


# Automatic Parallelization of Software Network Functions

Francisco Pereira  Fernando M. V. Ramos  Luis Pedrosa 

 INESC-ID, Instituto Superior Técnico, University of Lisbon

## Abstract

Software network functions (NFs) trade-off flexibility and ease of deployment for an increased challenge of performance. The traditional way to increase NF performance is by distributing traffic to multiple CPU cores, but this poses a significant challenge: *how to parallelize an NF without breaking its semantics?* We propose Maestro, a tool that analyzes a sequential implementation of an NF and automatically generates an enhanced parallel version that carefully configures the NIC’s Receive Side Scaling mechanism to distribute traffic across cores, while preserving semantics. When possible, Maestro orchestrates a shared-nothing architecture, with each core operating independently without shared memory coordination, maximizing performance. Otherwise, Maestro choreographs a fine-grained read-write locking mechanism that optimizes operation for typical Internet traffic. We parallelized 8 software NFs and show that they generally scale-up linearly until bottlenecked by PCIe when using small packets or by 100 Gbps line-rate with typical Internet traffic. Maestro further outperforms modern hardware-based transactional memory mechanisms, even for challenging parallel-unfriendly workloads.

## 1 Introduction

With the transition of Network Functions (or NFs) from custom, fixed-function devices to software running on commodity hardware came a well known performance challenge. As line-rates kept increasing, the networking community kept proposing new tools, techniques, and architectural enhancements to overcome individual bottlenecks. User-mode frameworks, like DPDK [38], bypass the kernel, avoiding costly context switches; DDIO [41] places incoming packets directly in the CPU cache as they arrive; and NICs implement Receive Side Scaling (RSS) [72] to consistently distribute traffic across multiple CPU cores using a configurable hash-function. Despite this wealth of tools, the challenge of developing performant software at these time scales is considerable, typically requiring parallelization [30] and, with it, a deep knowledge of low-level architectural details such as cache-friendly allocation, cache-coherence-aware coordination, and a deep understanding of the RSS hashing mechanism.

Although parallelization is paramount to achieving high performance, ensuring equivalence between parallel and sequential implementations is hard [22, 35, 50, 63, 67]. Thus, we argue that *developers need not shoulder the burden of fine-grained parallelization themselves*. Much like how developers typically do not write entire code-bases in assembly language, allowing a compiler to analyze their code, extract its function-

ality, and build an assembly implementation that is equivalent in semantics, we argue that the fine-scaled parallelization of NFs should follow a similar approach. Developers should implement sequential versions of their NFs, benefiting from the inherent simplicity of testing, debugging, and updating such systems, and when deploying to production they can “compile” the NF to obtain its parallelized version.

There are two key insights supporting the solution for this challenge. Due to the increasingly pervasive use of NF frameworks amenable to symbolic execution [1, 3, 10, 13, 17, 37, 45, 52, 68, 76], the first key insight is that this technique can be used to not only analyze the NF and infer how it maintains state, but also automatically generate modified versions of it. The second key insight is that by knowing how the NF maintains its state, we can configure the RSS mechanism to send packets accessing the same state to the same core, aiming to minimize inter-core coordination in a parallel implementation, thus maximizing performance.

With these key insights in mind, we propose **Maestro**, a tool that automatically analyzes a software NF and generates a new implementation that distributes the workload across multiple cores while preserving the semantics of the sequential implementation. This analysis builds a comprehensive symbolic model of how the NF stores and accesses state, and how that state is structured around flows. Flows (also called *flowspace* [50] and *scope* [22] by prior work) describe related packets—identified through packet header fields—that the NF logically tracks as an isolated unit. A firewall, for example, often tracks TCP/UDP flows, identified by the packet 5-tuple (source and destination IPs and ports and the IP protocol number), whereas a traffic monitor may identify flows by destination IP alone. As NFs typically store state on a per-flow basis [50, 69], Maestro learns how flows are defined in the NF by extracting the constraints that define how packets access state. We then use a solver to find an RSS configuration that distributes traffic across multiple CPU cores, in such a way as to minimize costly inter-core coordination. Our tool then automatically generates a new implementation of the NF that parallelizes its operation accordingly.

When possible, Maestro generates an implementation based on a *shared-nothing architecture*, wherein RSS is configured to forward packets of the same flow to the same CPU core, completely eliminating any inter-core coordination. When the NF is not compatible with such a model, Maestro can still generate a parallel implementation where cores share state but accesses to that state are coordinated by a read-write locking mechanism that, while not as performant

as a shared-nothing architecture, can still perform well under typical (Zipfian) Internet traffic.

Maestro draws inspiration from prior work in NF analysis [50] and verification [76, 77], as well as the wisdom of a wide body of research on NF performance [25, 30, 45, 60]. We also use the lessons learned by many before us that address the challenges of *manually* parallelizing NFs, including NUMA considerations [29], configuring RSS for symmetric flow handling [74], and rebalancing load with skew [8].

Maestro handles DPDK NFs which store state using the Vigor API [76]. For these NFs to be amenable to ESE, they are implemented under some constraints, which we describe in §5. These limitations, however, pertain only to NFs given as input to Maestro, and not to the generated parallel solutions.

We evaluate the performance of Maestro by parallelizing 8 DPDK NFs. Our experimental evaluation shows that NFs that can be parallelized using the shared-nothing architecture scale linearly with the number of cores used until bottlenecked by PCIe when using small packets or by 100 Gbps line-rate with typical Internet traffic [12]. The remaining NFs that require read-write locks to maintain their semantics vary their performance with the workload. High-churn traffic—where most packets establish a new flow—requires more writing to shared state, degrading performance. Fortunately, the majority of packets in typical Internet traffic belong to a minority of flows [12], requiring less state writing and allowing more concurrency. Under this read-heavy traffic, Maestro’s lock-based parallel NFs perform comparably to a shared-nothing model. Notably, when Maestro had to resort to locking, equivalent versions of the NFs that use hardware transactional memory [54] (TM) to preserve semantics (via the Restricted Transactional Memory interface [42]) were unable to outperform our optimized locks, as we show in §6.4. We also show that NFs automatically parallelized by Maestro rival in performance with ones manually parallelized using VPP [7].

In §2, we describe the inherent challenge of parallelizing NFs, to better motivate our work. We subsequently present the main contributions of our work, describing the Maestro architecture in §3 and several key optimizations in §4. In §5 we discuss Maestro’s inherent limitations. In §6, we evaluate Maestro and the performance of the parallel NFs it generates. Finally, we describe related work in §7 and conclude with final thoughts in §8.

## 2 Why Parallelization is Hard

Ideally, one would parallelize an NF by spinning up individual instances per core, each running independently, and using the NIC to evenly distributing traffic among them. NFs, however, typically store state that persists across packets. Sharing this state among cores requires coordinating access to it, but minimizing this coordination is crucial to achieving high performance. Parallel implementations that require no state sharing among their instances (and therefore no synchronization) are called *shared-nothing*. Implementing a shared-nothing imple-

mentation of a stateful NF requires carefully configuring the NIC to distribute traffic to each core in a way that aligns with how state is structured in the NF. With such a mechanism, state is *sharded* across cores and packets accessing the same state always find themselves on the same core.

The NIC can perform this traffic distribution in hardware using the Receive-Side Scaling (RSS) mechanism [72]. This mechanism hashes packet headers using a user-defined set of fields and a hash key. The computed hash is subsequently used to direct traffic to different queues which can deliver the packets to different cores. To send, for example, packets of the same TCP flow to the same core, one would configure RSS to hash the source and destination IP addresses, and TCP/UDP ports, and the IP protocol number (*i.e.* the 5-tuple), ensuring that any two packets with the same 5-tuple will have the same hash and will end up on the same core.

This leads us to the traditional method for building parallel shared-nothing NFs: first, developers shard state in the NF, building a full understanding of how state is accessed under all circumstances. They then use this sharding solution to construct an RSS configuration that distributes traffic accordingly. This approach, however, poses three big challenges:

**1. Finding the right sharding solutions is hard.** Though some NFs simply shard on the 5-tuple, many others require a more careful approach. One common use case involves symmetrical access to state based on the 5-tuple so that incoming traffic—that has the source and destination swapped—access the same state as outgoing traffic [74]. Other NFs require a more coarse-grained partitioning: some policers and traffic monitors only use the destination addresses to index state, connection limiters may only use source addresses, and network address translators (NATs) will typically shard on the WAN’s server address and port (as all the other addresses and ports are translated). Simply sharding on the 5-tuple here would require expensive coordination (*e.g.* locks), as cores are unable to act independently.

Arriving at sharding solutions is harder than generically using locks each time state is accessed. The developer needs intricate knowledge of the NF’s semantics and internals, particularly around how state is kept and manipulated. This thought process must not only take place upon initial implementation, but also as the NF code evolves over time. Augmenting a firewall with a connection limiter feature renders the previously configured 5-tuple sharding obsolete, requiring a complete rethink of how it should be sharded.

**2. Finding the right RSS configuration is hard.** Even if we take the sharding solution for granted, configuring RSS accordingly is difficult. For trivial cases, this is just a matter of selecting the right fields to hash but more complex scenarios can require carefully crafting the RSS key. Such an approach was used in [74] to handle symmetrical TCP/UDP flows, but manually tracking the sharding constraints and finding internal symmetries in the hash key that pair with those constraints quickly becomes unmanageable. For NFs with other sharding




Figure 1: Maestro’s architecture

requirements, the problem becomes even harder. Not all sets of fields are supported by NICs [39, 40], requiring a specific RSS key that cancels out some bits to circumvents this limitation. One might even require symmetry between different interfaces (when incoming and outgoing traffic use different NICs), which requires a separate but interrelated configuration and key for each NIC. More complex NFs can shard state in ways that do not neatly fit into any common case, requiring a custom formulation which, as before, may need to be completely rethought from scratch should the NF change over time. Some cases are outright infeasible, due to inherent NIC limitations, at which point a well-placed warning could help guide developers towards better solutions.

**3. Writing performant parallel code is hard.** Even if a developer correctly shards the NF and properly configures RSS to achieve a valid shared-nothing solution, they can still be leaving performance on the table. Though shared-nothing goes a long way towards ensuring good performance, many more minute details play a further role in parallel code. Packet buffers and state must now be cache-aligned to avoid false cache-line sharing. Memory allocation must be NUMA-aware to avoid slower remote accesses across the QPI bus. Even exogenous factors like traffic skew must now be considered [8] to fully realize the potential of a parallel implementation.

Getting any of these issues wrong can stand in the way of performance, correctness, or both, but are ultimately amenable to automation. Our tool—Maestro—tackles the first challenge by analyzing how the NF keeps its state and finding the constraints that packets that need to be sent to the same core must satisfy. It further tackles the second challenge by formulating an SMT problem and using a solver to find the right RSS keys that satisfy the sharding requirements. Finally, Maestro addresses the third challenge by automatically generating a parallel implementation that is semantically equivalent to its sequential counterpart. The generated code fully handles NIC initialization and RSS configuration, cache-alignment, load-balancing, and NUMA considerations. Even when a shared-nothing approach is not possible, Maestro can still help by generating an optimized lock-based parallel implementation that uses carefully crafted read-write locks to minimize inter-core coordination with typical Internet power-law traffic.

