Determination of Throughput Guarantees for Processor-based SmartNICs

Johannes Krude
RWTH Aachen University
krude@comsys.rwth-aachen.de

Jan Rüth
RWTH Aachen University
rueth@comsys.rwth-aachen.de

Daniel Schemmel
RWTH Aachen University
schemmel@comsys.rwth-aachen.de

Felix Rath
RWTH Aachen University
rath@comsys.rwth-aachen.de

Iohannes-Heorh Folbort
RWTH Aachen University
iohannes-heorh.folbort@rwth-aachen.de

Klaus Wehrle
RWTH Aachen University
wehrle@comsys.rwth-aachen.de

ABSTRACT

Programmable network devices are on the rise with many applications ranging from improved network management to accelerating and offloading parts of distributed systems. Processor-based SmartNICs, match-action-based switches, and FPGA devices offer on-path programmability. Whereas processor-based SmartNICs are much easier and more versatile to program, they have the huge disadvantage that the resulting throughput may vary strongly and is not easily predictable even to the programmer. We want to close this gap by presenting a methodology which, given a SmartNIC program, determines the achievable throughput of this SmartNIC program in terms of achievable packet rate and bit rate. Our approach combines incremental longest path search with SMT checks to establish a lower bound for the slowest satisfiable program path. By analyzing only the slowest program paths, our approach estimates throughput bounds within a few seconds. The evaluation with our prototype on real programs shows that the estimated throughput guarantees are correct with an error of at most 1.7% and provide a tight lower bound for processor- and memory-bottlenecked programs with only 8.5% and 18.2% underestimation.

CCS CONCEPTS

• **Networks** → **Programmable networks**; *Network Performance analysis*; • **Software and its engineering** → *Formal software verification.*

KEYWORDS

BPF/XDP, SmartNIC, packet rate, bit rate, longest path search

ACM Reference Format:
Johannes Krude, Jan Rüth, Daniel Schemmel, Felix Rath, Iohannes-Heorh Folbort, and Klaus Wehrle. 2021. Determination of Throughput Guarantees for Processor-based SmartNICs. In *The 17th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '21), December 7–10, 2021, Virtual Event, Germany*. ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3485983.3494842

CoNEXT '21, December 7–10, 2021, Virtual Event, Germany

1 INTRODUCTION

Data plane programmability promises the ability to add and change functionality on general-purpose network devices. Data plane programs are used in large-scale deployments to provide functionality such as load-balancing [41], DoS-traffic-scrubbing [1], and offloading packet processing from hypervisors [19]. More examples can be found in scientific literature ranging from in-network caching [32] to offloading parts of distributed systems such as Paxos [13], and accelerating machine learning within the network [36, 58, 59].

General-purpose data plane programmability bears the risk of slow programs causing bad throughput. Therefore, match-action pipelines in programmable switches were created to process packets at a fixed packet rate [3]. Match-action pipelines, however, come at the cost of complicated programming languages and reduced expressiveness [22, 26].

Another option are FPGA-based SmartNICs, as these also allow for data plane programmability with a fixed packet rate. However, FPGA NICs cost at least 8× the price of a regular NIC and require a dedicated team of hardware experts [4, 19] to write programs in hardware description languages. FPGAs can be used to implement a processor which is then much easier to program [4] but no longer processes packets at a fixed rate and is less performant than a hardware processor.

Processors are the common target when programming and allow for rich computation and control flow. For example, the Netronome Agilio CX SmartNIC can be programmed in C using a BPF/XDP toolchain [28, 31]. Although BPF limits the number of executed instructions per packet, the resulting throughput is not obvious [28] and can greatly vary between different packets processed by the same program. Measuring the throughput with a traffic trace can give some idea about the performance of a program, but does not help in predicting the performance in case the traffic changes. We want to close this gap in providing a methodology that determines throughput guarantees for processor-based SmartNICs.

Devices such as switches and NICs have bottlenecks which can be well described in terms of achievable throughput. Whenever the rate of incoming (packet-)data exceeds the throughput bottleneck, congestion forms that induces queuing delay and packet drops that then cause bad network performance. Device-induced latency on a fully loaded SmartNIC is dominated by queuing behavior [27, 37] instead of program execution time. We focus on throughput instead of latency and present a methodology to determine a lower bound for the achievable packet and bit rate of a program.

A program developer or network operator can use our fully automated approach to derive the worst-case guaranteed throughput of a program. If this guaranteed throughput is good enough to, e.g., not cause any congestion, the program can be safely executed on the data path. In case the throughput of the analyzed program does not yet meet the intended demand, she can try a different program variant or further optimize the identified worst-case.

Throughput guarantees are related to the worst-case execution time which is a well-established field of research (see [63] for an overview) and is a hard problem for general programs on typical processors. Packet processing programs are simpler to analyze, since they typically have no unbounded loops [30, 57, 65]. Existing packet processing performance analysis work targets general purpose processors [30, 52, 54] and determine only rough estimates such as the number of executed instructions and number of memory accesses [30] or use simplifying heuristics [52, 54]. They do not identify the worst-case [52] or require exhaustive symbolic execution [30, 54] which results in unfeasibly long analysis times. We instead target a SmartNIC without memory caches, analyze throughput instead of execution time, can determine both packet rate and bit rate guarantees, and achieve short analysis time due to incremental path enumeration.

To achieve short analysis time, we only analyze the slowest program paths. However, some paths cannot be triggered by any packet and are therefore irrelevant for the achievable throughput. Our approach is based on enumerating program paths ordered from the slowest path to the fastest path and uses satisfiability checks to exclude the unsatisfiable slowest paths. With incremental enumeration, the analysis can already be stopped on the first satisfiable path without enumerating all paths, resulting in short analysis time. In case this analysis time is still too long, e.g., because of path explosion, an incrementally improving lower bound for the throughput guarantee is produced with each enumerated unsatisfiable path. If one waits until the slowest satisfiable path is identified, our approach additionally yields an example packet and memory assignment which can then be used to measure the worst-case throughput on a real deployment.

We implemented a prototype that analyzes BPF/XDP programs compiled for the Netronome Agilio CX SmartNIC. The evaluation on real programs shows that a first lower throughput bound can be determined within 23.6 s and can be improved by up to 44% within 101.9 s. Throughput measurements show an error of up to 1.7% and a tight lower bound for processor- and memory-bottlenecked programs with only 8.5% and 18.2% underestimation. Our prototype yields useful results for real programs in a timely manner.

Structure. We start by explaining the targeted SmartNIC's architecture in § 2 and subsequently give an overview on our throughput analysis approach in § 3. Then, § 4 describes the per-path throughput capacity heuristics followed by § 5 which presents our incremental ordered path enumeration approach. § 6 evaluates the accuracy and analysis time of our prototype. Finally, we discuss our approach in § 7 followed by related work in § 8 and a conclusion in § 9.

