline e.g., “Rewrite Flow” and “Decrement IP TTL”). We  
 287 found redundancy in both types of operations. In the read operations, one IP classifier is sufficient to  
 288 accommodate the three traffic classes of this example and perform the routing. Thus, all the round-outlined  
 289 operations with solid lines (green) can be replaced by a single “Classify IP Traffic” operation.

290 Large savings are also possible with the rewrite operations. For example, the initial chain calculates  
 291 the TTL field 3 times and IP checksum 5 times, whereas only one computation for these fields suffices  
 292 in the synthesized chain. Based on our measurements on an Intel Xeon E5 processor the checksum  
 293 calculations cost 10-40 CPU cycles/packet. By integrating the “Decrement IP TTL” into the “Rewrite  
 294 Flow” operation and enforcing the checksum calculation only once, saves 237 CPU cycles/packet.

295 Figure 6 depicts a synthesized version of the NF chain shown in Figure 5. Following the SNF paradigm  
 296 presented in § 3, the synthesized chain forms a graph with two main parts. The left-most part (rounded  
 297 rectangles with solid outline in Figure 6) encodes all the read operations by composing paths that begin  
 298 from a specific interface and traverse the three traffic classes of this chain, until a packet is output or  
 299 dropped. Each path keeps a union of filters that represents the header space that matches the respective  
 300 traffic class. In this example, the filter for e.g., the allowed UDP packets is the union of the protocol and  
 301 destination port numbers. Such a filter is part of a classifier whose output port is linked with a set of write  
 302 operations (dashed vertices in Figure 6) associated with this traffic class (right-most part of the graph).  
 303 As shown in Figure 6, with SNF a packet passes through all the read operations once (guaranteeing  
 304 a single-read) and either the packet is discarded early or each header field is written once (ensuring a  
 305 single-write) before exiting the chain.

306 Synthesizing the counterpart of this example implies several code modifications to avoid the  
 307 redundancy caused by the design of each NF. To apply a per flow, per-field single-write operation we  
 308 ensure that the “Rewrite Flow” will smartly calculate the checksums once IP addresses, ports, and the IP  
 309 TTL fields are written. Therefore, in this example we saved four unnecessary operations (3 “Decrement IP  
 310 TTL” and 1 “Rewrite Flow”) and four checksum calculations (3 IP and 1 IP/UDP). Moreover, integrating  
 311 all decisions (i.e., routing, filtering) in one classifier caused this operation to be slightly heavier, but saved  
 312 another two redundant function calls to “Destination IP Lookup” and “Filter IP Traffic” respectively.

313 The final form of the synthesized chain requires only 5 processing operations to transfer the UDP  
 314 datagrams along the entire chain. The initial chain implements the same functionality using 18 processing

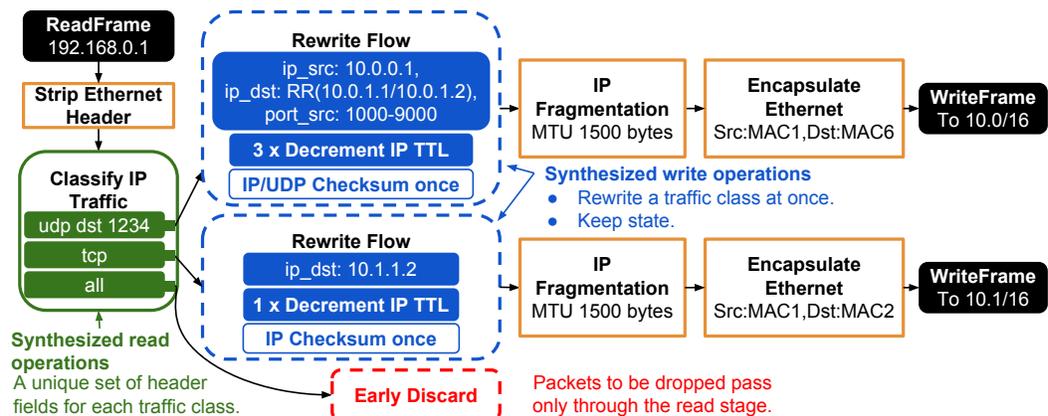


Figure 6. The synthesized chain equivalent to Figure 5. The SNF contributions are shown in floating text.

315 operations and two additional pairs of I/O operations. Based on our measurements the total *processing*  
316 cost of the initial chain is 2206 cycles/packet, while the synthesized chain requires 3x less (roughly 720)  
317 cycles/packet. If we account for the extra I/O cost per hop for the initial chain the difference becomes  
318 even greater. In production service chains, where packets arrive at high rates, this overhead can play a  
319 major role in limiting the throughput of the chain and the imposed latency; therefore, the advantages of  
320 synthesizing more complex service chains than this simple use case are expected to be even greater.

## 321 IMPLEMENTATION

322 As we stated earlier, SNF’s basic assumption is that each input service component (i.e., NF) is  
323 expressed as a graph (i.e., the NF DAG), composed of individual packet processing elements. This allows  
324 SNF to parse the NF DAG and infer the internal operations of each NF, producing a synthesized equivalent.  
325 Among the several candidate platforms that allow such a representation, we developed our prototype atop  
326 Click because it is the most widely used NFV platform in the academia. Many earlier efforts built upon it  
327 to improve its performance and scalability, hence we believe that this choice will maximize SNF’s impact  
328 as it allows direct comparison with state-of-the-art Click variants such as RouteBricks (Dobrescu et al.,  
329 2009), PacketShader (Han et al., 2010), Double-Click (Kim et al., 2012), SNAP (Sun and Ricci, 2013),  
330 ClickOS (Martins et al., 2014), and FastClick (Barbette et al., 2015).

331 We adopt FastClick as the basis of SNF as it uses DPDK, a state-of-the-art user-space I/O framework  
332 that exploits modern hardware amenities (including multiple CPU cores) and NIC features (including  
333 multiple queues and offloading mechanisms). Along with batch processing, non-uniform memory access  
334 support, and fine grained CPU core affinity techniques, FastClick scales a single router achieving line-rate  
335 throughput at 40 Gbps. **SNF aims for similar performance for an entire service chain.**

### 336 FastClick Extensions

337 We implemented SNF in C++11. The modules depicted in Figure 2 are 14376 lines of code.  
338 The integration with FastClick required another 1500 lines of code (modifications and extensions).  
339 Although FastClick improves a router’s throughput and latency, it lacks features required for broader NFV  
340 applications; therefore, we made the following extensions to target a service-oriented platform:

341 **Extension 1:** Stateful elements that deal with flow processing such as IP/UDP/TCPRewriter were not  
342 equipped with FastClick’s optimizations such as computational batching or cache prefetching. Moreover,  
343 these elements were not designed to be thread-safe hence they could cause race conditions when accessed  
344 by multiple CPU cores at the same time. We designed thread-safe data structures for these elements while  
345 also applying the necessary modifications to equip them with the FastClick optimizations.

346 **Extension 2:** We tailored several packet modification FastClick elements to comply with the synthesis  
347 principles, as we found that their implementation was not aligned with our single-write approach. For  
348 instance, we improved the IP/UDP/TCP checksum calculations by calling the respective functions only  
349 once all the header field modifications are applied. Moreover, we extended IP/UDP/TCPRewriter elements  
350 with additional input arguments. These arguments extend the elements’ packet modification capabilities  
351 (e.g., decrement IP TTL field to avoid unnecessary element calls) and guarantee that a packet entering  
352 these elements undergo a single-write operation per header field.