### 3 Maestro Architecture

Maestro uses symbolic analysis to extract information on how the NF maintains state, and with it infer possible dependencies between parallel instances. This analysis is crucial to achieve synchronization-free parallelization that shards state by carefully splitting traffic among cores. How this careful orchestration of packets can be used to avoid synchronization

among parallel instances is better explained via an example.

#### 3.1 Parallelizing a firewall

Consider a firewall NF connecting a LAN and a WAN that only forwards packets from the WAN that correspond to flows started in the LAN. To keep track of ongoing flows, it stores flow information in a map. Packets from the WAN lookup flow information symmetrically relative to packets from the LAN, naturally swapping source and destination fields.

Note that not all packets need access to all entries in the map: only the ones belonging to the packet’s flow. As such, in a parallel execution, making sure that *packets of the same flow are sent to the same core*, conjoined with the fact that packets of the same core are processed sequentially, allows us to parallelize this firewall without any synchronization between its instances—a *shared-nothing* architecture.

This orchestration of packets from the same flow to the same core requires a specific RSS configuration. Not only must we send LAN packets of the same flow to the same core, but also their (symmetric) WAN responses. A configuration partially fulfilling these requirements was already found by Woo and Park [74]<sup>1</sup>. By adapting their configuration to the firewalls’ needs, we ensure that every packet that needs access to the same memory region is sent to the same core.

#### 3.2 Generalizing NF parallelization

The above parallelization process is well tailored for our firewall, but different NFs keep state in different ways, and thus require different sharding solutions. Moreover, when access to specific state precludes flow-sharding, synchronization is necessary to maintain semantics.

Maestro deals with this parallelization process automatically by using the architecture shown in Figure 1. Maestro starts by analyzing the NF using Exhaustive Symbolic Execution (ESE) [18, 45, 76] to retrieve a sound and complete model of its behavior. Then, it hands the model over to a three stage pipeline: (1) the Constraints Generator, which uses this model to analyze how the NF keeps its state and arrive at a sharding solution; then (2) the RSS configuration generator stage—for which we built a library called RS3—that uses a solver to find an RSS configuration that steers packets following the sharding rules found by the previous stage to the same core; and finally (3) the Code Generator, that generates a parallel implementation that configures the RSS accordingly and adds additional synchronization mechanisms if needed.

<sup>1</sup>Woo and Park’s solution considers only a single RSS configuration, whereas our firewall deals with two ports (LAN and WAN), each requiring independent configurations. Although their findings are transposable to this scenario, it still requires expertise from the developers.

### 3.3 Extracting the NF’s model

Maestro uses ESE to extract the complete NF’s model. This allows us to not only analyze how the NF maintains its state, but also generate modified versions of its implementation.

The extracted model is an execution tree containing all the possible code execution paths a packet can trigger. Each node on this graph is either conditional (representing a branch condition), a stateful operation (representing a call to a stateful data structure, *e.g.* a map or a vector), or packet operation (*e.g.* forwarding, dropping, etc.). Both the packet and stateful data are traced as symbols, and every node contains a list of constraints on these symbols that can be given to a solver to query their possible values under any code path.

### 3.4 Finding the sharding solution

The NF model is passed to the Constraints Generator, which is tasked with finding a sharding dependency that allows shared-nothing parallelization. The idea is to find the constraints that hold true between packets that access the same state, *i.e.* packets that must be processed on the same core. This is intrinsically tied to how the NF maintains state. For example, in a map for two operations to access the same state they must use the same *key*. By symbolically tracking how such keys are derived from packets, we reason about the constraints on packets that access common state.

**Building a stateful report.** The Constraints Generator starts by analyzing the NF’s model and builds a stateful report (SR) of all the performed stateful operations. Each SR entry specifies the operation’s name (*e.g.* `map_put`), object instance, and other relevant arguments (*e.g.* the key used), and all the possible constraints on both the received packet and other stateful data when the operation was performed (*e.g.* `map_put` was called when a UDP packet arrived from interface 0).

**Filtering entries.** After building the SR, the Constraints Generator removes all entries related to read-only objects (*e.g.* routing tables that are filled on start-up and never updated). Such read-only accesses to shared state do not require coordination among cores and need not be reasoned about. Should all accesses be read-only, the SR will be left empty and Maestro asks the Code Generator to generate a parallel implementation that uses RSS with the sole purpose of load-balancing traffic among cores (we explain the RSS mechanism in §3.5).

**Analyzing the entries.** The use of any data structure can potentially preclude a shared-nothing approach, and therefore we need to infer the conditions under which it is safe to perform stateful operations concurrently for each of them (or if no such conditions exist). We present the analysis for one of the most predominant data structures: the map [2, 50, 50, 76].

The map stores data indexed by a key. This data can be accessed via the function `map_get`, and modified with `map_put`. Two map calls access the same memory region if and only if they are given the same key. For a shared-nothing approach, packets that trigger map calls to the same instance using the same key need to be steered to the same core. This alone is,

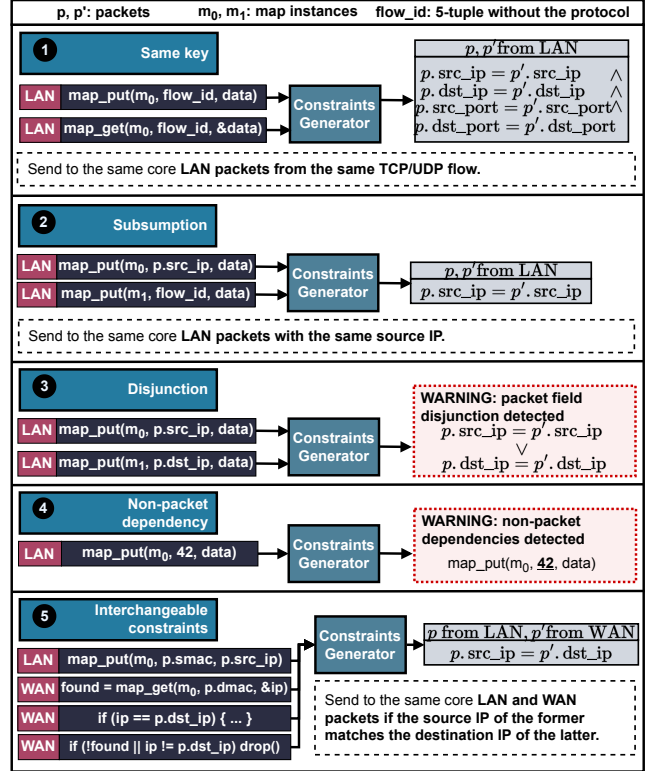


Figure 2: Example outputs of the Constraints Generator.

however, insufficient: we need to not only take into consideration any RSS limitations, but also reason about the use of multiple different map instances (or other data structures), each independently tied to the previous requirement. With this in mind, we designed a set of rules to guide Maestro towards finding correct shared-nothing sharding solutions:

- R1 Key equality.** The most obvious case is when two packets access the same map instance using the same key. In this case, the Constraints Generator builds the constraint from the formulas for the keys (1 in Figure 2).
- R2 Subsumption.** If a map instance is accessed using a subset of the packet fields used to access a second instance, then the subset takes precedence over its larger counterpart. That is, the coarser-grained requirement wins over the finer-grained one. This is exemplified in scenario 2 in Figure 2: sending packets with the same source address to the same core will also guarantee that packets with the same 5-tuple are also sent to the same core. More generally, we can always use a subset of the required packet fields. As we will see later, this rule can act further in concert with others to resolve incompatibilities.
- R3 Disjoint dependencies.** Accesses using disjoint sets of packet fields are problematic. An NF that keeps a pair of independent counters, one for source addresses and another for destination addresses, requires packets with the same source address *or* the same destination address to be sent to the same core. Due to limitations in the RSS mechanism, this is not possible: configuring it with

both the source and destination fields will guarantee that packets with the same source *and* destination will be sent to the same core. Maestro warns the user and provides the fundamental reason why the shared-nothing approach cannot be applied (3 in Figure 2).

**R4** *Incompatible dependencies.* RSS uses packet fields to steer packets to cores. This means that using keys containing (1) incompatible RSS packet fields or (2) no packet fields at all will completely block our attempt at correctly steering packets to cores. This is the case, for example, of NFs which index data with constant keys, as exemplified in case 4 of Figure 2. Again, in this case, Maestro provides feedback to the user as to why the shared-nothing approach is unfeasible<sup>2</sup>.

**R5** *Interchangeable constraints.* We define a pair of constraints as *interchangeable* if they trigger the same NF behavior. This allows us to completely replace constraints matching rules R3 or R4 with others that, if interchangeable, do not prohibit shared-nothing parallelization.

Example 5 of Figure 2 showcases this scenario. In this example, the packet is dropped when we fail to find the MAC address entry on the map, or whenever the incoming IP does not match with the stored address. Although the NF stores source addresses using an RSS-incompatible dependency on our NIC [39] (source MAC), the Constraints Generator finds that the NF’s behavior is *exactly the same* whether we shard on the MAC address or the destination address. In this case, these constraints are interchangeable, which allows Maestro to shard on either of them. Because the former uses an incompatible RSS field, the Constraints Generator opts for using the latter one for sharding.

By sharding with the IP address, changing solely this field can cause the packet to be sent to a different core. Although it may still find a matching entry of its MAC address on the map, it will find a different IP address stored on that same entry, and hence the packet will be dropped. Both not finding the MAC entry and mismatching the IP value result in the same behavior from the NF.

These rules allow Maestro to correctly find sharding solutions for a wide range of NFs (as we show in §6). Note that only R1 is specific to data structures that use a key to index state (e.g. maps, vectors, sketches). R2, R3, R4, and R5 are otherwise data structure agnostic, and Maestro employs them to all entries, regardless of their specific data structure. Though much of this analysis focuses on maps, it can be used as building blocks for others. Moreover, we need only reason about these details *once* per data-structure (or, at most, each time a breaking change is made). Once data-structure developers encode such properties into Maestro, NF developers can freely use these stateful data structures to build their NFs.

<sup>2</sup>Maestro behaves in a similar manner when finding global counters updated by every packet, as it bars it from implementing a shared-nothing parallel solution.

Name	Description
map	Stores integers indexed by arbitrary data.
vector	Stores arbitrary data indexed by integers.
dchain	Time-aware integer allocator.
sketch	Count-min sketch [21].

**Table 1: Stateful constructors currently supported by Maestro.**

Table 1 shows the stateful constructors currently supported by Maestro.

Even when Maestro fails to find a shared-nothing solution, it still provides the developer the fundamental reason why (e.g. constant keys or non-packet dependencies). When met with this result, the developer is faced with a decision: either use this feedback to tweak the NF implementation so that it becomes amenable to shared-nothing parallelism, or request a lock-based implementation from Maestro.

**Generating the constraints.** The next step in the Maestro pipeline is to generate the actual constraints, *i.e.*, the conditions that, if satisfied by a pair of packets, dictate that they must be sent to the same core. Towards this end, Maestro iterates over each pair of report entries of the same state instances, creating SMT formulas stating that both keys must be equal, and joining them all together with logical *ORs*.