2 PROCESSOR-BASED SMARTNICS

We analyze BPF/XDP [28, 31] programs executed on the Netronome Agilio CX 2x40 GbE SmartNIC. The Netronome Flow Processor (NFP) on this NIC is similar to its predecessor, the Intel IXP network processor. Both have been investigated in previous performance works [7, 12, 23, 25, 27, 50, 53]. Our work is based on the NFP's predictable cycle costs and the program properties ensured by the BPF/XDP toolchain.

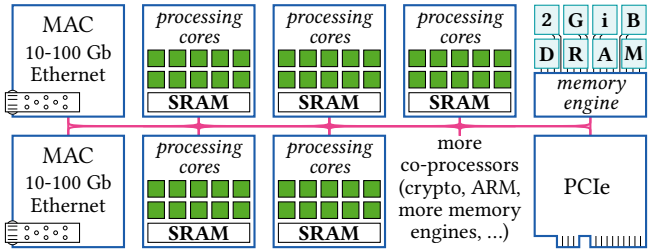

Figure 1: The Netronome Flow Processor architecture.

Islands. As shown in Figure 1, the NFP is organized into islands which communicate over a high-throughput switching fabric [31, 43, 44]. Some islands contain processing cores whereas others contain special functions such as Ethernet, PCIe, or a transactional memory engine with DRAM.

Many Simple Cores. Packet processing is parallelized onto a huge number of small cores lacking features such as branch prediction, out of order execution, and integer division [46]. Instead of caches, memory access latency is masked by cooperative hyper-threading whereby a thread may yield when waiting for a memory response.

An Explicit Memory Hierarchy. The NFP has different kinds of memory with varying access latencies [43]. Each processing core has fast access to its own instruction and data memory, medium latency when accessing the SRAM shared by all cores of its island, and some larger latency when waiting for a response from the memory engine which handles DRAM access. Unlike when using a cache hierarchy, pointers always explicitly encode which memory to access. When sending a memory request to the memory engine which handles DRAM access, transactional commands enable operations such as atomic increments without locks.

BPF/XDP on NFP. We analyze NFP Programs produced by the BPF/XDP [28, 31] toolchain since it supports high-level programming languages such as C and P4 and compiles programs to both x86_64 and NFP bytecode [31, 49]. A simplified example program is shown in Figure 2. The Linux kernel loads BPF/XDP programs onto the NIC and verifies program termination by calculating loop bounds and verifies that packet memory accesses are preceded by packet size checks [28]. The NIC's firmware [48] accepts packets over Ethernet and evenly distributes them to 50 processing cores where the BPF/XDP program is invoked for each packet. The program may modify an initial part of the packet in the island's SRAM, may access permanent state in the shared DRAM, and finally decides whether to drop a packet, to transmit it over Ethernet, or forward it over PCIe to the host.

Our goal is, given such a BPF/XDP program compiled to NFP bytecode, to determine a guaranteed throughput that the NIC will always achieve. We, therefore, estimate and compare the amount of processing and DRAM access of a program to identify the program-specific bottleneck throughput.

3 THROUGHPUT ANALYSIS

We want to establish throughput guarantees for SmartNIC programs to enable program developers and network operators to assess whether a given program on a given SmartNIC can achieve the required bit or packet rate. We do this with a fully automated analysis for a program's worst-case throughput capacity. Establishing a lower bound for the throughput capacity boils down to identifying which program path takes the longest time to execute.

Program Paths. The execution time and therefore throughput capacity fundamentally depends on the program path (i.e., the list of instructions and their execution time) that is imposed by the program's structure, the packet's as well as the memory's content. For example, when the program from Figure 2 receives an IPv4 packet of at least 100 byte, the program path through lines 2→4→5→6→8 is triggered and the throughput capacity depends on the execution time of the instructions on this path. However, looking at the example, we clearly see that it actually has four packet classes (`pkt.size` < 100, IPv4, IPv6, other). Each packet class results in a different program path and thus, different executed instructions, likely having a different throughput capacity. As such, any approach that wants to provide a lower bound on a program's throughput capacity must identify the slowest path through the instructions of a program.

Per-Path Throughput Heuristics. To identify the paths with the lowest throughput capacity, a heuristic is needed which estimates the execution time of instructions. The instruction's execution time on the processing cores is, however, not the only variable throughput limitation on the NFP. Instructions that issue memory operations to the shared DRAM may overload the memory engine. When the memory engine is overloaded, the packet throughput becomes a function of the memory engine's rate of executing memory operations. Depending on the ratio between memory and non-memory instructions, the achievable throughput of a program path is limited by either the execution time on the processing cores or the induced load on the memory engine. By using separate heuristics, the throughput capacity of the processing cores and memory engine can be independently estimated for each program path and then compared to identify the actual bottleneck.

With an overall throughput capacity number for each path, we can identify the path with the lowest throughput capacity independent of the individual path's bottleneck. In our example from Figure 2 we can therefore figure out whether the path through line 5 or the path through line 7 has the lower throughput capacity despite one being memory bottlenecked and the other being processing core bottlenecked.

Impossible Paths. When identifying paths through the program, we may encounter impossible paths that cannot be triggered by any packet. Looking at our example, the path with the highest number of executed instructions (2→4→5→6→7→8) cannot be triggered by any packet since the **`if`** conditions in lines 4 and 6 contradict. If such an impossible path is estimated to yield the lowest throughput capacity, its guarantee is not in itself wrong, however, as this execution can never occur in reality, the throughput bound may be far off from the actual (higher) lowest throughput capacity. As such, checking whether paths are possible has the potential of more closely estimating throughput guarantees.

```
1 int main(pkt) {
2     if (pkt.size < 100)
3         return XDP_DROP;
4     if (pkt[ethtype] == ETH_IPv4)
5         atomic_inc(&ip4_counter, 1);
6     if (pkt[ethtype] == ETH_IPv6)
7         for (i = 0; i < 10; i++) nop();
8     return XDP_PASS;
9 }
```

Figure 2: A simplified example of a BPF/XDP program.

Stateful Behavior. Program behavior may depend on permanent state stored in the NFP's shared DRAM. Whether a program path is possible, therefore, may depend on the content of the DRAM. By assuming, that the DRAM initially may contain any value, we analyze a broad range of program paths and establish a throughput guarantee which is valid independent of the actual DRAM content. Since BPF/XDP on NFP does not support reading and writing from the same DRAM location, the worst-case does not depend on any packet processed before or in parallel to the currently processed packet. We, therefore, do not consider sequences of packets, but only analyze a single run of the BPF/XDP program.