353 **Extension 3:** We developed a new element, called IPSynthesizer, in the heart of our execution model  
354 shown in Figure 1. This element implements per-core stateful flow tables that can be safely accessed in  
355 parallel allowing multiple TCUs to be processed at the same time. To avoid inter-core communication,  
356 thus keep the per-core cache(s) hot, we extended the RSS mechanism of DPDK (see Figure 1) using a  
357 symmetric approach proposed by Woo and Park (2012).

358 **Extension 4:** To make software-based classification more scalable, we implemented the lazy subtraction  
359 algorithm introduced in Header Space Analysis (HSA) (Kazemian et al., 2012). With this extension,  
360 SNF aggregates common IP prefixes in a filter and applies the longest one while building a TCU, thus  
361 producing shorter traffic class expressions. †

362 Our prototype supports a large variety of packet processing libraries, fully covering both native  
363 FastClick and hypervisor-based ClickOS deployments. Our prototype also takes advantage of FastClick’s  
364 computation batching with a processing core moving a group of packets between the classifier and the  
365 synthesizer with a single function call. New packet processing elements can be incorporated with minor  
366 effort. We made the FastClick extensions available at Katsikas, Georgios (2016).

† This extension is not a direct part of FastClick, since the optimized classification rules are computed by SNF beforehand; then, SNF uses these rules as arguments when calling FastClick’s Classifier or IPClassifier elements.

## PERFORMANCE EVALUATION

Recent efforts, such as ClickOS (Martins et al., 2014) and NetVM (Hwang et al., 2014), are unable to maintain constant high throughput and low latency for chains of more than 3 NFs when processing packets at high speed. This problem hinders large-scale hypervisor-based NFV deployments that could reduce network operators' expenses and provide more flexible network management and services (Cisco, 2014; SDX Central, 2015).

We envision SNF to be the key component of future NFV deployments, thus we evaluate the synthesis process using real service chains to exercise its true potential. In this section, we demonstrate SNF's ability to address three types of service chains:

**Chain 1:** Scale a long series of routers at the cost of a single router.

**Chain 2:** Nest multiple NAPT middleboxes.

**Chain 3:** Implement high performance ACLs of increasing cardinality at the borders of ISP networks.

We use the experimental setup described in § 6.1 to measure the performance of the above three types of chains and answer the following questions: Can we synthesize (stateful) chains *without* sacrificing throughput as we increase the chain length (see § 6.2, § 6.3)? What is the effect of different packet sizes on a system's throughput (see § 6.3)? What are the current limits of purely software-based packet processing (see § 6.4) and how can we overcome them (see § 6.5)?

### Testbed

We conducted our experiments on six identical machines each with a dual socket 16-core Intel<sup>®</sup> Xeon<sup>®</sup> CPU E5-2667 v3 clocked at 3.20 GHz. The cache sizes are: 2x32 KB L1, 256 KB L2, and 20 MB L3. Hyper-threading is disabled and the OS is the Ubuntu 14.04.1 distribution with Linux kernel v.3.13. Each machine has two dual-port 10 GbE Intel 82599 ES NICs.

Unless stated otherwise, we use two machines to generate and sink bi-directional traffic using MoonGen (Emmerich et al., 2015), a DPDK-based traffic generator. MoonGen allows us to saturate 10 Gbps NICs on a single machine using a set of cores, while receiving the same amount of traffic on another set of cores. To gain insight into the performance of the service chains, we measure the throughput and end-to-end latency to traverse the chains, at the endpoints. We use FastClick as a baseline and compare FastClick against SNF (which extends FastClick). We create service chains that run natively in a single process using RSS and multiple CPU cores, as this is the fastest FastClick configuration. We follow two different setups for our software-based and hardware-assisted SNF deployments as follows.

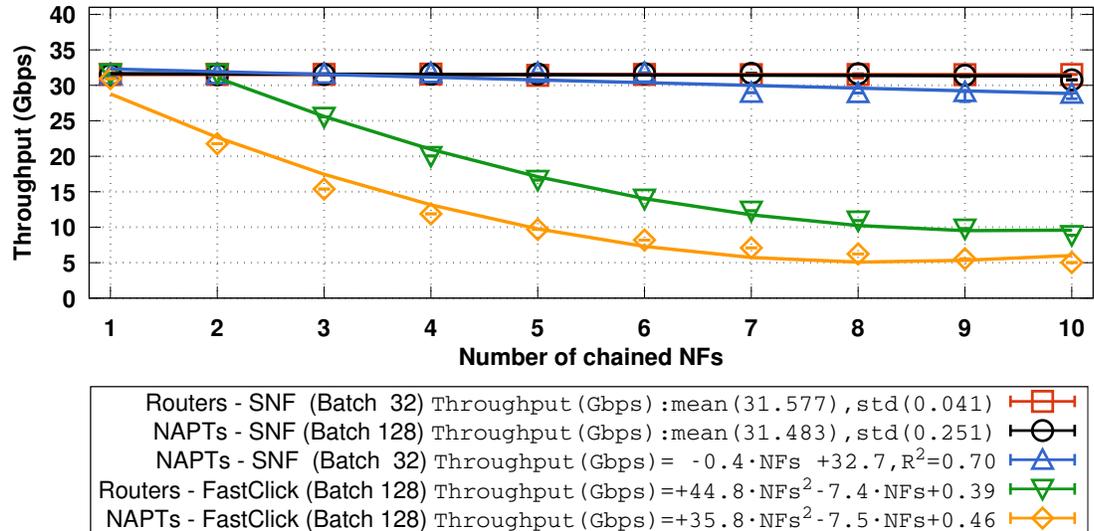
**Software-based SNF:** In § 6.2, § 6.3, and § 6.4 we stress different purely software-based NFV service chains that run in one machine following the execution model of Figure 1. This machine has 4 10 GbE NICs connected to the two traffic source/sink machines (two NICs on each machine), hence the total capacity of the NFV machine is 40 Gbps. The goal of this testbed is to show how much NFV processing FastClick and SNF can fit into a single machine and what processing limits this machine has.

**Hardware-assisted SNF:** For the complex NFV service chains, presented in § 6.4, we also deploy a testbed (see § 6.5) where we offload the traffic classification to a Noviflow 1132 OpenFlow switch with firmware 300.1.0. The switch is connected with two 10 GbE NICs to each of the two senders/receivers, and with one link to each of the four processing servers in our SNF cluster. This testbed has a total of 40 Gbps capacity (same as the software-based setup above), but the processing is distributed to more machines in order to show how our SNF system scales.

### A Chain of Routers at the Cost of One

This first use case targets a direct comparison with the state-of-the-art. Specifically, we chain a popular implementation of a software-based router that, after several years of successful research contributions (Dobrescu et al., 2009; Han et al., 2010; Kim et al., 2012; Sun and Ricci, 2013; Martins et al., 2014; Barbette et al., 2015), achieves scalable performance at tens of Gbps.

As we show in this section, a naive chaining of individual, fast NFs does not achieve high performance. To examine this we linearly connect 1-10 FastClick routers, where each router has four 10 Gbps ports (hence such a chain has a 40 Gbps link capacity). The down-pointing (green) triangular points of Figure 7 show the throughput achieved by these chains versus the increasing length of the chains, when we inject 60-bytes long frames, excluding the cyclic redundant check (CRC). The maximum throughput for this frame size size is 31.5 Gbps and this is the limit of our NICs, as reported earlier (Barbette et al., 2015).