Finally, we note that RSS must be independently configured on each interface. As such, the constraints generated by Maestro are interface-specific, reasoning about pairs of packets which may arrive from separate interfaces. Case 5 from Figure 2 exemplifies this. It requires LAN packets to be sent to the same core as packets from the WAN if the source address of the former equals the destination address of the latter.

Figure 3 shows the constraints found by the Constraint Generator when analyzing our firewall example. It finds that LAN packets with the same addresses and ports must be sent to the same core, and similarly for WAN packets. It also finds that WAN and LAN packets must be sent to the same core if they have the same, but swapped, sources and destinations.

### 3.5 Finding the right RSS configuration

The previous stage tackled the challenge of finding a shared-nothing sharding solution, producing constraints between packets that when true require the packets to be processed on the same core. We now focus on materializing this sharding solution by automatically finding RSS configurations that satisfy these constraints.

RSS is a hardware mechanism in the NIC that steers packets to core-specific queues. Once configured with an RSS key and a set of packet fields, it extracts from incoming packets the values of those fields and feeds them to a toeplitz-based hash-function [56]. This function, depicted in Figure 4, works by continuously left rotating the key  $k$  while iterating through the selected packet fields bits  $d$ . The running 32-bit hash value is *XOR*’ed with the current 32 least significant bits of the key whenever the current bit  $d_i$  is 1. The resulting hash is used to



**Figure 3: From the firewall’s SR to its sharding constraints.**

index an indirection table containing queue identifiers, and the packet is inserted in the corresponding queue.

Two packets with the same hash will be sent to the same core. Given the configurability of the RSS hashing function, we use it to ensure that packets that need to be processed on the same core will have the same hash. For simple constraints we can arrive at a satisfying RSS configuration solely by correctly choosing the packet field set (e.g., hashing only source and destination IPs and ports when requiring TCP packets with the same 5-tuple to be sent to the same core). However, what if (1) the NF requires a subset of packet fields that can only be used as a group in the RSS mechanism (e.g., a traffic monitor that shards solely the destination IP), (2) it requires complex constraints between packets (e.g., a Hierarchical Heavy Hitter sharding on multiple subnets of the source IP and/or source ports), or (3) there are constraints between packets arriving in different interfaces (which is the case for many NFs requiring both LAN and WAN interfaces, as in NATs, Firewalls, Connection Limiters, etc.)?

To address these scenarios in a generalized way, we built RS3, a C library capable of taking constraints as inputs and outputting RSS configurations that satisfy them. It uses the Z3 solver [23] to find suitable configurations by encoding the problem in a logical format. Maestro uses RS3 to generate RSS configurations that satisfy the constraints given by the Constraints Generator module.

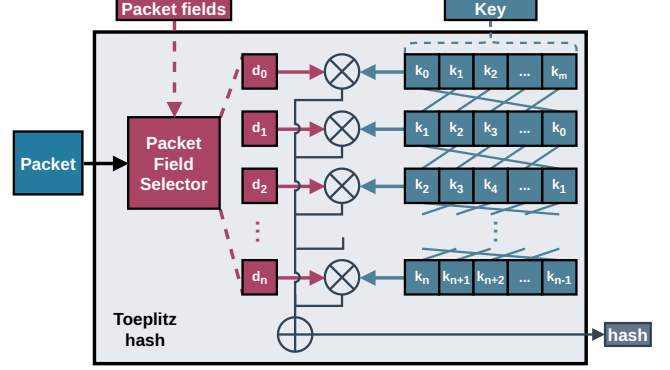
**Building the statement.** The query given to the solver needs to encode the following problem: *given set of constraints, find RSS keys that generate the same hash for every pair of packets that satisfy them.* To build this statement, we need to encode both the hash function and the constraints into an SMT format.

Let  $k$  be a 52 byte<sup>3</sup> RSS key,  $d$  and  $d'$  hash inputs for each of the packets (whose sizes depend on the extracted packet fields, e.g. 12 bytes for source and destination IPs and ports), and  $h(k, d)$  the 32 bit hash. Also, let  $|k| \geq |d| + |h|$ ,  $H(k, k', d, d')$  be true iff  $h(k, d) = h(k', d')$ , and  $C(d, d')$  be the constraint between  $d$  and  $d'$  provided by the constraint generator.

**Hash function.** As shown in Figure 4,  $H(k, k', d, d')$  can be represented as:

$$\bigwedge_{b=0}^{|h|-1} \left[ \bigoplus_{x=0}^{|d|} (d[x] \wedge k[x+b]) = \bigoplus_{y=0}^{|d'|} (d'[y] \wedge k'[y+b]) \right] \quad (1)$$

<sup>3</sup>Value for the Intel E810 100G NIC [39], but trivially adjustable in RS3.



**Figure 4: Toeplitz-based hash function.**

Note that although the size of the key is lower bounded, it should not have any influence on the feasibility of finding a suitable hash configuration. Only a subset of its bits are used on the hash function, and therefore constrained by our requirements, with the other bits being free to take any value.

**Base statement.** Initially, let us encode the following query: *find a single key  $k$  such that, given any two hash inputs  $d$  and  $d'$  that obey the constraints  $C$ , their corresponding hashes will always be equal.* That is:

$$\forall_{d, d'} \cdot k \neq 0 \wedge [(C(d, d') \wedge d \neq d') \rightarrow H(k, k, d, d')] \quad (2)$$

Having the key be 0 would always output 0 valued hashes, steering all packets to a single core, so we prevent the key from taking that value.

**Compatibility with multiple keys.** Each interface can have its RSS mechanism individually configured. With that in mind, let  $C_{ij}(d, d')$  be the constraint between a pair of packets coming from ports  $i$  and  $j$ , configured with the keys  $k_i$  and  $k_j$  respectively. Note that  $C_{ij} = C_{ji}$ , therefore it is enough to consider, for example, all the constraints  $C_{ij:\{j \leq i\}}$ . For Equation (2) to be multi-key aware, we simply conjunct the constraints across all  $i$  and  $j$ , allowing the solver to manage each key combination problem as a specific statement that must be true. That is, for  $n$  ports:

$$\forall_{d, d'} \cdot \bigwedge_{i=1}^n \bigwedge_{j=1}^i [(C_{ij}(d, d') \wedge d \neq d') \rightarrow H(k_i, k_j, d, d')] \quad (3)$$

**Compatibility with varying sets of RSS packet fields.** Just as different ports may need distinct RSS keys, we may also need to configure RSS to use different sets of packet fields depending on the interface. One way to address this would be to consider hash inputs  $d_0, \dots, d_{n-1}$  for  $n$  interfaces. This, however, highly increases the complexity of the query, making it harder for the solver to find a solution<sup>4</sup>. Another way to look at it would be to extend the hash inputs to include the union of both field-sets and to deal with any unused bits. To make the statement in Equation (3) consider constraints between packets arriving at different ports with different RSS packet field options, we again add more clauses to our large

<sup>4</sup>For  $n$  interfaces, and thus considering  $d_0, d'_0, \dots, d_{n-1}, d'_{n-1}$ , with 96 bit hash inputs we would have to deal with  $2 \times 96 \times n$  free bits.

conjunction, now considering all relevant RSS field sets, all while extracting for each one the required least significant bits of  $d$  and  $d'$  accordingly.

When given the constraints of our firewall, RS3 outputs two RSS keys, one for each NIC interface. The symmetry between the keys resembles the findings in [74], but generalized to two interfaces, rather than just one.

### 3.6 Code Generator

This stage takes the generated RSS configuration, as well as the NF's model, and outputs a parallel implementation of the original NF. Because the model is a sound and complete representation of the original NF, it can be used to generate an implementation identical in functionality to the original one. More importantly, it can be modified to employ shared-nothing parallelism by (1) configuring RSS, (2) allocating the state independently for each core, (3) making sure that each stateful call uses the data structures' instances of that particular core, and (4) launching the NF in multiple cores. Appendix A.1 contains adapted code excerpts from our Firewall example, showing both the sequential implementation used as input to Maestro and the final generated parallel shared-nothing implementation.

#### Parallel implementation with locking mechanisms.

When Maestro rules out a shared-nothing solution, it can fall back to generating parallel implementations that use locking mechanisms. In this scenario, it configures RSS with both a random key and all the available RSS-compatible packet fields, as now all cores share the same state.

Maestro also needs to carefully coordinate access to shared data using read/write locks. As such, we distinguish read-packets from write-packets: the former trigger only stateful read operations, and the latter trigger at least one write. To efficiently handle this scenario, we created a custom, highly optimized read/write lock implementation that entirely avoids cache-line sharing when acquiring read locks. We do this with a series of per-core, cache-aligned, atomic spin-locks that indicate whether the core has permission to proceed. Acquiring a read lock requires just locking the current core's lock. To perform a write, however, one must lock all core-specific locks (in order, to avoid deadlocks). With this in place, we speculatively process all packets as read-only until they attempt to perform a write operation, at which point we stop processing, release the local lock, acquire all core-specific locks, and restart processing the packet from the beginning.

The performance toll is minimized when an NF is subjected to read-heavy workloads (see §6.4), as read-only packets need only acquire a core-specific cache-aligned lock, and have no need to atomically write to any shared variable, or write to shared data. As all write-packets start out as read-packets before backtracking, starvation is not an issue.

## 4 Implementation challenges

**Finding good RSS keys.** The first set of keys found by the solver is often not ideal. If, for example, the solver finds a key

with all but the first bit set to zero, the hash, though semantically valid, will only ever be 0x0 or 0x80000000. This leads to packets being sent to only two cores.

The solution employed by RS3 involves setting the value 1 to as many bits as possible in the keys, so long as they still satisfy the given statement. This is known as a Partial MAXSAT problem [19]. We give the solver a statement that its corresponding solutions should always satisfy—Equation (3), hard constraints—and also a set of clauses that they should try to satisfy—soft constraints. The soft constraints correspond to a chain of logical *ANDs* setting each key bit to 1. There is no need for maximizing the number of satisfied soft constraints. Most of the times, a randomly selected set of bits with the value 1 is enough to avoid corner case problems like the one mentioned above. As such, Maestro uses a slightly modified version of the diagnosis-based approach introduced by Fu and Malik [33]. It begins by seeding the key with random bits. Then, if the combined hard and soft constraints are not satisfiable, we get the UNSAT core from the solver and randomly discard a subset of these soft constraints, repeating as necessary until either a key is found or no further soft constraints are left, indicating that no such key exists. Due to the randomized nature of this algorithm, we use multiple parallel solvers to independently find keys until one is found with an acceptable workload distribution.

**NUMA considerations.** In a NUMA environment, each possible combination of NIC, memory, and CPU pinning influences throughput. Our machines (see §6) have 100 Gbps NICs with 2 interfaces, thus both interfaces are pinned to the same NUMA node. Under these circumstances, pinning the packet buffers to the same NUMA node as the NIC is optimal [29].