Packet Sizes. There are two different commonly used metrics for throughput capacity: packet rate and bit rate. Many programs process only small headers independent of the actual packet size. For those programs, a bit rate guarantee is equivalent to a packet rate guarantee multiplied by the Ethernet minimum packet size of 60 byte (without CRC). Longer packets increase the actual bit rate, but cannot be considered for a bit rate guarantee as long as the same program paths can be triggered by small packets. This changes, if the program processes longer headers (e.g., tunneling, IPv6 options) or accesses the payload. Whenever a program successfully checks the packet size to access packet data beyond the 60 byte mark, we can infer that the actual packet size is at least the checked size. We can therefore use this knowledge on the packet sizes to establish higher bit rate guarantees.

In the example from Figure 2, all paths containing 2→4 require a minimum packet size of 100 byte. It is not obvious if the short path 2→3 triggered by a small packet, or one of the longer paths triggered by a longer packet, result in a lower achievable bit rate. To identify the path with the lowest bit rate, both the execution time of paths and the minimum packet size required by paths must be considered. Our approach can be used to analyze either a packet rate guarantee or a bit rate guarantee by ignoring or analyzing minimum packet sizes.

Path Explosion. When searching for the path with the lowest bit or packet rate, a naïve approach would simply enumerate *all* path and check each path for contradicting branch conditions and throughput capacity. However, the number of paths may be too large to enumerate them all. In our example, there are only 2^2 paths through the **`if`**s in lines 4 and 6. Yet, a program with n consecutive **`if`**s may produce 2^n paths rendering naïve enumerations quickly infeasible.

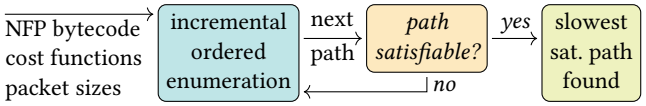

Figure 3: Searching for the slowest satisfiable path.

Naturally, we strive to enumerate only as few paths as possible while still producing valid throughput guarantees. By enumerating paths ordered from slowest path to fastest path (§ 5.1), we get valid throughput guarantees quickly and ignore all paths whose throughput capacity is too high to contribute to the worst-case throughput capacity. As shown in Figure 3, incremental ordered enumeration yields only the slowest paths which are then checked by an SMT solver for contradictions (§ 5.5). In case the enumerated paths are unsatisfiable, incremental ordered enumeration then yields some more paths until a satisfiable path is found. Due to the ordered enumeration, we can stop on the first satisfiable path without enumerating all paths. Thus, the runtime is primarily dominated by proving that low-throughput paths are impossible. This has a further upside, each enumerated path incrementally improves the lower bound, since all paths with a lower estimated throughput capacity have already been shown to be impossible. Our approach, therefore, produces valid intermediate results within a very short analysis time even when a program contains huge numbers of impossible paths with a low throughput capacity.

In the following, we provide details and design rationale for the different steps of our approach. Since the path enumeration builds upon the throughput costs, we start by analyzing the processing (§ 4.1) and DRAM throughput capacity (§ 4.2).

4 PER-PATH THROUGHPUT CAPACITY

Our approach enumerates program paths ordered by the throughput capacity of individual program paths. For that purpose, a heuristic that estimates the execution time of individual instructions can be used to determine the throughput capacity of individual program paths. We start with packet rate throughput since each received packet triggers one program execution. The resulting bit rates are determined at a later step (§ 5.3) by combining these packet rates with program path-specific packet size information.

In an ideal scenario, the SmartNIC manufacturer who has complete knowledge of the inner workings of the SmartNIC would provide a model which perfectly describes the throughput capacities. The documentation [43, 46] of the used SmartNIC contains only incomplete execution timing data and no throughput model. We, therefore, performed measurements on the Netronome Agilio CX SmartNIC to build throughput heuristics of the relevant parts.

We identified two NIC parts with a throughput capacity which varies based on the executed instructions: the processing cores and the DRAM memory engine. Whenever only one of these is overloaded, the other will spend some of its time idling. The actual throughput capacity of a path is the minimum throughput capacity over all parts. We therefore analyze a program's throughput capacity separately for each part and then use the minimum. Each instruction is therefore modeled by both a processing core execution time for the case that the processing cores are overloaded and a DRAM execution time for the case that the DRAM memory

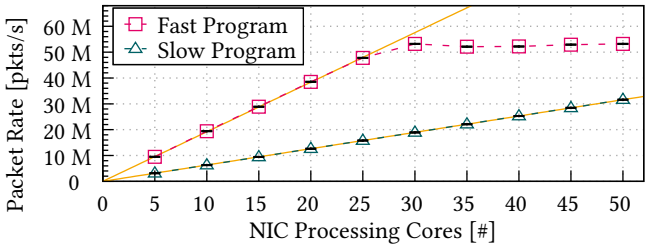

Figure 4: The throughput capacity when using different numbers of NIC cores is limited by the maximum packet rate of the NIC of 54.4M pkts/s. Below that, packet rates are proportional (orange lines) to the number of cores.
fast program: 1.92M pkts/s/core with $R^2 = 0.999949$
slow program: 0.63M pkts/s/core with $R^2 = 0.999985$

engine is overloaded. Our approach can be extended to handle more parts, but candidates like the per-island SRAM did not show any bottleneck behavior and BPF/XDP programs do not have access to any of the additional NFP co-processors.

Lastly, the NIC also has a fixed program-independent throughput limits such as the maximum rate at which the MAC part of the NIC accepts packets or the maximum bit rate of the used Ethernet variant (2x40 GbE in our testbed). For a program to run with maximum throughput, both the program's processing core throughput capacity and the program's DRAM throughput capacity need to be higher or equal than the fixed program-independent NIC limits.

We start with the processing core throughput heuristic for non-memory instructions followed by DRAM throughput and memory instruction timing.

4.1 Processing Cores Throughput

We want to estimate the throughput capacity of the processing cores for individual program paths. Since programs are executed in parallel on many processing cores, the resulting throughput capacity is influenced by the parallelization onto many cores and the execution time of the program path.

Many-Core Parallelization. The Netronome Agilio CX executes BPF/XDP programs on 50 processing cores. To investigate the impact of parallelization we measure the throughput while varying the number of processing cores by using multiple modified NIC firmware variants. We use BPF/XDP programs which do not access any memory, since in this first step we only investigate the processing cores. § 6.1 has more details on how we generate huge numbers of identical packets to always trigger the same program path.