**Figure 7.** Throughput (Gbps) of chained routers and NAPT's using (i) FastClick and (ii) SNF versus the numbers of chained NFs (60-byte frames are injected at 40 Gbps). Bigger batch sizes achieve higher throughput.

419 In our experiment, FastClick can operate at the maximum throughput only for a chain of 1 or 2  
 420 routers. As denoted by the equation in this fit to the graph, after this point there is a quadratic throughput  
 421 degradation that results in a chain of 10 routers achieving less than 10 Gbps of throughput.

422 SNF automatically synthesizes this simple chain (shown with red squares) to achieve the maximum  
 423 possible throughput using this hardware, despite the increasing length of the chain. The fitted equation  
 424 confirms that SNF operates at the speed of the NICs.

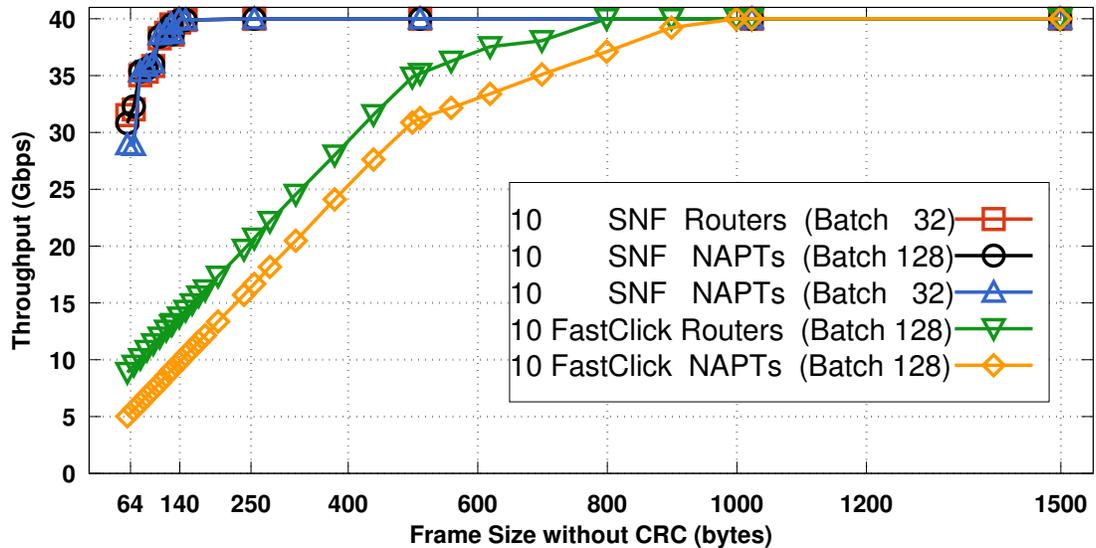
### 425 Stateful Service Chaining

426 The problem of Service Function Chaining has been recently investigated by Quinn and Nadeau (Quinn  
 427 and Nadeau, 2015) and several relevant use cases (Liu et al., 2014) have been proposed. In some of  
 428 these use cases, traffic needs to support distinct address families while traversing different networks. For  
 429 instance, within an ISP, IPv4/IPv6 traffic might either be directed to NAT64 (Bagnulo et al., 2011), or  
 430 a Carrier Grade NAT (Perreault et al., 2013). In more extreme cases, this traffic might originate from  
 431 different access networks such as fixed broadband, mobile, datacenters, or cloud customer premises (CPE),  
 432 thus causing the nested NAT problem (Penno et al., 2013).

433 The goal of this use case is to test SNF in such a stateful context using a chain of 1-10 NAPT's. Each  
 434 NAPT maintains a state table that stores the original and translated source and destination IP addresses  
 435 and ports of each flow, associated with the input interface where a flow was originated. The rhomboid  
 436 points of Figure 7 show that the chains of FastClick NAPT's suffer a steeper (according to the fitted  
 437 equation) quadratic degradation than the FastClick routers. Although we extended FastClick to support  
 438 thread-safe, parallelized NAPT operations across multiple cores, it is still unable to drive the NAPT chain  
 439 at line-rate, despite using 8 CPU cores and 128-packet batches.

440 SNF requires a certain batch size to realize the synthesized NAPT chains at the speed of hardware as  
 441 shown by the black circles of Figure 7. The curve with the up-pointing (blue) triangles indicates that a  
 442 batch size of 32 packets leads to a slight throughput degradation after the 6<sup>th</sup> NAPT in the chain. State  
 443 lookup and management operations executed for every packet cause this degradation. Depending on  
 444 the performance targets, a network operator might tolerate an increased latency to achieve the higher  
 445 throughput offered by an increased batch size.

446 Next, we explore the effect of different frame sizes on the chains of routers and NAPT's. We run the  
 447 longest chain (i.e., 10 NFs) for frame sizes in [60, 1500] (bytes). Figure 8 shows that SNF follows the  
 448 NICs' performance achieving line-rate forwarding at 40 Gbps for frames greater than 128 bytes. FastClick  
 449 catches up the line-rate performance for frame sizes greater than 800-1000 bytes.



**Figure 8.** Throughput of 10 routers and NAPT's chained using (i) FastClick and (ii) SNF versus the frame size in bytes (without CRC). The different frames are injected at 40 Gbps.

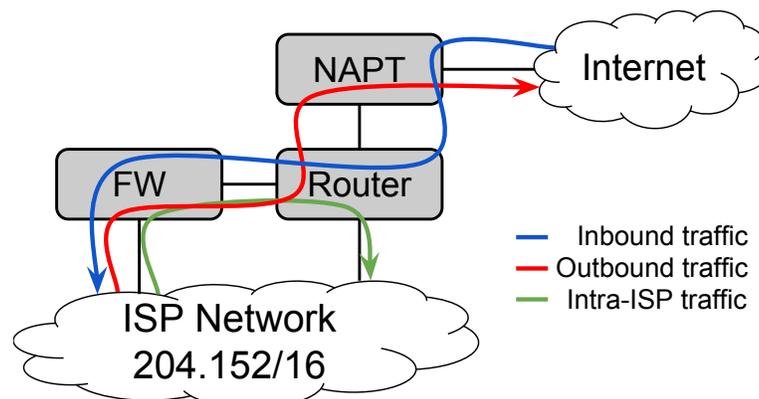
### Real Service Chain Deployments

Another common use case for an ISP is to deploy a service chain of a FW, a router, and a NAPT as depicted in Figure 9. The FW of such a chain may contain thousands of rules in its ACL causing serious performance issues for software-based NF implementations.