Another important consideration is that the dominant contention factor in parallel packet processing applications is the cache, specifically for Intel Data Direct I/O (DDIO) resources [25, 55]. Using DDIO, the packets coming from the NIC are directly placed in the last level cache (LLC) of the NUMA node. Contention happens when the number of concurrent packets exceeds the available reserved space for I/O in the LLC, at which point packets evict each other and performance suffers. Maestro allocates packet buffers close to the NIC, but keeps state local to each core's NUMA node. Deciding where to run each thread is, however, a deployment challenge, not an implementation one, and therefore out of scope for Maestro. Nevertheless, our experience has taught us a simple rule of thumb: if the LLC is large enough to hold all packet buffers at line-rate, then we should pin both the CPU and memory to the same NUMA node as the NIC. If, however, the LLC is too small, resulting in contention—as occurs with older processors—then it's better to distribute cores evenly across NUMA nodes, thus increasing the total available LLC. Though we have seen scenarios where using multiple NUMA nodes was best, in our testbed the LLC proved sufficiently large to justify using a single NUMA node, and all our experiments in this paper follow this guideline.

**Traffic skew.** The expression "mice and elephants" is typically used to describe packet flow distributions on the Internet [12, 36, 53]. These follow a Zipfian distribution, where a large fraction of packets relate to but a few flows, and the remaining ones share a small slice of traffic.


While traffic with a uniform distribution leads to packets being uniformly distributed to cores, traffic following a Zipfian distribution can overload a subset of cores, causing *skew*. This performance difference is shown in Figure 5, which demonstrates how the parallel firewall throughput varies with the traffic distribution. The Zipfian traffic was generated with parameters from [60], which were found by analyzing a real-world traffic sample from a University network in [12]. This generated traffic has 50k packets and 1k flows, 48 of which responsible for 80% of the traffic. RSS was configured with five different random keys and the error bars represent the min/max performance. Performance is influenced by both the RSS key and the indirection table, as more hash collisions cause more packets being sent to the same core. Under uniform traffic, the indirection table’s entries are expected to be equally accessed, and thus uniformly filling it leads to evenly spreading packets across cores. With Zipfian traffic, however, the higher density of certain flows leads to more accesses to some entries, overloading some cores. Note that when using a single core we see better performance under Zipfian traffic due to an increased cache hit-rate when accessing state [60], though the effect is less prominent when more cores are used.

RSS++ [8] fixes the distribution problem imposed by Zipfian traffic by dynamically adjusting the indirection table according to the traffic. It balances the indirection table by swapping entries associated with overloaded cores for ones associated with underloaded ones. It also provides us with mechanisms for state migration across cores which avoid both blocking and packet reordering. We implemented static versions of these mechanisms in Maestro, but their dynamic versions could be used to handle changes in skew over time.

**State sharding.** When applying shared-nothing parallelization, Maestro not only allocates each data structure instance on each core, but further adjusts each data-structure’s capacity, keeping approximately constant the total amount of memory used for all cores by reducing the per-core amount.

This raises an interesting question about the semantics of filling up state in a shared-nothing parallel version of an NF, which slightly differs from the sequential or lock-based parallel versions. As each core now has a reduced capacity, it is possible to exhaust the capacity of one core despite there being spare room in others. Ultimately, when a core becomes “full”, it will behave in the same way locally as the sequential NF would globally (*e.g.* by dropping packets from new flows). As the RSS++ mechanism redistributes flows across cores to counteract traffic skew, this also affects state distribution, making it harder to exhaust any one core.

This state sharding has the desirable side-effect of optimizing the NF’s cache utilization. If each core has a smaller



**Figure 5: Shared-nothing firewall under uniform and Zipfian traffic, with and without balanced tables.**

working-set, more of it will fit in the local L1+L2 data caches. This provides an extra performance advantage to the shared-nothing approach on top of that of parallelization on its own.

**Lock-based rejuvenation.** When following a read-write lock-based parallelization approach, flow rejuvenation can be a challenge. As simply reading state requires updating the flow entry aging data, a naive implementation would require a write lock for all packets, with dire consequences for performance. Maestro circumvents this issue by implementing an optimized rejuvenation algorithm that operates locally in each core for most cases. We first modify the data-structures to hold multiple cache-aligned copies of the entry aging data, one per core. Each core then manages state aging locally for each entry, allowing the age of the entries to deviate from core to core as packets from the same flow arrive at different cores at different times. When eventually one core believes it should expire an entry, only then does it acquire a write lock. At this point, the core inspects the aging data for that entry on all cores. If the flow indeed expired on all cores, it is cleared out globally. If, however, another core is found where the entry has not yet expired, the local timestamp is re-synced with the newest one. Ultimately, if packets from the same flow regularly hit all cores, no write-locks are ever needed.

**Implementation.** Maestro uses the KLEE symbolic execution engine, extending it with 14,859 lines of C++ code. We also implemented RS3 in 3,964 lines of C code, independently from Maestro<sup>5</sup>.

## 5 Assumptions and limitations

**NF limitations.** To allow ESE, NFs must fit within some limitations, much like the ones enumerated in [44]: i) there must be a clean separation between stateful and stateless operations, a constraint put in practice by only allowing state to persist within a set of well-defined data structures; ii) loops must be statically bounded; and iii) no pointer arithmetic is allowed outside the data-structures. These constraints are already enforced for safety reasons in commonly used packet processing framework like eBPF<sup>6</sup> [3], a widely used framework in both academia and industry<sup>7</sup> [1, 10, 13, 17, 37, 52, 68].

**RSS limitations.** For Maestro to consider other hash function besides the standard toepitz-based one, they would have to be formulated as an SMT problem and added to RS3. This requires having their implementation openly disclosed.

<sup>5</sup>Our code is openly available at [4].

<sup>6</sup>NFs developed in eBPF store their state in kernel-maintained maps [2].

<sup>7</sup>There is an ongoing effort in adapting Maestro to also accept eBPF NFs as input (an effort already set in motion by PIX [44]).



In practice, a more limiting factor is packet field selection: shared-nothing approaches can only be applied if state is sharded using RSS-compatible packet fields. DPDK’s API [43] reference includes all possible field combinations that RSS can use (*e.g.* IPv4/IPv6 TCP/UDP flow tuples), but each NIC only implements a subset of them [39, 40].

**Attacking state sharding.** We mentioned earlier that it would be possible to “fill-up” a single core with fewer flows in a shared-nothing parallel NF than would otherwise be needed in the sequential or lock-based parallel versions. This could potentially be used as a DoS attack vector, reducing the cost for an attacker to block new flows from being admitted. RSS++ flow redistribution addresses this for well-behaved traffic, but an attacker can subvert this by specifically using flows that induce exact RSS hash collisions. Colliding flows end up on the same entry within the RSS indirection table and thus cannot be split apart.

Though out-of-scope for this paper, Maestro provides some defense from such attacks due to the randomization used to generate RSS keys. Even assuming the attacker has access to the NF source code and understands how it can be sharded across cores, different random RSS keys that comply with the sharding constraints will still distribute different flows in a different way. Without access to the actual key generated in RS3, the attacker would have a harder time reverse-engineering a set of co-located flows, mitigating their ability to induce the kind of persistent skew needed in a successful attack.

## 6 Evaluation


In this section, we evaluate Maestro and the three different types of parallel implementations it can generate: (1) *shared-nothing*, (2) *lock-based*, and (3) parallel solutions using *hardware transactional memory* [54] via the Intel’s Restricted Transactional Memory interface [42]. We aim to answer four questions: (i) how long does it take Maestro to parallelize NFs? (ii) how well does the performance of these parallel implementations scale with the number of cores? (iii) what are the impacts on performance of the various parallelization strategies that Maestro can use? and (iv) how do Maestro’s automatic parallel implementations fare against highly-optimized manually parallelized versions?

### 6.1 Target NFs and Microbenchmarks

To evaluate Maestro we analyzed 8 NFs—a simple forwarder (NOP), a policer, a bridge, a firewall (FW), a port scan detector (PSD), a NAT, a load-balancer (LB), and a connection limiter (CL). These are open-source NFs, most are non-trivial in complexity, and all have been used by a body of previous work [44, 45, 76]<sup>8</sup>. In this section, we present a brief description of each, and show how Maestro parallelizes them<sup>9</sup>. For each NF, we measured how much time Maestro took to gener-

<sup>8</sup>As mentioned in §5, the requirement that NFs be amenable to ESE can prevent Maestro from analyzing many existing codebases.

<sup>9</sup>Every automatically generated parallel solution can be found on [4].



**Figure 6: Time (in minutes) to generate parallel implementations for each NF (averaged over 10 runs).**

ate a parallel implementation (shared-nothing when possible, lock-based otherwise), summarizing the results in Figure 6.

**NOP.** This is a simple forwarding no-operation NF, *i.e.* a stateless NF that simply forwards all packets that arrive from one interface to the other. Maestro finds that this NF has no state, and provides no constraints between packets arriving at the same core. RSS is thus configured with all available packet fields and a random key on both ports.

**Policer.** This NF aims to limit each user’s download rate, identifying users by their IPv4 address. When Maestro analyzes this NF, it finds that state is indexed by the destination IP address, implying that packets with the same destination address must be sent to the same core. Because this constraint uses the destination IP address, the chosen RSS packet field options must contain this field. Although DPDK allows RSS packet field options containing only IP addresses, our NICs do not support this option. Maestro thus chooses a packet field option that includes IP addresses and TCP/UDP ports. This increases the complexity of the constraints on the key, increasing the generation time in Figure 6.

**Bridge.** A bridge associates MAC addresses with interfaces, and redirects packets accordingly. In a typical MAC learning bridge, the association between source MAC addresses and input interface is learned dynamically. When analyzing this NF, Maestro detects that state is indexed by a MAC address, which is a field not supported by RSS on our NIC. As such, Maestro warns the user that it cannot generate a shared-nothing implementation, opting for read/write locks instead.

By modifying the NF to disable dynamic MAC learning, leaving only statically configured MAC-Port bindings, the NF becomes more amenable to parallelization (as all state is read-only), albeit with reduced functionality. This further illustrates the ability of Maestro to inform developers and help guide the development process by pointing out relevant trade-offs between functionality and performance. With this in mind, we created two versions of this NF: the standard bridge with dynamic MAC learning (DBridge) and a static one with fixed bindings (SBridge). When analyzing SBridge, Maestro encounters only read-only data structures, requiring no specific constraints on the RSS configuration. As with NOP, Maestro generates a random RSS key and uses all the available packet fields on all ports.

**FW.** This is the same firewall we have been using as a running example throughout the paper (§3.1). It indexes state with typical flow information on the LAN (source and destination addresses and ports), and symmetrically on the WAN. Maestro generates a shared-nothing implementation




Figure 7: Testbed for our experiments.

that shards state by the flow information, sending WAN packets corresponding to symmetric LAN sessions to the same core as these (as shown in Figure 3).

**PSD.** A Port Scan Detector (PSD) counts how many distinct destination TCP/UDP ports each host (source IP) has touched within a given time frame. Above a threshold, connections to new ports are blocked, preventing port scans. Maestro analyzes the PSD and finds that it uses only the source IP to access one map, but also the source IP and destination port to access another. As such, the constraints for accessing the first map subsume those of the second (R2) and Maestro finds an RSS key that shards based only on source IPs.

**NAT.** A NAT translates addresses between a LAN and a WAN, allowing multiple clients in the LAN to share a single public IP in the WAN [70]. It keeps track of flows initiated in the LAN, but to aid with translation it associates a unique external port with each flow. Reply packets from the WAN are checked to see if their address and port match those on record before subsequently translating the destination address and port to match those of the client.