Figure 4 shows the resulting packet rates for two programs, a fast program performing few calculations on each packet and a slower program performing more calculations. As can be seen with the black bars showing the 99% confidence intervals, there is only little variation between multiple runs of the same configuration. No configuration exceeds a throughput of 54.4M pkts/s, which was confirmed by Netronome to be roughly the maximum rate at which the MAC part of the NIC can receive packets. Below this limit, the packet rate is strongly proportional to the number of cores, which can be seen by the fitted lines with a resulting R^2 close to 1. Since

the throughput is proportional to the number of cores and the clock frequency of the cores is fixed, the throughput capacity can be calculated as:

$$\#cores \times \frac{clock\ frequency}{cycles\ per\ packet}$$

Clock Cycles per Packet. The number of clock cycles that a processing core spends per packet is composed of the instructions executed inside the program and overhead in the firmware when moving from one packet to the next. Since Netronome provides a cycle accurate firmware simulator, we were confident that an accurate model of instructions costs is possible. The NFP reference manual [46] states that most non-memory instructions take a single cycle. The cycle costs of a branch instruction is higher if the branch is taken, but does not depend on previous executions since the NFP has no branch prediction. We confirmed and extended the cycle costs with microbenchmarks to build a cycle-accurate model of the relevant non-memory NFP instructions. Given the instruction trace of a program path, the model gives the number of cycles to execute this program path. To calculate the resulting throughput capacity, we additionally need the number of cycles between returning from the program until the program is invoked with the next packet.

To quantify the per-packet firmware overhead we used the smallest BPF/XDP program, overloaded the NIC with packets, and measured the resulting packet rate. The NIC firmware [48] however contains variable packet processing, as it parses multiple headers to assign packets from the same flow to the same host queue. Whenever a processing core processes a packet, it first selects a host queue for the packet and then calls the BPF/XDP program which can arbitrarily override the selected queue. Therefore, queue selection can safely be removed from the firmware, since the identical functionality can be implemented within a BPF/XDP program (or even be replaced by more advanced queue selection [29]) for which we then can determine a throughput guarantee. As an alternative, we could have extracted the queue selection part from the firmware and include it in the program analysis. By removing the queue selection decision from the firmware, we obtained a fairly constant per-packet firmware overhead of approximately 224 cycles which we found to be independent of packet sizes and content. [1] The per-packet cycle overhead is then calculated by converting the measured packet rate into mean cycles per packet and subtracting the calculated cycle costs of our benchmark program. When combining this overhead with an instruction trace, we can calculate the throughput capacity.

4.2 Memory Access

So far, we have looked at non-memory instructions. To analyze programs that access packet data in the per-island SRAM or permanent state in the shared DRAM, we assess the cycle costs and memory bottleneck of memory instructions.

The closed source variant of the NIC firmware [45] accesses the shared DRAM through a hash table abstraction with hidden code which we cannot analyze, whereas the open source NIC firmware [48] does not support DRAM access from BPF/XDP programs. Since raw memory instructions are easier to analyze, we modified the open source NIC firmware and the NFP Linux kernel driver to expose the shared DRAM as raw memory through BPF array maps. More complex memory access schemes can then be implemented within BPF/XDP programs and will then be analyzed by our approach together with the rest of the program.

[1] All modifications are open-sourced as described in Appendix § A.

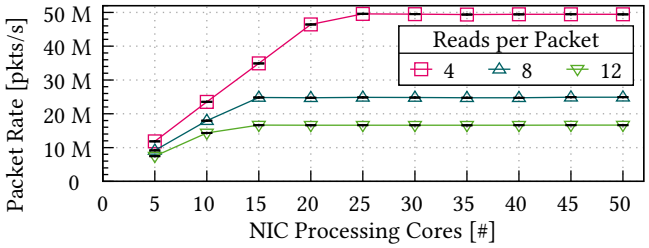

Figure 5: The DRAM bottleneck is observable when enough processing cores are used. The memory engine then performs at a rate of 197.4M ops/s to 199.7M ops/s independent of the number of reads per packet.

The NIC's documentation contains only coarse memory latency information [43] and no memory throughput data. We instead derive a throughput capacity heuristic from measurements. Since BPF and NFP pointers always explicitly encode the targeted memory region in the memory hierarchy, per-island SRAM and shared-DRAM performance can be independently modelled. As stated before, we observed no bottleneck behavior on the per-island SRAM but a varying shared-DRAM throughput capacity which depends on the executed memory operation and the accessed memory locations.

The observed DRAM throughput is lowest when spreading the accessed locations by no more than 16 byte and increases by four times when spreading accesses over large ranges. Since we determine throughput guarantees, we must analyze the worst-case which is the case were only a small range of memory is accessed. A factor of up to 4 may cause a huge underestimation of the actually achievable throughput and we are unable to analyze which memory access patterns a program may experience. However, our evaluation (§ 6.1) shows a much smaller gap between estimated worst-case and measured throughput, since the analyzed programs repeatedly access the same memory locations when repeatedly receiving the same packet.

DRAM Throughput Capacity. As shown in Figure 5, we measured the achievable packet rate for small programs which perform different numbers of read operations to the same location in the shared DRAM. When using few processing cores, the processing cores are the bottleneck, as can be seen by the initial proportional increase in packet rate when increasing the number of cores. Once there are enough cores to overload the memory engine with read operations, the packet rate remains constant since the memory engine throughput capacity now dominates the resulting packet rate. The resulting memory throughput, which is calculated by multiplying the packet rate with the reads per packet, is in the range of 197.4M ops/s to 199.7M ops/s for all program variants. We conclude that a read operation incurs a constant worst-case cost on the DRAM memory engine.

We repeated the same measurements with the second DRAM memory operation supported by the BPF/XDP to NFP compiler [49], atomic increment, and observed a constant throughput capacity of

248.5M ops/s to 250.5M ops/s. With the derived DRAM cost functions for memory instructions, we can calculate the DRAM throughput capacity of a program path. The overall throughput capacity of a program path is then limited by the minimum over its DRAM throughput capacity and processing core throughput capacity.

Memory Cycle Costs. Executing memory instructions puts load onto both the memory engine and processing cores. Although the NFP processing cores have no caches, memory instructions still take a variable number of cycles. If the memory engine is not overloaded, an atomic increment takes a single cycle on a processing core since it does not wait for a response from the memory engine. However, reading from DRAM or SRAM pauses execution until the result is available, even when the memory engine is not overloaded. Hyperthreading masks the throughput impact of waiting for a response since another instance of the same program is scheduled. This works well for a single memory access, but can still lead to all threads waiting, if all threads on a core issue memory requests within short succession. We found that the resulting throughput can be estimated well by a minimum number of clock cycles between two memory operations within an instruction trace. We empirically determined the minimum cycles between blocking memory instructions and found different values for SRAM and DRAM access.

Our processing core and DRAM throughput capacity heu-ristics are complete for all instructions issued by the BPF/XDP to NFP compiler and can be used to determine the throughput capacity for all program paths. In the next step, we build upon these heuristics and use the derived cost functions to enumerate paths ordered by achievable bit rate or packet rate and identify or underestimate the slowest satisfiable path.