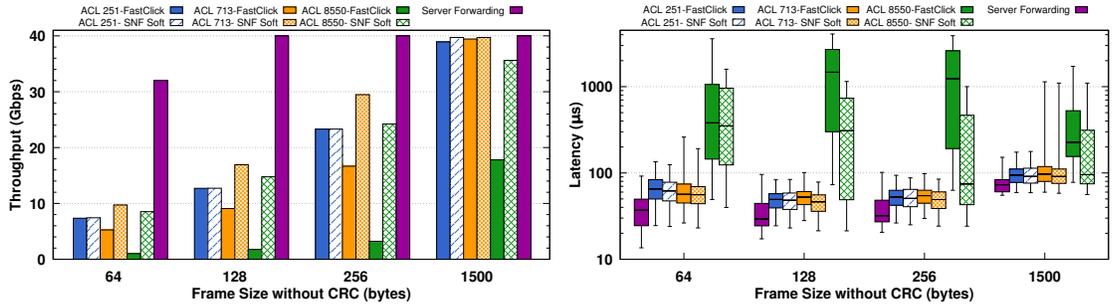
In this section we measure the performance of SNF using actual FW configurations of increasing cardinality and complexity, while exploring the limits of software-based packet processing on our hardware. We utilize a set of three actual ACLs (Taylor and Turner, 2007), taken from several ISPs, to deploy the service chain of Figure 9. The FW implements one ACL with 251, 713, or 8550 entries. The second NF is a standards-compliant IP router that redirects packets either towards the ISP's domain (intra-ISP traffic with prefix 204.152.0.0/16) or to the Internet. For the latter traffic, the third NF interconnects the ISP with the Internet by performing source and destination NAPT.

We use the above ACLs to generate traces of 64-byte frames that systematically exercise all of their entries. The generated packets emulate intra-ISP, inbound and outbound Internet traffic (see Figure 9). Figure 10 presents the performance of the 3 chains versus the different frames sizes (64, 128, 256, and 1500 bytes). We implemented the chains in FastClick and a purely software-based SNF using the full capacity of our processor's socket (i.e., 8 cores in one machine), symmetric RSS, and a batch size of 128 packets.

Figure 10a shows that the small ACL (251 rules), executed as a single FastClick instance, achieves satisfactory throughput, equal to its synthesized counterpart. This indicates that a small ISP or a chain



**Figure 9.** An ISP's service chain that serves inbound and outbound Internet traffic as well as intra-ISP traffic using three NFs.



(a) Throughput (Gbps)

(b) Latency ( $\mu\text{s}$ ) in logarithmic scale. The lower and upper percentiles are 1% and 99% respectively.

**Figure 10.** System’s performance versus 4 frame sizes (64, 128, 256, and 1500 bytes) of three different ISP-level chains with 251, 713, and 8550 rules in their ACLs. FastClick and SNF implement these chains in software using 8 CPU cores (in a single machine with four NICs), symmetric RSS, and batch size of 128 packets. Input rates are 40 Gbps for the throughput test and 5 Gbps for the latency test.

469 deployment in small subnets (e.g., using links with capacity equal or less than 10 Gbps) may not fully  
 470 benefit from SNF. As depicted in Figure 10b, the latency is also bounded below  $100 \mu\text{s}$ . This time is  
 471 dominated by the fact that our traffic flows as follows: traffic originating from one machine enters an SNF  
 472 server and, after being processed, sent back to the origin server. We believe that the observed latency  
 473 values are realistic for such a topology.

474 However, for the ACLs with 713 and 8550 rules the combination of all possible traffic classes  
 475 among the FW, router, and NAPT boxes causes the classification tree of the chain to explode in size,  
 476 hence **synthesis is a powerful yet necessary solution**. This causes three problems to FastClick: (i)  
 477 the throughput when executing the last two ACLs (713, and 8550 rules) is reduced by almost 1.5x-10x  
 478 respectively (on average), (ii) the median latency of the largest ACL is at least an order of magnitude  
 479 greater than the median latencies of the smaller ACLs (see Figure 10b), and consequently (iii) the 99<sup>th</sup>  
 480 percentile of the latency increases (up to almost 4 ms).

481 In contrast, SNF effectively synthesizes the large ACLs (i.e., 713 and 8550 rules) maintaining high  
 482 throughput despite their increasing complexity. In the case of 713 rules, the synthesis is so effective that  
 483 leads to better throughput than the 251-rule case. Regarding latency, SNF demonstrates 1.1-10x lower  
 484 median latency (bounded below  $500 \mu\text{s}$ ) and 2-3.5x lower latency variance (slightly above 1 ms in some  
 485 cases). The throughput gain of SNF is up to 8.5x greater than the FastClick chains.

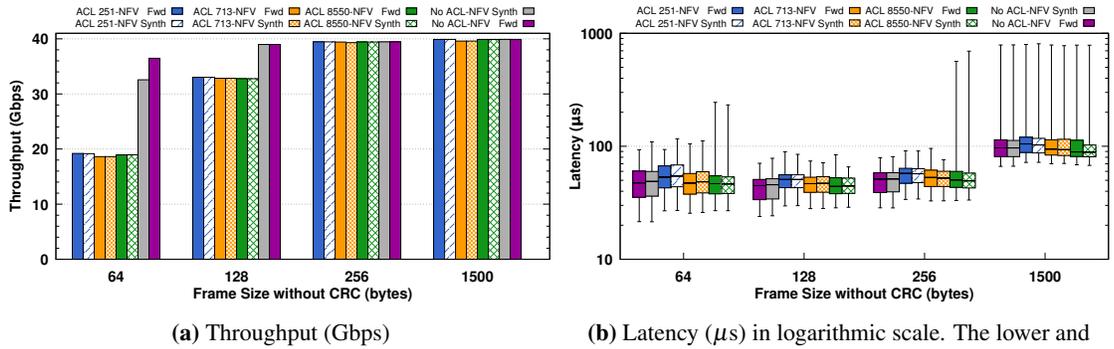
#### 486 Hardware-accelerated SNF

487 The results presented in the previous section show that software-based SNF cannot handle packet  
 488 processing at a high enough rate when the NFs are complex. We analyzed the root cause and concluded  
 489 that the packet classifier (that dispatches incoming packets to synthesized NFs) is the bottleneck. To  
 490 overcome this problem, we run additional experiments, in which we offload packet classification to a  
 491 hardware OpenFlow switch (since commodity NICs do not offer sufficient programmability). By doing  
 492 so, we showcase SNF’s ability to scale to high data rates with realistic NFs. In addition, we hint at the  
 493 performance that is potentially achievable by offloading packet classification to a programmable interface.

#### 494 Throughput Measurements

495 This extended version of SNF includes a script that converts the classification rules computed by the  
 496 original SNF to OpenFlow 1.3 rules. The translation is not straightforward because the switch rules are  
 497 less expressive than the ones accepted by the NFs. Specifically, rules that match on TCP and UDP port  
 498 ranges are problematic. While OpenFlow does allow only matches on concrete values of ports, naive  
 499 unrolling of ranges into multiple OpenFlow matches leads to an unacceptable number of rules. Instead,  
 500 we solve the problem by utilizing a pipeline of flow tables available in the switch. The first two tables  
 501 match only on the source and destination ports respectively, assign them to ranges, and write metadata that  
 502 defines the range. Further tables include the real ACL rules and also match on the metadata previously  
 503 added to a packet. Moreover, since the rules in the NFs are explored in the top-to-bottom order, we  
 504 emulate the same behavior by assigning decreasing priorities to the OpenFlow rules.

505 We use the same sets of ACLs as before, and evaluate throughput and latency in the hardware-  
 506 accelerated SNF. We first measure the throughput that SNF can achieve leveraging OpenFlow classification.



**Figure 11.** Hardware-assisted SNF’s performance versus 4 frame sizes (64, 128, 256, and 1500 bytes) of three different ISP-level chains with 251, 713, and 8550 rules in their ACLs. SNF’s classification is offloaded to an OpenFlow switch, while stateful processing occurs in 4 servers connected to the switch. Input rates are 40 Gbps for the throughput test and 5 Gbps for the latency test.