Maestro notices that the NAT associates flows with external ports using a map, fitting case R4 in §3.4. However, it also finds an additional constraint fitting case R5: packets from the WAN are only translated if they target the hosts that started the session in the first place. This constraint allows for sharding based on the external server’s IP address and port.

Much like its sequential implementation, the parallelized NAT enforces unique ports inside each core. It does not, however, enforce this uniqueness *across* cores, a feature that does not break semantic equivalence. Whereas on the sequential implementation the allocated ports were used to distinguish between sessions, now the sharding solution allows for packets sent to different cores (hence pertaining to different external servers) to have the same allocated ports.

**CL.** A Connection Limiter (CL) aims to limit how many connections any single client (source IP) can make to any single server (destination IP) over a wider time frame (*e.g.* several days). Given the longer time frames involved, this NF uses a memory-efficient count-min sketch [21] to estimate the connection count from each client to each server. For new connections, the source and destination IPs are used to index the sketch, indexing a configurable number of entries based on different hashes (5 by default in our case). If all entries surpass the connection limit, the packet is dropped, preventing the new connection. Otherwise, each entry is incremented.

As with the PSD, Maestro finds two different access patterns: the 5-tuple indexes a connection tracking map, while the source and destination IPs index the sketch. Again, the latter constraint subsumes the former and Maestro shards based

on source and destination IPs.

**LB.** LB is a Maglev-like load balancer [27]. Its main goal is to distribute traffic coming from the WAN to a series of identical servers on the LAN. LB registers new servers when it receives their packets coming from the LAN, and matches packets coming from the WAN with previously registered servers, keeping track of flows to ensure the same server handles packets from the same flow.

In order to maintain semantic equivalency between a shared-nothing parallel implementation and a sequential implementation, packets that find an available server in the sequential implementation must also find it available in the other. This ultimately means that all cores would need to have all backends registered in their local state. That said, packets coming in from the LAN in such a parallel implementation would only be able to be registered in a single core, preventing packets that arrive at other cores from seeing it. With this limitation in mind, it becomes impossible for multiple cores to hold an identical set of backend servers without coordination, thus preventing the use of a shared-nothing model. The Maestro analysis detects this issue when analyzing the LB SR. Lacking a better alternative, Maestro issues a warning and opts for a read/write lock based approach.


## 6.2 Performance Benchmarking Methodology

To benchmark the NFs, we use a standard testbed topology [16], connecting a traffic generator (TG) and a device under test (DUT), as shown in Figure 7. Both devices connect through a top-of-rack (TOR) switch from which we collect packet counters at the end of each experiment. Both TG and DUT are equipped with dual socket Intel Xeon Gold 6226R @ 2.90GHz, 96 GB of DRAM, and Intel E810 100 Gbps NICs. Turbo Boost, Hyper-Threading, and power saving features were disabled, as recommended by DPDK.

To measure throughput, the TG replays a given traffic sample (a PCAP file) in a loop at a given rate via the outbound cable for 10s per experiment. The DUT receives this traffic, processes it, and sends it back via the return cable, allowing the TG to measure latency. We further use the TOR to infer loss at the DUT, and—through comparison with the TG report—to also detect when packets were lost within the TG as well. We use DPDK-Pktgen [49] on the TG to find the maximum rate with less than 0.1% loss. We exclude and repeat sporadic experiment runs where loss within the TG—as opposed to the DUT—limited the results. When studying scalability, we repeatedly reevaluate the NF, while varying the number of cores it may use. We perform 10 measurements per experiment for statistical relevance and show error bars with min/max values. Our experiments properly handle NUMA considerations and indirection table rebalancing (§4).

## 6.3 Picking the Workload

In this section we analyze how different workloads impact performance, and ultimately establish the right workload con-



**Figure 8: Throughput in Gbps (blue) and Mpps (red) of the parallel NOP running on 16 cores for different packet sizes.**

figuration to evaluate all Maestro’s parallelization solutions.

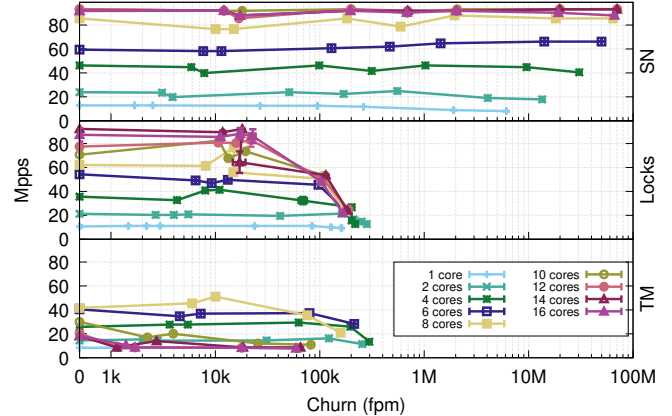
**Packet size.** To measure the impact of packet size on the performance of NFs, we ran the NOP on all cores and generated traffic with fixed-sized packets (40k uniformly distributed flows), varying the size on each iteration. The results (Figure 8) show that typical Internet traffic [12] and large packets easily achieve line-rate (100G), but that smaller packets struggle to keep up, reaching only ~45Gbps with 64B packets—even with such a trivial NF. Prior work [6, 57, 65] has pointed out that this bottleneck comes from PCIe 3.0 x16 and cannot be overcome without improved hardware. Unless stated otherwise, further experiments in this paper use 64B packets. As we measure more complex NFs that limit throughput below the 90Mpps shown in Figure 8, the bottleneck shifts from PCIe to the CPU, illustrating the NF’s intrinsic performance.

**Churn.** The performance of parallel NFs can vary significantly for read or write workloads. In networking terms, this typically relates to *churn*, or the rate at which new flows are added and expired. This is particularly important for lock and TM based implementations, where creating new flows can lead to costly aborted transactions or exclusive write locks.

We start by studying these churn effects on performance by focusing on the read/write lock-based parallel firewall, and comparing it to its shared-nothing counterpart. To conduct churn experiments, ideally one would generate traffic live that changes flows periodically in an online manner. We found it challenging to generate such traffic programmatically at line-rate so we followed an alternative solution: generating PCAPs with different levels of *relative churn*—measured in flows/Gbit. As Pktgen varies the replay rate of the PCAP to probe the NF, the resulting *absolute churn*—measured in flows/minute or fpm—changes in tandem. This guarantees that our experiments converge to an equilibrium where the highest rate is found for the given churn. Once we find this rate, we can multiply the PCAP’s relative churn with the experimental rate to compute the absolute churn.

With this in mind, we built PCAPs which (i) were small enough to fit in memory; (ii) changed enough flows to produce the desired relative churn; (iii) evenly spread these changes throughout the traffic; and (iv) were cyclic (*i.e.* the flows that expire at the start of the PCAP are created at the end). We then replay these files in a loop for 10s as in all other experiments.

Figure 9 shows how the FW—parallelized with different approaches—scales under varying amounts of churn. As absolute churn is computed based on the achieved rate, note that it too has error bars. Under low or no churn, the lock-based FW



**Figure 9: Churn study of the shared-nothing (top), lock-based (middle), and TM (bottom) parallel firewall.**

scales well until bottlenecked by PCIe. At a churn of ~100k fpm we start observing the collapse of performance as the use of more cores just wastes more cycles busy-waiting under exclusive write locks. Under heavy churn, performance is abysmal as all cores end up contending for write locks. Note that the churn limit of an NF depends on the size of packets—Figure 9 uses 64B packets but for Internet traffic [12] the lock-based FW handles churn up to 400k fpm.

The results also show just how badly the FW parallelized with transactional memory handles churn. Although a useful tool in other domains, it proves ineffective when dealing with networked applications under churn.

The shared-nothing approach, unlike the lock-based one, suffers almost no performance variation with churn up to at least ~100M fpm, a great advantage over the lock-based implementation. Benson *et al.* [12] tell us to expect up to 6M fpm in typical data-center traffic—within the ability of our shared-nothing FW, but not the lock-based one. University networks—typically with less than 15k fpm—could easily be handled even by our lock-based FW.


We focus the rest of this evaluation on studies without churn, giving the lock and TM based approaches the benefit of the doubt and illustrating their *best-case* performance.

## 6.4 Performance benchmarks

With parallel versions of each of the above 8 NFs generated, we now evaluate their performance and scalability. By default, Maestro generates a shared-nothing implementation when possible, falling back to read/write locks otherwise. This choice can, however, be overridden, and Maestro can specifically generate parallel implementations using read/write locks and TM for any of the NFs, upon request.

**Parallelization technologies.** We now study the performance and scalability of each NF while being parallelized for each of the three approaches. As per §6.3, the workload used is composed of uniformly-distributed, read-heavy, small packets<sup>10</sup>. Figure 10 shows throughput as a function of the

<sup>10</sup>Experimental results using Zipfian traffic are shown in Appendix A.2



**Figure 10: Parallel NF scalability with uniformly-distributed, read-heavy, small packets, using a shared-nothing approach when possible, read/write locks, and TM. Maestro cannot do a shared-nothing DBridge or LB.**

number of cores. Our raw performance is comparable to measurements from other recent works [30], but we focus our attention on *scalability*. Though most NFs top out their performance before using all 16 cores due to bottlenecks in the PCIe bus or the memory controller, the takeaway here is the relative performance of the different approaches.

For all NFs where a shared-nothing approach was feasible, this option scales linearly until bottlenecked by the PCIe bus and then plateaus—an ideal outcome. The lock-based implementations—though slower than their shared-nothing counterparts when available—still scale fairly well but do not always reach the PCIe bottleneck with 16 cores<sup>11</sup>. The Policer shows what happens to these locks when writes are inevitable: as every packet must update the token bucket state, every packet requires an exclusive write lock, and performance suffers catastrophically. Fortunately, this NF can be sharded by IP address, so is amenable to the shared-nothing approach.

The benefits of state sharding (§4) become clear when we compare the shared-nothing approaches with the lock-based ones for the more state intensive NFs, *i.e.* the FW, NAT, CL, and PSD. When each core holds less state due to sharding, more of it fits in the core-local L1+L2 cache. In a shared-

<sup>11</sup>Eventually, all lock-based NFs except for the Policer and CL can reach the PCIe bottleneck using extra cores from the remote NUMA node.

nothing approach where cores work independently on different working-sets this leads to an added performance improvement due to better caching, in addition to the benefits of parallelization. As a result, performance for few ( $< 4$ ) cores can be worse than linear scalability would predict and using many cores can have an added boost in comparison. Running these experiments with a workload of only 256 flows—which fits entirely in L1 cache—nullifies this effect.

A surprising takeaway is that TM does not work well with the kinds of workloads found in more complex NFs, even in the absence of churn. For simpler NFs it performs quite well, scaling linearly with the number of cores, though still operating more slowly than both shared-nothing and lock-based alternatives. In these cases TM eventually catches up with the other approaches, albeit needing more cores to do so. However, for more complex NFs TM performs abysmally, as the likelihood of a transaction aborting increases.

Ultimately, the clear winner is the shared-nothing approach, with the best backup option consistently being our read/write locks. The PSD—our most CPU intensive NF which stands to gain the most from parallelization—performs  $19\times$  better with 16 cores than a single-core version, due to the *compound effects* of parallelization and improved cache efficiency.