5 FINDING THE SLOWEST SATISFIABLE PROGRAM PATH

Given that we can estimate the runtime costs of individual program paths, we now need to find the slowest path. However not all paths are actually possible to execute. As such we are looking for the satisfiable path that gives the lowest throughput capacity, which we will simply refer to as the slowest satisfiable path (SSP). This SSP yields a valid throughput guarantee, since all other paths have either a higher throughput capacity or cannot be executed.

5.1 Incremental Sorted Path Enumeration

To mitigate path explosion, we avoid analyzing fast paths, since only the SSP determines the throughput guarantee. Instead, we enumerate paths ordered from lowest to highest throughput and stop analyzing on the first satisfiable path. With each analyzed path, we get an improved lower bound until the SSP yields the final throughput guarantee.

Before identifying the SSP, it is unknown how many impossible paths need to be checked. Enumerating a fixed number of slowest paths may not suffice to find the SSP. Therefore, a mechanism is needed to efficiently enumerate additional paths in case all already enumerated paths are unsatisfiable.

As discussed in the previous section, the SmartNIC has multiple throughput limiting components and we use separate cost functions for each component, e.g., processing core cycle costs and DRAM cycle costs. The resulting throughput capacity is always the minimum over the throughput capacities of the individual components. Because of its simplicity, we choose the incremental longest path algorithm [35] to enumerate paths for a single component ordered by packet rate. We then combine multiple instances of this algorithm for different components and different packet sizes.

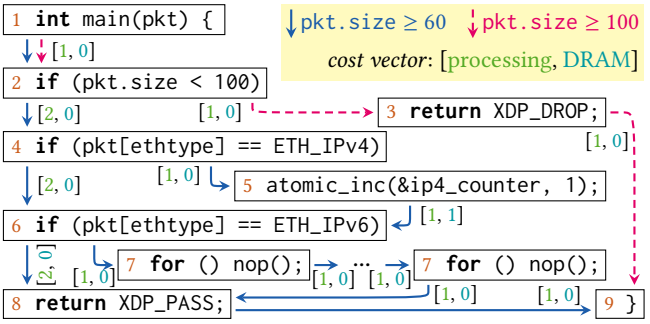

Figure 6: The resulting CFG of the program from Figure 2. A realistic example would use NFP bytecode.

The incremental longest path algorithm [35] was initially proposed to find the maximum delay in integrated circuit (IC) designs. Similar to our problem, IC designs have "non-functional" paths which cannot be triggered and therefore do not contribute to the highest possible propagation delay through the IC. The incremental longest path algorithm is suitable for our needs because, after it has already enumerated the n longest paths it can enumerate the $n+1$ longest path with a time complexity independent of n.

A single instance of this algorithm suffices to determine packet rate guarantees for a single component, e.g., processing cores or DRAM. We continue describing how to transform a program into a graph suitable for this algorithm.

5.2 Preparing a Suitable CFG

To search for the SSP, a control flow graph (CFG) of the analyzed program is needed which is loop-free and has constant costs.

Loop Unrolling. Packet processing programs typically have only bounded loops [30, 57, 65] and the BPF in-kernel verifier only loads NFP programs onto the NIC if it can prove the loop bound [28]. We, therefore, unroll all loops as shown in Figure 6 which has multiple copies of the **for** loop in line 7.

Cost Vector. We search for the SSP according to the processing core cycles and DRAM cycles of individual instructions. Therefore, each edge in Figure 6 has a cost vector although all edges except 5→6 do not incur any DRAM costs. Paths for either the processing cores or DRAM can then separately be enumerated by using a single entry from the cost vectors.

Edge Costs. We use edge costs instead of node costs, since the number of processing cycles to execute a branch instruction depends on the branch result as can be seen at 2→3 and 2→4. This does however not yet suffice to have constant costs since our cost functions for memory instructions depends on the number of cycles since the previous memory access. We overestimate the cycle costs for each edge according to the maximum possible with any path to this edge, which gives a valid underestimated throughput capacity.

The preprocessed CFG can be used to estimate a packet rate guarantee for either processing cores or DRAM. We continue with combining multiple search instances to estimate overall packet rate guarantees.

5.3 Combining Multiple Cost Functions

We separately enumerate paths for processing cores throughput capacity and DRAM throughput capacity and the minimum over both components gives the overall throughput guarantee. Identifying the SSP for each component separately in succession, however, may result in unnecessary analysis work and does not yield valid overall intermediate results. Instead, we simultaneously use multiple instances of incremental longest path search and interleave their results in ascending throughput capacity order.

Multiple Search Instances. We enumerate paths by their overall throughput capacity with the following procedure. For each cost function, a separate search instance is initialized and asked for the slowest path to form an initial set of candidate paths. Although these paths are enumerated using different cost functions, they can be compared by their throughput capacity and the slower path is the path with the lowest overall throughput capacity. In case this path is found to be unsatisfiable, the originating search instance is asked for the next slowest path such that the set of candidate paths again includes a path from each search instance and can yield the next overall slowest path. With this procedure, paths are enumerated according to the overall throughput capacity and each enumerated path gives a valid intermediate result for the overall throughput guarantee.

Each search instance produces each path, but each path should be enumerated only once and only for its bottleneck throughput. In our example from Figure 6, the path 1→2→4→5→6→7…7→8→9 is a slow path for both the processing cores and DRAM, but we are only interested in the bottleneck of this path. Each path is enumerated first for its bottleneck component and later enumerated again for the other components. We, therefore, only consider paths when they are enumerated for their bottleneck and discard them in all other cases. Now, all paths are enumerated once and ordered by their minimum over the processing core and DRAM throughput capacity.

Up to this point, we can determine packet rate guarantees, but not yet bit rate guarantees. Next, we analyze packet size requirements and use even more search instances.

5.4 Enumerating by Bit Rate

Achievable bit rates depend on the cycle costs of a path and on the minimum packet size (MPS) required to trigger the path. Some CFG edges require larger packets (e.g., 2→4), but this packet size information cannot be mapped to constant edge costs. Instead, we determine the set of possible MPSs for a program and then enumerate paths for each distinct MPS by a separate search instance. In our example from Figure 6 one search instance enumerates paths with an MPS of 60 byte and another search instance enumerates paths with an MPS of 100 byte. For each distinct MPS we additionally need separate search instances for processing core and DRAM throughput capacity. Our example needs a total of four search instances to enumerate paths ordered by their overall bit rate capacity.

Packet Size Analysis. For each distinct MPS, only a subset of the CFG edges is needed to cover all paths with that MPS. We statically analyze the MPS requirement for each CFG edge and then collect the set of edges needed for each distinct size. In Figure 6, the edge 2→4 requires a packet of at least 100 byte and additional predecessor and successor edges are needed to cover all paths which have an MPS of 100 byte. In this example, the solid edges are used to enumerate 100 byte paths and the dashed edges to enumerate 60 byte paths. Edge 1→2 is needed for both 60 byte and 100 byte.