507 We design an experiment where two machines use a total of four 10 Gbps links to send traffic. The packets  
 508 are crafted so that they uniformly exercise all visible classification rules (some rules from the original  
 509 data set are fully covered by other rules). We use the same frame sizes as in § 6.4. The switch classifies  
 510 the packets and forwards them across four SNF servers that are using 10 Gbps links to connect to the  
 511 switch. The servers work in two modes: (i) forward only, where they do not implement any NFs and  
 512 simply forward packets (the first bar in each pair in Figure 11a), and (ii) synthesized mode, where they  
 513 implement the real NF chain (the second bar in each pair in Figure 11a). Additionally, for comparison,  
 514 we created an experiment where the switch installs only four basic classification rules (to do simple  
 515 forwarding) to measure the performance of the NFs themselves (the last pair of bars in Figure 11a).

516 We observe that throughput depends mostly on the frame size. The system can operate at almost 20  
 517 Gbps for small frames (i.e., 64 bytes), and it reaches the full line-rate for 256-byte frames. Interestingly,  
 518 the rule set size does not affect the throughput.

519 In the real data sets, the second bar in each pair is almost as high as the first one, which shows that the  
 520 *software part of SNF does not limit the performance*. Finally, with simple forwarding rules in the switch  
 521 (the first pair of bars in Figure 11a) the overall throughput is high even for small frames, which confirms  
 522 that packet processing at the switch is the bottleneck of the whole system. To further prove this point, we  
 523 run an experiment with only 2 ports sending traffic at an aggregate speed of 20 Gbps. In this case, SNF  
 524 processes packets at the line-rate except for the smallest frames, where it achieves 15 Gbps.

### 525 **Latency Measurements**

526 A middlebox chain should induce low, bounded packet processing delays. In this set of experiments,  
 527 we send traffic at a lower rate and measure latency. The setup is the same as in the previous scenario.  
 528 Thus, the latency we show includes the time for frames to be: (i) transmitted out of the network interface  
 529 of the traffic generating machines, (ii) received, processed, and forwarded by the OpenFlow switch, (iii)  
 530 received, processed, and forwarded by the SNF machines, and (iv) received by the destination server (the  
 531 same machine as the sender).

532 Figure 11b shows the latency depending on the frame size and the synthesized function (results for  
 533 the input rate of 20 Gbps are very similar). Our results show that the median latencies are low and stable  
 534 across all frame sizes and chains. There are several main observations here. First, the 75<sup>th</sup> percentiles  
 535 (marked by the top horizontal line of the boxplots) are close to the median latencies and we find this result  
 536 to be encouraging. Second, large frames (i.e., 1500 bytes) face two times greater median latency than the  
 537 smaller ones regardless of the rule configuration. Third, there are outliers that are an order of magnitude  
 538 less/greater than the medians (e.g., 10  $\mu$ s at the 1<sup>st</sup> and 100  $\mu$ s at 99<sup>th</sup> percentiles for 64-byte frames and  
 539 80  $\mu$ s at the 1<sup>st</sup> and 800  $\mu$ s at 99<sup>th</sup> percentiles for MTU-sized frames). Part of this latency variance is  
 540 due to the batch I/O and processing techniques of the FastClick framework; as shown in Figure 11, these  
 541 techniques offer high throughput, but have a well-studied effect on the latency variance.

## 542 VERIFICATION

543 In this section we discuss possible tools that could be utilized to *systematically* verify the correctness  
544 of the synthesis proposed by SNF.

545 Recent efforts have employed model checking (Canini et al., 2012; Kim et al., 2015a) techniques  
546 to explore the (voluminous) state space of modern networked systems in an attempt to find state  
547 inconsistencies due to etc. bugs or misconfigurations. Symbolic execution has also been utilized either  
548 alone (Kuzniar et al., 2012; Dobrescu and Argyraki, 2014) or combined with model checking (Canini  
549 et al., 2012), to systematically identify representative input events (i.e., packets) that can adequately  
550 exercise code paths without requiring to exhaust the input space (hence bound the verification time).

551 Specifically, Software Dataplane Verification (Dobrescu and Argyraki, 2014) is a close fit for verifying  
552 NFV service chains. The authors proposed a scalable approach to verifying complex NFV pipelines, by  
553 verifying each internal element of the pipeline in isolation; then by composing the results the authors  
554 proved certain properties about the entire pipeline. One could use this tool to systematically verify a  
555 complex part of SNF, which is the traffic classification. However, this tool might not be able to provide  
556 sound proofs regarding all the stateful modifications of SNF, since the authors verified only two simple  
557 stateful cases (i.e., a NAT and a traffic monitor) and did not generalize their ideas for a broader list of  
558 NFV flow modification elements.

559 SOFT (Kuzniar et al., 2012) could also be employed to test the interoperability between a chain  
560 realized with and without SNF. In other words, SOFT could inject a broad set of inputs to test whether  
561 the SynthesizedNF defined in § 3.1.3 outputs packets that are identical with the packets delivered by the  
562 original set of NFs. Similarly, HSA (Kazemian et al., 2012) could be used to statically verify loop-freedom,  
563 slice isolation, and reachability properties of SNF service chains. Unfortunately, HSA statically operates  
564 on a snapshot of the network configuration, hence is unable to track dynamic state modifications caused  
565 by continuous events. Similarly, SOFT is a special-purpose verification engine for software-defined  
566 networking (SDN) agent implementations. Therefore, both works require significant additional effort to  
567 verify stateful NFV pipelines.

568 Finally, translating an SNF processing graph into a finite state machine understandable by Kinetic (Kim  
569 et al., 2015a) would potentially allow Kinetic to use its model checker to verify certain properties for the  
570 entire pipeline. However, Kinetic does not systematically verify the actual code that runs in the network,  
571 but rather builds and verifies a model of this code. Therefore, we are concerned (i) whether a Kinetic  
572 model can sufficiently cover complex service chains such as the ISP-level chains presented in § 6.4 and  
573 (ii) whether Kinetic's located packet equivalence classes (LPECs) can handle the complex TCUs of SNF  
574 without causing state space explosion.

575 To summarize, although the works above have provided remarkable advancements in software  
576 verification, a substantial amount of additional research is required to provide strong guarantees about the  
577 correctness of SNF. For this reason, in this paper we focus our attention on delivering ultra high speed  
578 pipelines for complex and stateful NFV service chains and leave the verification of SNF as a future work.

## 579 LIMITATIONS

580 We do not attempt to provide a solution that can synthesize arbitrary software components, but rather  
581 target a broad but finite set of middlebox-specific NFs that operate on the entire space of a packet's header.  
582 SNF makes two assumptions:

- 583 1. An NFV provider must specify an NF as an ensemble of abstract packet processing elements (i.e.,  
584 the NF DAG defined in § 3.2.1). We believe that this is a reasonable assumption, followed also  
585 by other state-of-the-art approaches such as Click, Slick, and OpenBox. However, if a middlebox  
586 provider does not want to share this information under non-disclosure or via a licensing agreement,  
587 SNF can synthesize the middleboxes before and after this provider's middlebox. This is possible  
588 by omitting the processing graph of this middlebox from the inputs given to the Service Chain  
589 Configurator (see § 3.2.1).
- 590 2. No further decision (i.e., read) utilizes an already rewritten field, therefore, an LB that splits traffic  
591 based on source port after a source NAPT, might not work. Similarly, in this case, SNF can exclude  
592 the LB from the synthesis.