Maestro does not deeply affect latency. We subjected all NFs to a 1Gbps uniform background traffic of 64B packets and collected 1000 latency probes within 10 seconds. We detected no noticeable differences on the average and tail latency values between the sequential NFs and their respective parallel implementations, regardless of the adopted parallelization strategy. Pktgen measured an average of  $12 \pm 2\mu\text{s}$  for CL and  $11 \pm 1\mu\text{s}$  for the remaining NFs.

**VPP comparison.** Finally, we compare Maestro with the Vector Packet Processing framework (VPP) [7, 31], which extends the concept of batch processing to the entire packet processing pipeline with the purpose of increasing performance by minimizing instruction cache misses. VPP follows a converse approach to Maestro: packets are processed in batches in a shared-memory parallel environment where packets can end-up on any core without regard to flows or locality. Developers must then adapt the way they implement the NF to those assumptions. This approach can require more expertise and development effort, but once NFs are built in this way the framework handles many of the low-level details.

To compare the performance of a Maestro parallelized NF with an expertly developed one for VPP, we pitch our NAT against the VPP nat44-ei with the DPDK plugin. These two NFs are the most similar we found between the VPP distribution and our corpus. We further removed a number of features from the VPP NAT to bring their implementations even closer together<sup>12</sup>.

Figure 11 shows the performance comparison between

<sup>12</sup>We removed statistical counters, disabled IPv4 checksum checking, completely removed the IPv4 reassembly feature, and finally replaced the IPv4 lookup with static forwarding.



**Figure 11: VPP and Maestro NAT comparison.**

the parallel Maestro NAT (shared-nothing and lock-based) and nat44-ei, all under uniformly distributed 64B packets. Though all approaches scale well, Maestro’s shared-nothing decisively outperforms VPP, reaching the PCIe bottleneck with 10 cores. This is due to the shared-memory design that VPP follows. A fairer comparison would be between VPP and the lock-based Maestro NAT, as both use shared-memory. Here both scale more slowly, never fully reaching the PCIe bottleneck up to 16 cores. Maestro slightly outperforms VPP. Further investigation with the perf [62] tool showed us that although the Maestro lock-based NAT and the VPP one perform very similar numbers of memory reads and writes per packet, the Maestro NAT more frequently finds the data on L1 cache (Maestro’s 55% vs. VPP’s 46%) and has to access RAM less frequently (Maestro’s 3% vs. VPP’s 4%). The key takeaway though, is that Maestro’s *automatically* parallelized NFs perform competitively with expertly developed, manually parallelized NFs, without as much of a hassle.

## 7 Related Work

**Fast packet processing.** To address the performance challenges associated with software NFs, new packet I/O frameworks were proposed [3, 15, 24, 64]. To achieve high packet processing rates these solutions explore several types of optimizations including zero-copy, kernel bypass, I/O batching, and multi-queue support [9]. VPP [7] even expands batching to the whole packet processing pipeline in order to reduce instruction cache misses. Most implementations of network functions today [28, 71, 75, 77], including those from Maestro, rely on Intel DPDK [38], a kernel-bypass packet processing framework that provides a set of software libraries and drivers for fast packet processing, providing multi-core and NUMA-aware functionalities.

**NF acceleration.** PacketMill [30] accelerates NFs by carefully managing packet metadata and performing code-optimizations across the whole network stack. Another approach to improve the performance of a software NF is to leverage the platform hardware. Previous work [26, 46, 59, 71, 73] has explored multi-core CPU architectures, showing the significant improvements they can achieve on an NF’s performance, but also the challenges involved. Papadogiannaki *et al.* [59], for instance, explored the advantages of a shared-nothing model over a lock based implementation. The goal of Maestro is to offer the advantages of parallelization to NFs, for free. Although their work focused on the most efficient utilization of available resources, we use their shared-nothing model as guidance for automated generation of parallel network functions. These solutions are *manual*, requiring

extensive expertise and painstaking effort from the developer.

De Carli *et al.* [22] proposed a concurrency model for software IDSEs that uses program analysis to infer the NF’s flow semantics, feeding that information to a software scheduler that steers packets to shared-nothing threads. Though the concepts share similarities, Maestro’s approach differs from theirs by (1) considering a wider class of NFs more generally, rather than IDSEs in particular; (2) using ESE to extract fine-grained state access patterns, as opposed to their less granular program-slicing approach; and (3) handling packet steering entirely in hardware by generating RSS configurations for NICs, avoiding the bottleneck of the software scheduler and allowing Maestro parallelized NFs to scale better.

**Flow steering.** Although some NICs support rich flow-steering configurable features [39, 58], these are orthogonal to RSS and do not replace it. Using them to assure semantic equivalence on a shared-nothing implementation may require frequently adding/deleting a large amount of rules (specially under high churn), which can heavily affect performance [47]. **NF verification and synthesis.** In recent years, verification techniques have started to be applied to network functions. Some of the most relevant work includes verification of network properties [48, 51], configurations [11, 32], and NFs [76]. More recently, the research community has started exploring synthesis approaches for SDN-based control [20], data plane programs [34, 61, 78], and BGP configurations [14, 66]. Our work fits into this line, by analyzing sequential NFs to automatically generate accelerated versions.

## 8 Conclusions

In this paper we presented Maestro, a tool to automatically parallelize sequential network functions. Maestro judiciously configures the NIC’s RSS mechanism to distribute traffic across cores, while preserving semantics, resorting to locking mechanisms only when necessary. Maestro significantly improved performance for all the NFs we analyzed—scaling-up performance linearly until hitting fundamental bottlenecks in PCIe, the memory controller, or line-rate—while reducing developer effort to the push of a button.

## Acknowledgments

We are grateful to our shepherd, Tom Barbette, and the anonymous NSDI’24 reviewers. We thank Hugo Sadok for his comments on earlier drafts of the paper. We also thank Paolo Romano and Daniel Castro for their help on handling the TM approaches. This work was supported by the European Union (ACES project, 101093126), INESC-ID (via UIDB/50021/2020), and the SALAD-Nets CMU-Portugal/FCT project (2022.15622.CMU). Francisco Pereira was supported by the FCT scholarship PRT/BD/152195/2021.

## References

- [1] Cilium project, 2023. <https://cilium.io/>.
- [2] eBPF maps, 2023. [https://prototype-kernel.readthedocs.io/en/latest/bpf/ebpf\\_maps.html](https://prototype-kernel.readthedocs.io/en/latest/bpf/ebpf_maps.html).
- [3] Express Data Path, 2023. [https://en.wikipedia.org/wiki/Express\\_Data\\_Path](https://en.wikipedia.org/wiki/Express_Data_Path).
- [4] Maestro source code, 2023. <https://github.com/snaplab-dpss/maestro/tree/nsdi24>.
- [5] Maestro’s test suit, 2023. <https://github.com/snaplab-dpss/maestro-eval/tree/nsdi24>.
- [6] Saksham Agarwal, Rachit Agarwal, Behnam Montazeri, Masoud Moshref, Khaled Elmeleegy, Luigi Rizzo, Marc Asher de Kruijf, Gautam Kumar, Sylvia Ratnasamy, David Culler, and Amin Vahdat. Understanding host interconnect congestion. In *Proceedings of the 21st ACM Workshop on Hot Topics in Networks, HotNets ’22*, page 198–204, New York, NY, USA, 2022. Association for Computing Machinery.
- [7] David Barach, Leonardo Linguaglossa, Damjan Marion, Pierre Pfister, Salvatore Pontarelli, and Dario Rossi. High-Speed Software Data Plane via Vectorized Packet Processing. *IEEE Communications Magazine*, 56(12):97–103, 2018.
- [8] Tom Barbette, Georgios P. Katsikas, Gerald Q. Maguire, and Dejan Kostić. RSS++: Load and State-Aware Receive Side Scaling. In *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies, CoNEXT ’19*, page 318–333, New York, NY, USA, 2019. Association for Computing Machinery.
- [9] Tom Barbette, Cyril Soldani, and Laurent Mathy. Fast Userspace Packet Processing. In *Proceedings of the Eleventh ACM/IEEE Symposium on Architectures for Networking and Communications Systems, ANCS ’15*, page 5–16, USA, 2015. IEEE Computer Society.
- [10] David Beckett, Jaco Joubert, and Simon Horman. ACM SIGCOMM 2018 Morning Tutorial on Host Data-plane Acceleration (HDA). *New York, NY, USA, 2018*. <https://conferences.sigcomm.org/sigcomm/2018/tutorial-hda.html>.
- [11] Ryan Beckett, Aarti Gupta, Ratul Mahajan, and David Walker. A General Approach to Network Configuration Verification. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication, SIGCOMM ’17*, page 155–168, New York, NY, USA, 2017. Association for Computing Machinery.
- [12] Theophilus Benson, Aditya Akella, and David A. Maltz. Network traffic characteristics of data centers in the wild. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement, IMC ’10*, page 267–280, New York, NY, USA, 2010. Association for Computing Machinery.
- [13] Gilberto Bertin. XDP in practice: integrating XDP into our DDoS mitigation pipeline. In *Technical Conference on Linux Networking, Netdev*, volume 2, pages 1–5. The NetDev Society, 2017.
- [14] Rudiger Birkner, Dana Drachler-Cohen, Laurent Vanbever, and Martin Vechev. Config2Spec: Mining network specifications from network configurations. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*, pages 969–984, Santa Clara, CA, February 2020. USENIX Association.
- [15] Nicola Bonelli, Stefano Giordano, and Gregorio Procissi. Network Traffic Processing With PFQ. *IEEE Journal on Selected Areas in Communications*, 34(6):1819–1833, 2016.
- [16] S. Bradner and J. McQuaid. Benchmarking Methodology for Network Interconnect Devices. RFC 2544, RFC Editor, 03 1999. <https://tools.ietf.org/rfc/rfc2544.txt>.
- [17] Marco Spaziani Brunella, Giacomo Belocchi, Marco Bonola, Salvatore Pontarelli, Giuseppe Siracusano, Giuseppe Bianchi, Aniello Cammarano, Alessandro Palumbo, Luca Petrucci, and Roberto Bifulco. hXDP: Efficient software packet processing on FPGA NICs. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*, pages 973–990. USENIX Association, November 2020.
- [18] Cristian Cadar, Daniel Dunbar, and Dawson Engler. KLEE: Unassisted and Automatic Generation of High-Coverage Tests for Complex Systems Programs. In *Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation, OSDI’08*, page 209–224, USA, 2008. USENIX Association.
- [19] Byungki Cha, Kazuo Iwama, Yahiko Kambayashi, and Shuichi Miyazaki. Local search algorithms for partial maxsat. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence and Ninth Conference on Innovative Applications of Artificial Intelligence, AAAI’97/IAAI’97*, page 263–268. AAAI Press, 1997.
- [20] Haoxian Chen, Anduo Wang, and Boon Thau Loo. Towards example-guided network synthesis. In *Proceedings of the 2nd Asia-Pacific Workshop on Networking, APNet ’18*, page 65–71, New York, NY, USA, 2018. Association for Computing Machinery.