Since we use static analysis to determine packet size requirements for edges, we have to underestimate them to get valid lower bounds for the bit rate guarantee. Therefore, a search instance for a particular MPS may produce some paths which require a larger packet size. Using the smaller size still yields valid lower bounds and the bit rate guarantee is further improved by enumerating additional paths up to a bit rate which matches the actually larger MPS.

Improving Overestimated Costs. We overestimate edge costs and underestimate packet sizes, both of which lead to an underestimation of a path's throughput capacity. We enumerate ordered by this underestimated and the underestimation for a satisfiable path gives a valid throughput guarantee. This underestimation can be further improved by enumerating additional paths. The non-underestimated throughput capacity of a path can be used once all paths with a higher underestimate have been enumerated. Thereby, lower bounds of underestimated paths can be improved by enumerating a few more paths.

With packet size analysis, paths can be enumerated ordered by their achievable bit rate. Each enumerated path is then checked for satisfiability and the first satisfiable path establishes the bit rate guarantee.

5.5 Checking Paths for Satisfiability

Some program paths cannot be triggered by any packet since they contain contradicting branch conditions. We use an SMT solver to check each enumerated path for such contradictions. In case a path is satisfiable, the SMT solver additionally produces an accurate MPS and a minimally sized packet and DRAM assignment to trigger the path.

Symbolic Memory and Pointers. We track register and memory assignments with quantifier-free bitvector and array logic, resulting in branch conditions that depend on a symbolic packet and symbolic DRAM content. The memory region addressed by a symbolic pointer can be ambiguous. Since the BPF in-kernel verifier ensures that pointers always stay within their memory region, we can assume a segmented memory model [2] where no operation on a pointer can change the memory region it points to.

For each satisfiable path, the SMT solver additionally produces a DRAM assignment and minimally sized packet which triggers the path. We evaluate the estimation accuracy by measuring throughput with these example packets.

6 EVALUATION

We evaluate the estimation accuracy, the time to compute the throughput guarantees, some of the design choices, and use cases.

Table 1: Overview over the analyzed programs.

Analyzed Program	Language	Bottleneck	Loops	# NFP Instructions
switch.p4 (parser) [10]	P4	processing	–	1 559
Cloudflare DoS [1, 9]	C	processing	–	406
QUIC LB (IPv4) [17]	C	processing	–	599
QUIC LB (IPv6) [17]	C	processing	✓	662
RTP a→μ-law [21, 56]	C	processing	✓	233
RTP a→μ-law (opt) [21, 56]	C	processing	✓	205
DNS Cache [55]	C	DRAM	–	2 137
Count-Min (5) [11]	C	processing	✓	668
Count-Min (10) [11]	C	DRAM	✓	991
Count-Min (15) [11]	C	DRAM	✓	1 403
Count-Min (20) [11]	C	DRAM	✓	1 743
Path Explosion	C	processing	–	1 035

Implementation. Our prototype fully implements our approach, analyzes real BPF/XDP programs and is open source as described in the Appendix § A. We use the Z3 [40] SMT solver and enumerate batches of program paths to parallelize satisfiability checking onto CPU cores. Unlike KLEE-based [6] approaches such as CASTAN [52], BOLT [30], and SymPerf [54], our implementation does not analyze LLVM bytecode. We instead choose to directly analyze NFP bytecode, since LLVM bytecode lacks many important subtleties of the performance and behavior of the NFP.

Analyzed Programs. We estimate and measure throughput on real BPF/XDP programs shown in Table 1. The number of NFP instructions cannot be directly translated to throughput, since not every program path triggers each instructions, some of the programs have loops, and the individual instructions incur non-uniform processing and DRAM costs. Most programs are written in C, whereas switch.p4 (parser) is implemented in the P4 language and preprocessed by p4c-xdp [5] which produces BPF/XDP-compatible C. We removed everything but the parser from the original switch.p4 [10] and reduced the nesting depth for some protocols in order to fit the program onto the NFP. The Cloudflare DoS program filters unwanted packets during a DoS attack and is reconstructed from code [1] and tools [9] published by Cloudflare. Most other programs were created by different team members using existing documentation. The QUIC LB computes on the QUIC connection id as described in an Internet draft [17], the RTP a→μ-law transcodes up to 160 bytes of audio payload according to the standards [21, 56], the DNS Cache responds with precomputed standard-compliant [55] DNS responses, and the Count-Min sketches [11] counts the number of UDP and TCP flows, using a varying number of hash functions. Finally, we created a program to resist our analysis with 2^{64} unsatisfiable paths which are slower than the SSP.

Some of the analyzed programs can exceed the maximum achievable packet rate of the NICs MAC part of 54.4M pkts/s (26 G Bit/s at 60 byte packet size) when executed on all 50 processing cores. Although this is usually a desired result when developing a program, it limits our ability to evaluate the estimation accuracy, as we can no longer measure the program's throughput capacity. For this evaluation, we, therefore, estimate and measure throughput capacities of processing core limited programs at 5 processing cores and then scale these numbers to 50 processing cores. We still estimate and measure DRAM throughput limited programs at 50 processing cores at the cost of being unable to measure all satisfiable paths through these programs.

Table 2: The throughput guarantees are improved by up to 44% by identifying the SSP and increase by up to 13% by measuring identified paths. The estimated slowest paths are correct with an error of at most 1.0%.

Analyzed Program	Naïve Bound [Bit/s]	Estimated Slowest Sat. Path [Bit/s]	Slowest Measured Path [Bit/s]
switch.p4 (parser)	17.1G	**+44%** 24.7G	**+4.1%** 25.8G ±0.05G
Cloudflare DoS	32.1G	**+10%** 35.2G	**-1.0%** 34.7G ±0.07G
QUIC LB (IPv4)	22.8G	**+0%** 22.8G	**+2.9%** 23.5G ±0.04G
QUIC LB (IPv6)	21.4G	**+29%** 27.6G	**+2.9%** 28.5G ±0.05G
RTP a→μ-law	2.97G	**+0%** 2.97G	✓ 2.97G ±0.01G
RTP a→μ-law (opt)	4.70G	**+0%** 4.70G	✓ 4.70G ±0.01G
DNS Cache	6.4G	**+41%** 9.0G	**+13.1%** 10.4G ±0.02G
Count-Min (5)	21.5G	**+0%** 21.6G	**+2.4%** 22.2G ±0.01G
Count-Min (10)	11.9G	**+0%** 11.9G	✓ 12.0G ±0.06G
Count-Min (15)	8.0G	**+0%** 8.0G	✓ 8.0G ±0.03G
Count-Min (20)	6.0G	**+0%** 6.0G	✓ 6.0G ±0.04G
Path Explosion	1.2G	–	–

6.1 Estimation Accuracy

To assess the accuracy of our estimates, we measure the throughput of individual program paths.