593 Moreover, our tool does not support network-wide placement of the chain's components, but we  
594 envision SNF being integrated in controllers, such as E2 or Slick.

## 595 RELATED WORK

596 Over the last decade, there has been considerable evolution of software-based packet processing  
597 architectures that realize wireline throughputs, while providing flexible and cost effective in-cloud  
598 network processing.

599 **Monolithic middlebox implementations.** Until recently, most NFV approaches have treated NFs as  
600 monolithic entities placed at arbitrary locations in the network. In this context, even with the assistance  
601 of state-of-the-art OSs, such as the Click-based (Kohler et al., 2000) ClickOS (Martins et al., 2014) as  
602 well as fast network I/O (Rizzo, 2012; DPDK, 2016) and processing (Kim et al., 2012, 2015b; Barbette  
603 et al., 2015) mechanisms, chaining more than 2 NFs leads to serious performance degradation as stated  
604 by the authors of both ClickOS and NetVM (Hwang et al., 2014). The main reason, also shown in our  
605 experiments, for this poor performance is the I/O overhead due to forwarding packets along physically  
606 remote and virtualized NFs. More recently, OpenNetVM (Zhang et al., 2016) showed that VM-based  
607 NFV deployments do not scale with increasing number of chained instances, hence opted for NFs running  
608 in lightweight Docker (Docker, 2016) containers interconnected with shared memory segments.

609 **Consolidation at the machine level.** Concentrating network processing into a single machine is a  
610 logical way to overcome the limitations stated above. CoMb (Sekar et al., 2012) consolidates middlebox-  
611 oriented flow processing into one machine, mainly at the session layer. Similarly, OpenNF (Gember-  
612 Jacobson et al., 2014) provides a programming interface to migrate NFs, which can in turn be collocated  
613 in a physical server. DPIaaS (Bremner-Barr et al., 2014) reuses the costly deep packet inspection (DPI)  
614 logic across multiple instances. RouteBricks (Dobrescu et al., 2009) exploits parallelism to scale software  
615 routers across multiple servers and cores within a single server, while PacketShader (Han et al., 2010) and  
616 NBA (Kim et al., 2015b) take advantage of cheap and powerful auxiliary hardware components such as  
617 GPUs to provide fast packet processing. All of these works only partially exploit the benefits of sharing  
618 common middlebox functionality, thus they are far from supporting optimized service chains.

619 **Consolidation at the individual function level** is the next level of composition of scalable and  
620 efficient NF deployments. In this context, Open Middleboxes (xOMB) (Anderson et al., 2012) proposes  
621 an incrementally scalable network processing pipeline based on triggers that pass the flow control from  
622 one element to another in a pipeline. The xOMB architecture allows great flexibility in sharing parts of the  
623 pipeline; however, it only targets request-oriented protocols and services, unlike our generic framework.

624 Slick (Anwer et al., 2015) operates on the same level of packet processing as SNF to compose  
625 distributed, network-wide service chains driven by a controller. Slick provides its own programming  
626 language to achieve this composition and unlike our work, it addresses placement requirements. Slick is  
627 very efficient when deploying service chains that are not necessarily collocated. However, we argue that  
628 in many cases all the NFs of a service chain need to be deployed in one machine and effectively being  
629 dispatched across cores in the same socket. Slick does not allow all the NF elements to be physically  
630 placed into a single process. Our work goes beyond Slick by trading the flexibility of placing NF elements  
631 on demand for extensive consolidation of the chain processing. Our synthesized SNF realizes such chains  
632 with zero context switching and zero redundancy of individual packet operations.

633 Very recently, Bremner-Barr et al. (2016) applied the SDN control and dataplane separation paradigm  
634 to OpenBox; a framework for network-wide deployment and management of NFs. OpenBox applications  
635 input different NF specifications to the OpenBox controller via a north-bound application programming  
636 interface. The controller communicates the NF specifications to the OpenBox Instances (OBIs) that  
637 constitute the actual dataplane, ensuring smart NF placement and scaling. An interesting feature of the  
638 OpenBox controller is its ability to merge different processing graphs, from different NFs, into a single  
639 and shorter processing graph, similar to our SNF. The authors of OpenBox made a similar observation  
640 with us regarding the need to classify the traffic of a service chain only once, and then apply a set of  
641 operations that originate from the different NFs of the chain.

642 However, OpenBox does not highly optimize the result chain-level processing graph for two reasons:

643 (i) The OpenBox merge algorithm can only merge homogeneous packet modification elements (i.e.,  
644 elements with the same type). For example, two “Decrement IP TTL” elements, that each decrements  
645 the TTL field by one, can be merged into a single element that directly decrements the TTL field by two.  
646 Imagine, however, the case where OpenBox has to merge the NFs of Figure 5. In this example, OpenBox  
647 cannot merge the “Rewrite Flow” element (that modifies the source and destination IP addresses as well  
648 as the source port of UDP packets) with the 3 “Decrement IP TTL” elements, since these elements do not  
649 belong to the same type. This means that the final OpenBox graph will have 2 distinct packet modification

elements (i.e., 1 “Rewrite Flow” and 1 “Decrement IP TTL”) and each element has to compute the IP and UDP checksums separately. Therefore, OpenBox does not completely eliminate redundant operations. In contrast, SNF effectively synthesized the operations of all these elements into a single element (see Figure 6) that computes the IP and UDP checksums only once. Consequently, SNF produces both a shorter processing graph and a synthesized chain with *no redundancy*, hence achieving lower latency.

(ii) Although OpenBox can merge the classification elements of a chain into a single classifier, the authors have not addressed how they handle the increased complexity of the final classifier. Our preliminary experiments showed that in complex use cases, such as the ISP-level traffic classification presented in § 6.4, the complexity of the chain-level classifier dramatically increases with increasing number of ACL rules. Therefore, SNF implements the lazy subtraction optimization proposed by Kazemian et al. (2012). The benefits of this algorithm are stated in § 5.1.

Finally, the authors of OpenBox did not stress the limits of the OpenBox framework in their performance evaluation. An input packet rate of 1-2 Gbps cannot adequately stress the memory utilization of the OBIs. Moreover, there is limited discussion related to how OpenBox exploits the multi-core capacities of modern NFV infrastructures. In contrast, in § 6.2, § 6.3, and § 6.4 we demonstrated how SNF realizes complex, purely software-based service chains at 40 Gbps line-rate. This is possible by exploiting multiple CPU cores and by fitting most of the data of an entire service chain into those cores’ L1 caches.

**Scheduling NFs for high throughput.** Recently, the E2 NFV framework (Palkar et al., 2015) demonstrated a scalable way of deploying NFV services. E2 mainly tackles placement, elastic scaling, and service composition by introducing pipelets. A pipelet defines a traffic class and a corresponding DAG of NFs that should process this traffic class. SNF’s TCUs are somewhat similar to E2’s pipelets but SNF aims to make them more efficient. Concretely, an SNF TCU is not processed by a DAG of NFs, but rather by a highly optimized piece of code (produced by the synthesizer) that directly applies a set of operations to this specific traffic class.