- [21] Graham Cormode and S. Muthukrishnan. An improved data stream summary: the count-min sketch and its applications. *Journal of Algorithms*, 55(1):58–75, 2005.
- [22] Lorenzo De Carli, Robin Sommer, and Somesh Jha. Beyond pattern matching: A concurrency model for stateful deep packet inspection. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, CCS '14*, page 1378–1390, New York, NY, USA, 2014. Association for Computing Machinery.
- [23] Leonardo de Moura and Nikolaj Bjørner. Z3: An Efficient SMT Solver. In C. R. Ramakrishnan and Jakob Rehof, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, pages 337–340, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.
- [24] Luca Deri. Improving Passive Packet Capture : Beyond Device Polling. *Proceedings of SANE*, 2004:85, 2004.
- [25] Mihai Dobrescu, Katerina Argyraki, and Sylvia Ratnasamy. Toward Predictable Performance in Software Packet-Processing Platforms. In *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation, NSDI'12*, page 11, USA, 2012. USENIX Association.
- [26] Mihai Dobrescu, Norbert Egi, Katerina Argyraki, Byung-Gon Chun, Kevin Fall, Gianluca Iannaccone, Allan Knies, Maziar Manesh, and Sylvia Ratnasamy. RouteBricks: Exploiting Parallelism to Scale Software Routers. In *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles, SOSP '09*, page 15–28, New York, NY, USA, 2009. Association for Computing Machinery.
- [27] Daniel E. Eisenbud, Cheng Yi, Carlo Contavalli, Cody Smith, Roman Kononov, Eric Mann-Hielscher, Ardas Cilingiroglu, Bin Cheyney, Wentao Shang, and Jinah Dylan Hosein. Maglev: A Fast and Reliable Software Network Load Balancer. In *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation, NSDI'16*, page 523–535, USA, 2016. USENIX Association.
- [28] Paul Emmerich, Sebastian Gallenmüller, Daniel Raumer, Florian Wohlfart, and Georg Carle. MoonGen: A Scriptable High-Speed Packet Generator. In *Proceedings of the 2015 Internet Measurement Conference, IMC '15*, page 275–287, New York, NY, USA, 2015. Association for Computing Machinery.
- [29] Paul Emmerich, Maximilian Pudelko, Simon Bauer, and Georg Carle. User Space Network Drivers. In *Proceedings of the Applied Networking Research Workshop, ANRW '18*, page 91–93, New York, NY, USA, 2018. Association for Computing Machinery.
- [30] Alireza Farshin, Tom Barbette, Amir Roozbeh, Gerald Q. Maguire Jr., and Dejan Kostić. PacketMill: Toward per-Core 100-Gbps Networking. In *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '21*, page 1–17, New York, NY, USA, 2021. Association for Computing Machinery.
- [31] FD.io. Vector Packet Processing - One Terabit Software Router on Intel Xeon Scalable Processor Family Server. Technical report, 2017. <https://fd.io/docs/whitepapers/FDioVPPwhitepaperJuly2017.pdf>.
- [32] Ari Fogel, Stanley Fung, Luis Pedrosa, Meg Walraed-Sullivan, Ramesh Govindan, Ratul Mahajan, and Todd Millstein. A General Approach to Network Configuration Analysis. In *Proceedings of the 12th USENIX Conference on Networked Systems Design and Implementation, NSDI'15*, page 469–483, USA, 2015. USENIX Association.
- [33] Zhaohui Fu and Sharad Malik. On Solving the Partial MAX-SAT Problem. In Armin Biere and Carla P. Gomes, editors, *Theory and Applications of Satisfiability Testing - SAT 2006*, pages 252–265, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- [34] Xiangyu Gao, Taegyun Kim, Michael D. Wong, Divya Raghunathan, Aatish Kishan Varma, Pravein Govindan Kannan, Anirudh Sivaraman, Srinivas Narayana, and Aarti Gupta. Switch Code Generation Using Program Synthesis. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication, SIGCOMM '20*, page 44–61, New York, NY, USA, 2020. Association for Computing Machinery.
- [35] Aaron Gember-Jacobson, Raajay Viswanathan, Chaithan Prakash, Robert Grandl, Junaid Khalid, Sourav Das, and Aditya Akella. OpenNF: Enabling Innovation in Network Function Control. In *Proceedings of the 2014 ACM Conference on SIGCOMM, SIGCOMM '14*, page 163–174, New York, NY, USA, 2014. Association for Computing Machinery.
- [36] Liang Guo and I. Matta. The war between mice and elephants. In *Proceedings Ninth International Conference on Network Protocols. ICNP 2001*, pages 180–188, 2001.
- [37] Toke Høiland-Jørgensen, Jesper Dangaard Brouer, Daniel Borkmann, John Fastabend, Tom Herbert, David Ahern, and David Miller. The EXpress Data Path: Fast Programmable Packet Processing in the Operating System Kernel. In *Proceedings of the 14th International Conference on Emerging Networking EXperiments and*

- Technologies*, CoNEXT '18, page 54–66, New York, NY, USA, 2018. Association for Computing Machinery.
- [38] Intel. Data Plane Development Kit, 2010. <https://www.dpdk.org>.
- [39] Intel. *Intel® Ethernet Controller E810 Datasheet*, 10 2022. Version 2.4. <https://www.intel.com/content/www/us/en/content-details/613875/intel-ethernet-controller-e810-datasheet.html>.
- [40] Intel. *Intel® Ethernet Controller X710/XXV710/XL710 Datasheet*, 6 2022. Version 4.1. <https://www.intel.com/content/www/us/en/content-details/332464/intel-ethernet-controller-x710-xxv710-xl710-datasheet.html>.
- [41] Intel. Intel data-direct I/O technology, 2023. <https://www.intel.com/content/www/us/en/io/data-direct-i-o-technology.html>.
- [42] Intel. Restricted Transactional Memory Overview, 2023. <https://www.intel.com/content/www/us/en/docs/cpp-compiler/developer-guide-reference/2021-8/restricted-transactional-memory-overview.html>.
- [43] Intel. RSS compatible packet fields on the DPDK RSS API, 2023. [https://github.com/DPDK/dpdk/blob/4fceeed5b5e9bf04acffd66239c79d81e79260/lib/ethdev/rte\\_ethdev.h#L572](https://github.com/DPDK/dpdk/blob/4fceeed5b5e9bf04acffd66239c79d81e79260/lib/ethdev/rte_ethdev.h#L572).
- [44] Rishabh Iyer, Katerina Argyraki, and George Candea. Performance Interfaces for Network Functions. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI '22)*, pages 567–584, Renton, WA, April 2022. USENIX Association.
- [45] Rishabh Iyer, Luis Pedrosa, Arseniy Zaostrovnykh, Solal Pirelli, Katerina Argyraki, and George Candea. Performance Contracts for Software Network Functions. In *Proceedings of the 16th USENIX Conference on Networked Systems Design and Implementation*, NSDI'19, page 517–530, USA, 2019. USENIX Association.
- [46] Muhammad Jamshed, Jihyung Lee, Sangwoo Moon, Deokjin Kim, Sungryoul Lee, and Kyoungsoo Park. Kargus: A Highly-scalable Software-based Intrusion Detection System Categories and Subject Descriptors. *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, pages 317–328, 2012.
- [47] Georgios P. Katsikas, Tom Barbette, Marco Chiesa, Dejan Kostić, and Gerald Q. Maguire. What You Need to Know About (Smart) Network Interface Cards. In Oliver Hohlfeld, Andra Lutu, and Dave Levin, editors, *Passive and Active Measurement*, pages 319–336, Cham, 2021. Springer International Publishing.
- [48] Peyman Kazemian, George Varghese, and Nick McKeown. Header Space Analysis: Static Checking for Networks. In *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation*, NSDI'12, page 9, USA, 2012. USENIX Association.
- [49] Keith Wiles. Pktgen - Traffic Generator powered by DPDK, 2023. <https://github.com/pktgen/Pktgen-DPDK>.
- [50] Junaid Khalid, Aaron Gember-Jacobson, Roney Michael, Anubhavnidhi Abhashkumar, and Aditya Akella. Paving the Way for NFV: Simplifying Middle-box Modifications Using StateAlyzr. In *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation*, NSDI'16, page 239–253, USA, 2016. USENIX Association.
- [51] Ahmed Khurshid, Wenxuan Zhou, Matthew Caesar, and P. Brighten Godfrey. Veriflow: Verifying Network-Wide Invariants in Real Time. volume 42, page 467–472, New York, NY, USA, sep 2012. Association for Computing Machinery.
- [52] Marios Kogias, Rishabh Iyer, and Edouard Bugnion. Bypassing the Load Balancer without Regrets. In *Proceedings of the 11th ACM Symposium on Cloud Computing*, SoCC '20, page 193–207, New York, NY, USA, 2020. Association for Computing Machinery.
- [53] Kun-chan Lan and John Heidemann. A Measurement Study of Correlations of Internet Flow Characteristics. *Comput. Netw.*, 50(1):46–62, jan 2006.
- [54] James R Larus and Ravi Rajwar. Transactional memory. *Synthesis Lectures on Computer Architecture*, 1(1):1–226, 2007.
- [55] Antonis Manousis, Rahul Anand Sharma, Vyas Sekar, and Justine Sherry. Contention-Aware Performance Prediction For Virtualized Network Functions. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication*, SIGCOMM '20, page 270–282, New York, NY, USA, 2020. Association for Computing Machinery.
- [56] Microsoft Inc. RSS Hashing Functions, 2023. <https://learn.microsoft.com/en-us/windows-hardware/drivers/network/rss-hashing-functions>.
- [57] Rolf Neugebauer, Gianni Antichi, José Fernando Zazo, Yury Audzevich, Sergio López-Buedo, and Andrew W. Moore. Understanding PCIe Performance for End Host Networking. In *Proceedings of the 2018 Conference*