Testbed. We use a Barefoot Tofino Switch to generate huge numbers of identical packets, similar as proposed by P4pktgen [51]. Each program path is measured separately by repeating a single packet which always triggers this path. For most program paths, the throughput capacity of a single path can be measured by filling 2x40 GbE with packets back to back. We then determine the rate of actually processed packets by reading NIC counters at fixed intervals over 30 s runs and calculating 99% confidence intervals.

Due to a bug in the MAC part of the NIC firmware (confirmed by Netronome) we had to measure some of the program paths differently. When a program has a throughput capacity between ~36 M pkts/s and ~50 M pkts/s (but not above ~50 M pkts/s) and is overloaded with packets, the NIC accepts packets at a rate of only ~30 M pkts/s. Since the NIC processes packet rates up to the throughput capacity of the program path, as long as the program is not overloaded, we measure these paths by shaping the rate of transmitted packets to determine the maximum rate the NIC can handle without breaking down.

Per-Path Accuracy. To asses the limits on our estimation accuracy, we measure the throughput of many paths. We, therefore, enumerate not only the SSP but continue enumerating slower paths for one hour, thereby discovering a total of 21 470 measurable paths. The estimate matches the measurement for 9.9% of paths, underestimates 89.6% of paths and is too high for 0.4% of paths. No processing- and

Table 3: The time it takes to calculate naïve and worst-path throughput guarantees compared to the time it takes to enumerate and check all possible paths.

Analyzed Program	Naïve Bound	Slowest Sat. Path	All Sat. Paths
switch.p4 (parser)	**1.6 s** ±16 ms	**28 s** ±24 ms	**68 s** ±69 ms
Cloudflare DoS	**0.6 s** ±15 ms	**0.9 s** ±14 ms	**2.5 s** ±13 ms
QUIC LB (IPv4)	**0.3 s** ±63 ms	**0.4 s** ±63 ms	**0.5 s** ±65 ms
QUIC LB (IPv6)	**1.1 s** ±21 ms	**33 s** ±55 ms	≥ **1 h**
RTP a→μ-law	**23.5 s** ±57 ms	**102 s** ±57 ms	≥ **4 m**
RTP a→μ-law (opt)	**23.5 s** ±59 ms	**95 s** ±70 ms	≥ **1 h**
DNS Cache	**2.5 s** ±25 ms	**38 s** ±64 ms	**13 m** ±0.4 s
Count-Min (5)	**0.2 s** ±1 ms	**0.2 s** ±1 ms	**0.5 s** ±2 ms
Count-Min (10)	**0.3 s** ±7 ms	**0.4 s** ±7 ms	**4.2 s** ±9 ms
Count-Min (15)	**0.5 s** ±3 ms	**0.7 s** ±4 ms	**51 s** ±0.1 s
Count-Min (20)	**0.8 s** ±4 ms	**1.0 s** ±4 ms	**33 m** ±5.6 s
Path Explosion	**0.5 s** ±4 ms	≥ **47 m**	≥ **47 m**

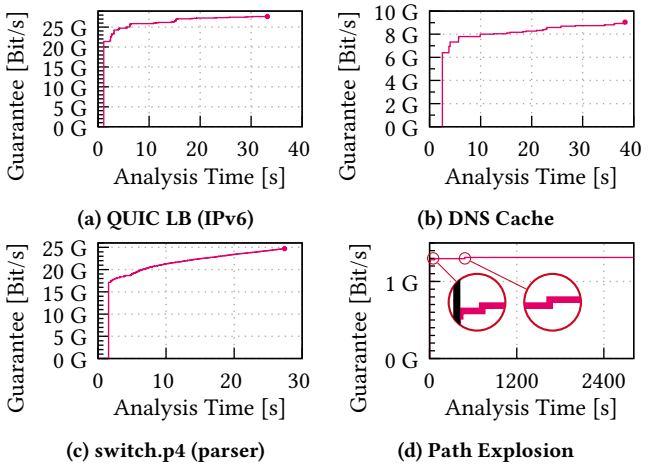

Figure 7: The throughput guarantee improves until a satisfiable path is found (a-c) or the the analysis is aborted (d).

memory-bottlenecked paths is underestimated by more than 8.5% and 18.2%. For the paths with a too high estimate, no estimate exceeds the measured throughput by more than 1.7%, possibly caused by inaccuracies in our per-path throughput heuristic. Despite our per-path throughput heuristic being based on measurements, it still produces mostly accurate and tight lower throughput bounds.

Slowest Satisfiable Path. We establish throughput guarantees for programs by identifying the SSP. The estimated SSP is indeed also the slowest measured path for all except one program. For the DNS Cache, the slowest measured path was wrongly estimated to be the seventh slowest path and has a measured bit rate 2.4% lower than the measured bit rate of the estimated SSP. Such inaccuracies are expected, since our example packets do not produce a worst-case memory access pattern. As shown in Table 2, the slowest measured bit rate for DNS Cache is still 13.1% higher than the estimated worst-case throughput capacity. For all analyzable example programs, the estimated worst-case throughput capacity is close to the slowest measured path.

Naïve Lower Bound. For each program, we calculate different throughput guarantees: a naïve lower bound which is the throughput estimate for the slowest, possibly unsatisfiable, path, and a throughput estimate of the SSP. As can be seen in Table 2, this search for the SSP improves the throughput guarantees by up to 44%. However for some programs, the naïve bound cannot be improved, since for these programs the overall slowest path is already satisfiable. Satisfiability checking of paths has the potential of significantly improving throughput guarantees, but is not needed for all programs and prolongs the analysis time.

6.2 Analysis Time

For a useful approach, analysis results have to be computed within a reasonable time, even when path explosion happens.

Analysis Setup. We executed our prototype on a desktop computer with an Intel Core i7-7700 CPU with 4 cores (8 threads) and 16 GiB of RAM. Every program analysis was repeated over 20 runs with non-terminating runs being aborted after one hour. The results of our analysis time evaluation are realistic since we fully implemented the approach as a working prototype, ran this prototype on real programs, and used a typical desktop computer.

Analysis Time. As can be seen in Table 3, the naïve bound is computed on all example programs within 23.5 s and except for the Path Explosion program, the SSP is found within 102 s. Analyzing a SmartNIC program takes only little time, enabling developers to regularly check throughput guarantees. The analysis times are so short, it is even feasible to integrate our prototype into regularly executed regression tests.