**Impact.** E2 can use SNF to fit more service chains into one machine, hence postpone its elastic scaling. Existing approaches can transparently use our extensions to provide services such as (i) lightweight Xen VMs that run synthesized ClickOS instances using the netmap network I/O, (ii) parallelized service chains using the multi-server, multi-core RouteBricks architecture, and (iii) synthesized chains that are load balanced across heterogeneous hardware components (i.e., CPU and GPU) using NBA.

## CONCLUSION

We have addressed the problem of synthesizing chains of NFs with SNF. SNF requires minimal I/O interactions with the NFV platform and applies single-read-single-write operations on the packets, while early-discarding irrelevant traffic classes. SNF maintains state across NFs. To realize the above properties, we parse the chained NFs and build a classification graph whose leaves represent unique traffic class units. In each leaf we perform a set of packet header modifications to generate an equivalent configuration that implements the same functionality as the initial chain using a minimal set of elements.

SNF synthesizes stateful chains that appear in production ISP-level networks realizing high throughput and low latency, while outperforming state-of-the-arts works.

## REFERENCES

- Anderson, J. W., Braud, R., Kapoor, R., Porter, G., and Vahdat, A. (2012). xOMB: Extensible Open Middleboxes with Commodity Servers. In *Proceedings of the Eighth ACM/IEEE Symposium on Architectures for Networking and Communications Systems*, ANCS ’12, pages 49–60, New York, NY, USA. ACM.
- Anwer, B., Benson, T., Feamster, N., and Levin, D. (2015). Programming Slick Network Functions. In *Proceedings of the 1st ACM SIGCOMM Symposium on Software Defined Networking Research*, SOSR ’15, pages 14:1–14:13, New York, NY, USA. ACM.
- Bagnulo, M., Matthews, P., and van Beijnum, I. (2011). Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers. RFC 6146 (Proposed Standard).
- Barbette, T., Soldani, C., and Mathy, L. (2015). Fast Userspace Packet Processing. In *Proceedings of the Eleventh ACM/IEEE Symposium on Architectures for Networking and Communications Systems*, ANCS ’15, pages 5–16, Washington, DC, USA. IEEE Computer Society.
- Bremner-Barr, A., Harchol, Y., and Hay, D. (2016). OpenBox: A Software-Defined Framework for

702 Developing, Deploying, and Managing Network Functions. In *Proceedings of the 2016 Conference on*  
703 *ACM SIGCOMM 2016 Conference*, SIGCOMM '16, pages 511–524, New York, NY, USA. ACM.

704 Bremler-Barr, A., Harchol, Y., Hay, D., and Koral, Y. (2014). Deep Packet Inspection as a Service. In  
705 *Proceedings of the 10th ACM International on Conference on Emerging Networking Experiments and*  
706 *Technologies*, CoNEXT '14, pages 271–282, New York, NY, USA. ACM.

707 Canini, M., Venzano, D., Perešini, P., Kostić, D., and Rexford, J. (2012). A NICE Way to Test Openflow  
708 Applications. In *Proceedings of the 9th USENIX Conference on Networked Systems Design and*  
709 *Implementation*, NSDI'12, pages 10–10, Berkeley, CA, USA. USENIX Association.

710 Cisco (2014). Scaling NFV - The Performance Challenge. [http://blogs.cisco.com/  
711 enterprise/scaling-nfv-the-performance-challenge](http://blogs.cisco.com/enterprise/scaling-nfv-the-performance-challenge).

712 Dobrescu, M. and Argyraki, K. (2014). Software Dataplane Verification. In *Proceedings of the 11th*  
713 *USENIX Conference on Networked Systems Design and Implementation*, NSDI'14, pages 101–114,  
714 Berkeley, CA, USA. USENIX Association.

715 Dobrescu, M., Argyraki, K., Iannaccone, G., Manesh, M., and Ratnasamy, S. (2010). Controlling  
716 Parallelism in a Multicore Software Router. In *Proceedings of the Workshop on Programmable Routers*  
717 *for Extensible Services of Tomorrow*, PRESTO '10, pages 2:1–2:6, New York, NY, USA. ACM.

718 Dobrescu, M., Egi, N., Argyraki, K., Chun, B.-G., Fall, K., Iannaccone, G., Knies, A., Manesh, M., and  
719 Ratnasamy, S. (2009). RouteBricks: Exploiting Parallelism to Scale Software Routers. In *Proceedings*  
720 *of the ACM SIGOPS 22Nd Symposium on Operating Systems Principles*, SOSP '09, pages 15–28, New  
721 York, NY, USA. ACM.

722 Docker (2016). Docker Containers. <https://www.docker.com/>.

723 DPDK (2016). Data Plane Development Kit (DPDK). <http://dpdk.org>.

724 Emmerich, P., Gallenmüller, S., Raumer, D., Wohlfart, F., and Carle, G. (2015). MoonGen: A Scriptable  
725 High-Speed Packet Generator. In *Proceedings of the 2015 ACM Conference on Internet Measurement*  
726 *Conference*, IMC '15, pages 275–287, New York, NY, USA. ACM.

727 Enguehard, M. (2016). Hyper-NF: synthesizing chains of virtualized network functions. *Master*  
728 *Thesis, KTH School of Information and Communication Technology (ICT)*. [http://urn.kb.se/  
729 resolve?urn=urn%3Anbn%3Ase%3Aakth%3Adiva-180397](http://urn.kb.se/resolve?urn=urn%3Anbn%3Ase%3Aakth%3Adiva-180397).

730 European Telecommunications Standards Institute (2012). NFV Whitepaper. [https://portal.  
731 etsi.org/NFV/NFV\\_White\\_Paper.pdf](https://portal.etsi.org/NFV/NFV_White_Paper.pdf).

732 Evans, D. (2011). The internet of things: How the next evolution of the internet is changing everything.  
733 *Cisco Internet Business Solutions Group (IBSG)*, pages 1–11. [https://www.cisco.com/c/  
734 dam/en\\_us/about/ac79/docs/innov/IoT\\_IBSG\\_0411FINAL.pdf](https://www.cisco.com/c/dam/en_us/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf).

735 Gember-Jacobson, A., Viswanathan, R., Prakash, C., Grandl, R., Khalid, J., Das, S., and Akella, A.  
736 (2014). OpenNF: Enabling Innovation in Network Function Control. In *Proceedings of the 2014 ACM*  
737 *Conference on SIGCOMM*, SIGCOMM '14, pages 163–174, New York, NY, USA. ACM.

738 Han, S., Jang, K., Park, K., and Moon, S. (2010). PacketShader: A GPU-accelerated Software Router.  
739 *SIGCOMM Comput. Commun. Rev.*, 40(4):195–206.

740 Hwang, J., Ramakrishnan, K. K., and Wood, T. (2014). NetVM: High Performance and Flexible  
741 Networking Using Virtualization on Commodity Platforms. In *Proceedings of the 11th USENIX*  
742 *Conference on Networked Systems Design and Implementation*, NSDI'14, pages 445–458, Berkeley,  
743 CA, USA. USENIX Association.

744 Intel (2016). Receiver-Side Scaling (RSS). [http://www.intel.com/content/dam/support/  
745 us/en/documents/network/sb/318483001us2.pdf](http://www.intel.com/content/dam/support/us/en/documents/network/sb/318483001us2.pdf).

746 Katsikas, Georgios (2016). SNF extensions of FastClick's stateful flow processing elements. [https://  
747 github.com/gkatsikas/fastclick/tree/snf](https://github.com/gkatsikas/fastclick/tree/snf).