- of the ACM Special Interest Group on Data Communication, SIGCOMM '18, page 327–341, New York, NY, USA, 2018. Association for Computing Machinery.
- [58] NVIDIA. Mellanox ConnectX-5 Ethernet Adapter Card, 2020. <https://network.nvidia.com/files/doc-2020/pb-connectx-5-en-card.pdf>.
- [59] Eva Papadogiannaki, Lazaros Koromilas, Giorgos Vasiladiadis, and Sotiris Ioannidis. Efficient Software Packet Processing on Heterogeneous and Asymmetric Hardware Architectures. *IEEE/ACM Transactions on Networking*, 25(3):1593–1606, 2017.
- [60] Luis Pedrosa, Rishabh Iyer, Arseniy Zaostrovnykh, Jonas Fietz, and Katerina Argyraki. Automated Synthesis of Adversarial Workloads for Network Functions. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication, SIGCOMM '18*, page 372–385, New York, NY, USA, 2018. Association for Computing Machinery.
- [61] Francisco Pereira, Gonçalo Matos, Hugo Sadok, Daehyeok Kim, Ruben Martins, Justine Sherry, Fernando M. V. Ramos, and Luis Pedrosa. Automatic Generation of Network Function Accelerators Using Component-Based Synthesis. In *Proceedings of the Symposium on SDN Research, SOSR '22*, page 89–97, New York, NY, USA, 2022. Association for Computing Machinery.
- [62] Perf. perf: Linux profiling with performance counters, 2023. <https://perf.wiki.kernel.org>.
- [63] Shriram Rajagopalan, Dan Williams, Hani Jamjoom, and Andrew Warfield. Split/Merge: System Support for Elastic Execution in Virtual Middleboxes. In *Proceedings of the 10th USENIX Conference on Networked Systems Design and Implementation, NSDI'13*, page 227–240, USA, 2013. USENIX Association.
- [64] Luigi Rizzo. Netmap: A Novel Framework for Fast Packet I/O. In *Proceedings of the 2012 USENIX Conference on Annual Technical Conference, USENIX ATC'12*, page 9, USA, 2012. USENIX Association.
- [65] Hugo Sadok, Nirav Atre, Zhipeng Zhao, Daniel S. Berger, James C. Hoe, Aurojit Panda, Justine Sherry, and Ren Wang. Ensō: A Streaming Interface for NIC-Application Communication. In *17th USENIX Symposium on Operating Systems Design and Implementation (OSDI 23)*, pages 1005–1025, Boston, MA, July 2023. USENIX Association.
- [66] Tibor Schneider, Rüdiger Birkner, and Laurent Vanbever. Snowcap: Synthesizing Network-Wide Configuration Updates. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference, SIGCOMM '21*, page 33–49, New York, NY, USA, 2021. Association for Computing Machinery.
- [67] Justine Sherry, Peter Xiang Gao, Soumya Basu, Aurojit Panda, Arvind Krishnamurthy, Christian Maciocco, Maziar Manesh, João Martins, Sylvia Ratnasamy, Luigi Rizzo, and Scott Shenker. Rollback-recovery for middleboxes. volume 45, page 227–240, New York, NY, USA, aug 2015. Association for Computing Machinery.
- [68] Nikita Shirokov and Ranjeeth Dasineni. Open-sourcing Katran, a scalable network load balancer, 2018. <https://engineering.fb.com/2018/05/22/open-source/open-sourcing-katran-a-scalable-network-load-balancer/>.
- [69] Vishal Shrivastav. Stateful Multi-Pipelined Programmable Switches. In *Proceedings of the ACM SIGCOMM 2022 Conference, SIGCOMM '22*, page 663–676, New York, NY, USA, 2022. Association for Computing Machinery.
- [70] P. Srisuresh and K. Egevang. Traditional IP Network Address Translator (Traditional NAT). RFC 3022, RFC Editor, 01 2001. <https://www.rfc-editor.org/rfc/rfc3022>.
- [71] Martino Trevisan, Alessandro Finamore, Marco Mellia, Maurizio Munafo, and Dario Rossi. Traffic Analysis with Off-the-Shelf Hardware: Challenges and Lessons Learned. *Comm. Mag.*, 55(3):163–169, mar 2017.
- [72] Amy Viviano. Introduction to Receive Side Scaling, 2023. <https://docs.microsoft.com/en-us/windows-hardware/drivers/network/introduction-to-receive-side-scaling>.
- [73] Andreas Voellmy, Junchang Wang, Y Richard Yang, Bryan Ford, and Paul Hudak. Maple: Simplifying SDN Programming Using Algorithmic Policies. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM, SIGCOMM '13*, page 87–98, New York, NY, USA, 2013. Association for Computing Machinery.
- [74] Shinae Woo and Kyoungsoo Park. Scalable TCP session monitoring with symmetric receive-side scaling. *KAIST, Daejeon, Korea, Tech. Rep*, 144, 2012.
- [75] Xiaoban Wu, Peilong Li, Yongyi Ran, and Yan Luo. Network measurement for 100 GbE network links using multicore processors. *Future Generation Computer Systems*, 79:180–189, 2018.
- [76] Arseniy Zaostrovnykh, Solal Pirelli, Rishabh Iyer, Matteo Rizzo, Luis Pedrosa, Katerina Argyraki, and George Candea. Verifying Software Network Functions with No Verification Expertise. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles, SOSP '19*,

page 275–290, New York, NY, USA, 2019. Association for Computing Machinery.

- [77] Arseniy Zaostrovnykh, Solal Pirelli, Luis Pedrosa, Katerina Argyraki, and George Candea. A Formally Verified NAT. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication, SIGCOMM '17*, page 141–154, New York, NY, USA, 2017. Association for Computing Machinery.
- [78] Kaiyuan Zhang, Danyang Zhuo, and Arvind Krishnamurthy. Gallium: Automated Software Middlebox Offloading to Programmable Switches. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication, SIGCOMM '20*, page 283–295, New York, NY, USA, 2020. Association for Computing Machinery.

## A Appendix

### A.1 Code excerpts from Maestro

We present here the pseudo-code of the firewall NF used throughout the paper, both its sequential and parallel shared-nothing implementations. These serve to provide a sense of what the Maestro pipeline both accepts as input (Figure 12) and automatically generates as output (Figure 13). As such, we reiterate that these are *not* complete examples, but only pseudo-code, as they were shortened and simplified for clarity purposes. The complete solutions can be found on our public GitHub repository [4].

Notice the symmetry of the RSS hashes (lines 7 to 25 in Figure 13), as it is what ultimately enables its shared-nothing approach. As explained in §6.1, this symmetry allows packets coming from the WAN to be sent to the same core as their corresponding symmetric packets from the LAN.

### A.2 Macrobenchmarks with Zipfian traffic

While in Figure 10 we show how throughput varies for different parallelization techniques under uniform traffic, here we repeat the experiment with Zipfian traffic instead [12] (we describe this Zipfian traffic in §4). We balanced the indirection table for each implementation to better handle the skew, as described in §4. The results are shown in Figure 14.


```
01: struct Map *flows;
02:
03: #define LAN 0
04: #define WAN 1
05:
06: // Run once.
07: int init() {
08:     map_init(flows, 65536);
09: }
10:
11: // Run by each packet.
12: void process(int port, pkt_t pkt) {
13:     if (port != WAN) {
14:         flow_t flow = {
15:             pkt.src_ip, pkt.dst_ip,
16:             pkt.src_port, pkt.dst_port
17:         };
18:
19:         map_put(flows, flow, port);
20:         forward(WAN);
21:     } else {
22:         flow_t inv_flow = {
23:             pkt.dst_ip, pkt.src_ip,
24:             pkt.dst_port, pkt.src_port
25:         };
26:
27:         int out_port;
28:         bool found = map_get(
29:             flows, inv_flow, &out_port);
30:
31:         if (!found) {
32:             drop();
33:         } else {
34:             forward(out_port);
35:         }
36:     }
37: }
```

Figure 12: Pseudo-code of the sequential firewall used as an example throughout the paper.

```
01: // One map for each thread.
02: struct Map** flows;
03:
04: #define LAN 0
05: #define WAN 1
06:
07: uint8_t RSS_HASH_PORT_0[52] = {
08:     0xa1, 0x24, 0x00, 0x15, 0x00, 0x14, 0xa1, 0x24,
09:     0xa1, 0x24, 0x00, 0x14, 0xa1, 0x24, 0x00, 0x15,
10:     0xa7, 0xfa, 0x11, 0x22, 0x6f, 0xd3, 0xf0, 0x42,
11:     0x1b, 0x6c, 0xeb, 0x14, 0x62, 0x02, 0xa3, 0x44,
12:     0x23, 0x90, 0xfb, 0x1c, 0x43, 0x99, 0xe7, 0xaf,
13:     0x80, 0x73, 0x15, 0xfe, 0x29, 0x5a, 0x73, 0xd0,
14:     0x55, 0x85, 0xf2, 0xc4
15: };
16:
17: uint8_t RSS_HASH_PORT_1[52] = {
18:     0x00, 0x14, 0xa1, 0x24, 0xa1, 0x24, 0x00, 0x15,
19:     0x00, 0x14, 0xa1, 0x24, 0x00, 0x14, 0xa1, 0x24,
20:     0x6a, 0xe3, 0xac, 0x86, 0x3e, 0xcb, 0x7e, 0x73,
21:     0x83, 0x15, 0xcb, 0x75, 0xc4, 0x73, 0x2c, 0xda,
22:     0xdb, 0x05, 0x31, 0x46, 0xdb, 0xd4, 0x76, 0x5a,
23:     0xa8, 0x20, 0x9d, 0x0a, 0x44, 0x7a, 0xc6, 0xae,
24:     0x5d, 0x72, 0x34, 0x9c
25: };
26:
27: // Run by each worker thread.
28: int init() {
29:     unsigned core_id = rte_lcore_id();
30:
31:     if (core_id == rte_get_main_lcore()) {
32:         rss_configure(
33:             LAN, RSS_HASH_PORT_0, IP_TCP | IP_UDP);
34:         rss_configure(
35:             WAN, RSS_HASH_PORT_1, IP_TCP | IP_UDP);
36:     }
37:
38:     map_init(flows[core_id], 65536);
39: }
40:
41: // Run by each packet on a specific worker thread.
42: void process(int port, pkt_t pkt) {
43:     unsigned core_id = rte_lcore_id();
44:
45:     if (port != WAN) {
46:         flow_t flow = {
47:             pkt.src_ip, pkt.dst_ip,
48:             pkt.src_port, pkt.dst_port
49:         };
50:
51:         map_put(flows[core_id], flow, port);
52:         forward(WAN);
53:     } else {
54:         flow_t inv_flow = {
55:             pkt.dst_ip, pkt.src_ip,
56:             pkt.dst_port, pkt.src_port
57:         };
58:
59:         int out_port;
60:         bool found = map_get(
61:             flows[core_id], inv_flow, &out_port);
62:
63:         if (!found) {
64:             drop();
65:         } else {
66:             forward(out_port);
67:         }
68:     }
69: }
```

Figure 13: Pseudo-code of the firewall, but now parallelized by Maestro with a shared-nothing architecture (and described in §6.1).

The key takeaways are the same as in Figure 10: when available the shared-nothing approach is always preferred; the lock-based solutions frequently do not scale as well as their



**Figure 14: Parallel NF implementation scalability with Zipfian, read-heavy, small packet traffic, using a shared-nothing approach when possible, read/write locks, and TM.**

shared-nothing alternatives and suffer in more state-intensive NFs; and TM-based approaches perform unreliably.

We do, however, find differences between these results and their counterparts under uniform traffic. Although under uniform traffic it is rather clear that throughput scales up with the number of cores when using the shared-nothing approach, with Zipfian traffic this scaling is not always consistently monotonic. This is to be expected, as the efficacy of balancing load across cores may not consistently improve when more cores are added. Indeed, when many cores are used, a single elephant flow can bottleneck a single core, limiting the maximum throughput we will be able to achieve in our experimental setup. This is particularly limiting for computationally and state intensive NFs—such as the the Connection Limiter—which are unable to perform as well with Zipfian traffic as they do with uniform. These results nevertheless confirm that Maestro generated NFs almost always perform as well with Zipfian traffic as they do with uniform.

### A.3 Reproducibility

We make Maestro’s code publicly available in [4]. In that repository, one can find not only the source code for the entire pipeline, but also the complete set of NFs we mention on this paper, along with their corresponding parallel solutions found by Maestro and described in §6.1.

We also make available our test suit in [5]. It contains all the required scripts to generate Figures 5, 8 to 11 and 14. They were tested on 2 machines with dual socket Intel Xeon Gold 6226R @ 2.90GHz, 96 GB of DRAM, and e810 Intel NICs [39], and running Ubuntu 22.04.