748 Kazemian, P., Varghese, G., and McKeown, N. (2012). Header Space Analysis: Static Checking  
749 for Networks. In *Proceedings of the 9th USENIX Conference on Networked Systems Design and*  
750 *Implementation*, NSDI'12, pages 9–9, Berkeley, CA, USA. USENIX Association.

751 Kim, H., Reich, J., Gupta, A., Shahbaz, M., Feamster, N., and Clark, R. (2015a). Kinetic: Verifiable  
752 Dynamic Network Control. In *Proceedings of the 12th USENIX Conference on Networked Systems*  
753 *Design and Implementation*, NSDI'15, pages 59–72, Berkeley, CA, USA. USENIX Association.

754 Kim, J., Huh, S., Jang, K., Park, K., and Moon, S. (2012). The Power of Batching in the Click Modular  
755 Router. In *Proceedings of the Asia-Pacific Workshop on Systems*, APSYS '12, pages 14:1–14:6, New  
756 York, NY, USA. ACM.

- 757 Kim, J., Jang, K., Lee, K., Ma, S., Shim, J., and Moon, S. (2015b). NBA (Network Balancing Act): A  
758 High-performance Packet Processing Framework for Heterogeneous Processors. In *Proceedings of the*  
759 *Tenth European Conference on Computer Systems*, EuroSys '15, pages 22:1–22:14, New York, NY,  
760 USA. ACM.
- 761 Kohler, E., Morris, R., Chen, B., Jannotti, J., and Kaashoek, M. F. (2000). The Click Modular Router.  
762 *ACM Trans. Comput. Syst.*, 18(3):263–297.
- 763 Kuzniar, M., Peresini, P., Canini, M., Venzano, D., and Kostic, D. (2012). A SOFT Way for Openflow  
764 Switch Interoperability Testing. In *Proceedings of the 8th International Conference on Emerging*  
765 *Networking Experiments and Technologies*, CoNEXT '12, pages 265–276, New York, NY, USA. ACM.
- 766 Liu, W., Li, H., Huang, O., Boucadair, M., Leymann, N., Fu, Q., Sun, Q., Pham, C., Huang, C., Zhu, J.,  
767 and He, P. (2014). Service Function Chaining (SFC) General Use Cases. Internet-Draft draft-liu-sfc-  
768 use-cases-08, IETF Secretariat. [https://tools.ietf.org/html/draft-liu-sfc-use-](https://tools.ietf.org/html/draft-liu-sfc-use-cases-08)  
769 [cases-08](https://tools.ietf.org/html/draft-liu-sfc-use-cases-08).
- 770 Martins, J., Ahmed, M., Raiciu, C., Olteanu, V., Honda, M., Bifulco, R., and Huici, F. (2014). ClickOS  
771 and the Art of Network Function Virtualization. In *Proceedings of the 11th USENIX Conference*  
772 *on Networked Systems Design and Implementation*, NSDI'14, pages 459–473, Berkeley, CA, USA.  
773 USENIX Association.
- 774 McKeown, N., Anderson, T., Balakrishnan, H., Parulkar, G., Peterson, L., Rexford, J., Shenker, S., and  
775 Turner, J. (2008). OpenFlow: Enabling Innovation in Campus Networks. *SIGCOMM Comput. Commun.*  
776 *Rev.*, 38(2):69–74.
- 777 Palkar, S., Lan, C., Han, S., Jang, K., Panda, A., Ratnasamy, S., Rizzo, L., and Shenker, S. (2015). E2:  
778 A Framework for NFV Applications. In *Proceedings of the 25th Symposium on Operating Systems*  
779 *Principles*, SOSP '15, pages 121–136, New York, NY, USA. ACM.
- 780 Penno, R., Wing, D., and Boucadair, M. (2013). PCP Support for Nested NAT Environments. Internet-  
781 Draft draft-penno-pcp-nested-nat-03, IETF Secretariat. [https://tools.ietf.org/html/](https://tools.ietf.org/html/draft-penno-pcp-nested-nat-03)  
782 [draft-penno-pcp-nested-nat-03](https://tools.ietf.org/html/draft-penno-pcp-nested-nat-03).
- 783 Perreault, S., Yamagata, I., Miyakawa, S., Nakagawa, A., and Ashida, H. (2013). Common Requirements  
784 for Carrier-Grade NATs (CGNs). RFC 6888 (Best Current Practice).
- 785 Quinn, P. and Nadeau, T. (2015). Problem Statement for Service Function Chaining. RFC 7498  
786 (Informational).
- 787 Rizzo, L. (2012). Netmap: A Novel Framework for Fast Packet I/O. In *Proceedings of the 2012 USENIX*  
788 *Conference on Annual Technical Conference*, USENIX ATC'12, pages 9–9, Berkeley, CA, USA.  
789 USENIX Association.
- 790 SDX Central (2015). Performance - Still Fueling the NFV Discussion. [https://www.sdxcentral.](https://www.sdxcentral.com/articles/contributed/vnf-performance-fueling-nfv-discussion-kelly-leblanc/2015/05/)  
791 [com/articles/contributed/vnf-performance-fueling-nfv-discussion-](https://www.sdxcentral.com/articles/contributed/vnf-performance-fueling-nfv-discussion-kelly-leblanc/2015/05/)  
792 [kelly-leblanc/2015/05/](https://www.sdxcentral.com/articles/contributed/vnf-performance-fueling-nfv-discussion-kelly-leblanc/2015/05/).
- 793 Sekar, V., Egi, N., Ratnasamy, S., Reiter, M. K., and Shi, G. (2012). Design and Implementation  
794 of a Consolidated Middlebox Architecture. In *Proceedings of the 9th USENIX Conference on*  
795 *Networked Systems Design and Implementation*, NSDI'12, pages 24–24, Berkeley, CA, USA. USENIX  
796 Association.
- 797 Sherry, J., Hasan, S., Scott, C., Krishnamurthy, A., Ratnasamy, S., and Sekar, V. (2012). Making  
798 Middleboxes Someone else's Problem: Network Processing As a Cloud Service. In *Proceedings of the*  
799 *ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for*  
800 *Computer Communication*, SIGCOMM '12, pages 13–24, New York, NY, USA. ACM.
- 801 Sun, W. and Ricci, R. (2013). Fast and Flexible: Parallel Packet Processing with GPUs and Click. In  
802 *Proceedings of the Ninth ACM/IEEE Symposium on Architectures for Networking and Communications*  
803 *Systems*, ANCS '13, pages 25–36, Piscataway, NJ, USA. IEEE Press.
- 804 Taylor, D. E. and Turner, J. S. (2007). ClassBench: A Packet Classification Benchmark. *IEEE/ACM*  
805 *Trans. Netw.*, 15(3):499–511.
- 806 Woo, S. and Park, K. (2012). Scalable TCP Session Monitoring with Symmetric Receive-side Scaling.  
807 KAIST Technical Report. pages 1–7.
- 808 Zhang, W., Liu, G., Zhang, W., Shah, N., Loppreiato, P., Todeschi, G., Ramakrishnan, K., and Wood, T.  
809 (2016). OpenNetVM: A Platform for High Performance Network Service Chains. In *Proceedings of the*  
810 *2016 ACM SIGCOMM Workshop on Hot Topics in Middleboxes and Network Function Virtualization*.  
811 ACM.