A major advantage of our approach is the ordered enumeration of program paths. Exhaustive symbolic execution approaches such as BOLT [30] and SymPerf [54] always analyze all paths through a program, whereas our approach only analyzes the slowest paths. For comparison with such approaches, we do not stop on the SSP but continue enumerating all satisfiable paths as shown in Table 3. When enumerating all paths, the analysis time increases by a factor of up to × 2 039 and becomes infeasible for some programs within an hour or because we ran out of memory before that. Our approach, therefore, provides significantly lower analysis times. By focusing on only the slowest paths, we enable the analysis of many programs which otherwise would have too many paths to analyze. Note that enumerating additional paths or directly estimating a programmer-defined path is possible and may give further insights. Incremental sorted path enumeration identifies the SSP in significantly shorter time compared to exhaustive symbolic execution.

Path Explosion. On the Path Explosion program, our prototype checked 241 174 paths before running out of memory without having discovered a single satisfiable path. However, the naïve bound, which is a valid throughput guarantee, can always be computed independently of path explosion.

Intermediate Results. In case it takes too long to identify the SSP, ordered enumeration produces valid intermediate results for the throughput guarantee. Each plot in Figure 7 shows one analysis run where a first throughput guarantee is established through the naïve bound and then improved until the SSP is found or the analysis is

Table 4: Range of change in analysis time, $> \times 1.00$ is slower.

Alternative Implementation	SSP Analysis Time
Check Satisfiability on each Branch	$\times 1.19 - \geq 1$ h
No Static Analysis	$\times 0.35 - \geq 1$ h
Separate Processor & DRAM Analysis	$\times 1.07 - \times 2.36$
Packet Rate Analysis	$\times 0.02 - \times 1.37$

aborted. If for example, the QUIC LB (IPv6) program needs to process 25 GBit/s, the analysis can already be stopped after 5.7 s instead of 33.2 s. There is however no guarantee that a useful intermediate result is produced in a significantly shorter time, as can be seen with the Path Explosion program. The ability to produce intermediate results before identifying the SSP is a direct result of our choice to perform incremental ordered enumeration.

To summarize, our prototype finds the SSP within minutes on all useful example programs and yields intermediate results before that. In case the SSP cannot be found, the naïve lower bound and additional intermediate results still produce valid throughput guarantees for any program.

6.3 Influence of Design Choices

The analysis time is influenced by several design choices.

Satisfiability Checking. Unlike our approach, symbolic execution checks the satisfiability each time the searcher crosses a branch instruction. We only check the satisfiability of *completed* paths ordered by throughput capacity to avoid any checks on fast paths. For comparison, we modified our prototype to perform checks on each branch and repeated the search for the SSP on all example programs except Path Explosion. Performing an SMT check on each branch increased the number of required SMT checks by a factor of up to $\times 23.5$ and increases the CPU time spent in the Z3 SMT checker by a factor of up to $\times 272.9$. The impact on the overall analysis time is shown in Table 4. The time to find the SSP increases on the example programs by at least a factor of $\times 1.19$ and for some programs increases the time from previously minutes to beyond 1 hour, confirming the advantage of our choice over symbolic execution- and KLEE-based approaches (e.g., CASTAN [52], BOLT [30], SymPerf [54]).

Static Analysis. We use static analysis to remove impossible CFG edges and to determine the per-edge minimum packet size. Some example programs do not benefit from this static analysis and can be analyzed faster without it, for other programs it becomes unfeasible to analyze them in a reasonable time without static analysis. Static analysis, therefore, is an important step in our approach.

Combined Processor & DRAM Anlysis. Instead of interleaving the throughput capacity analysis for the processing cores and DRAM, these could be analyzed separately in succession. Separate analysis not only has the disadvantage of no valid intermediate results but also is slower in all cases.

Packet- vs. Bit-Rate Analysis. Depending on the program and use case, one might be interested in packet rates or bit rates. Packet rate analysis is simpler since it is independent of packet size information. Packet rate analysis can indeed be significantly faster, but this is not the case for all programs.

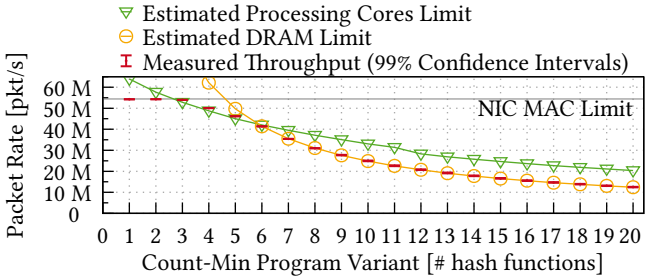

Figure 8: Packet rate estimations and measurements for different variants of the Count-Min flow counter.

6.4 Use Cases

Our approach can have additional uses while developing a SmartNIC program.

Program Optimization. We were unsatisfied with the bit rate of the RTP a→μ-law transcoder. Upon inspecting where the SSP spends most of its execution time, we were able to create a program variant with identical behavior but a 58% higher bit rate. Our approach, therefore, can be used as a tool to aid the optimization of programs, although this still requires human effort.

Program Parametrization. The Count-Min program overcounts the number of packets for individual network flows. When increasing the number of different hash functions used for the count-min sketch, the counting accuracy increases at the cost of a decreased achievable packet rate. We, therefore, analyze differently parameterized variants of this program. As shown in Figure 8, the Count-Min program achieves a perfect throughput with up to 2 hash functions, is processing core limited up to 5 hash functions, and DRAM throughput limited beyond. A developer of such a program can use our approach to analyze the accuracy vs. throughput tradeoff and better decide on a suitable parametrization.

7 DICUSSION & FUTURE WORK

Our approach works quite well, but can still be improved.

Search Strategy. The time to find the SSP is dominated by the work of the satisfiability checker. Improvements in the search strategy should, therefore, focus on reducing the number of satisfiability checks. This may be achieved by reusing satisfiability results for similar paths or reducing the number of to-be-checked paths. Especially, inaccurate minimum packet size estimates may result in incorrect ordering of paths and therefore too many satisfiability checks. Replacing the static analysis of packet size requirements may therefore improve the analysis time for some programs.

Packet Sizes. We analyze for minimum packet sizes, but real packets will often be larger. Since our packet size analysis is based on accessed packet memory, typical packet payloads that are forwarded but not accessed, are ignored. Better incorporating actual or typical packet sizes might lead to throughput guarantees which are closer to the actual throughput.

DRAM Access Patterns. Although our throughput capacity estimates are close to our measurements, the estimates will sometimes be much lower than the real throughput. We assume a worst-case DRAM access pattern (§ 4.2), but we are unable to analyze if a

path can experience such a bad access pattern and our generated example packets sometimes do not match this bad access pattern. The throughput guarantees could, therefore, be further improved by analyzing programs for their accessed DRAM locations.

Stateful Programs. BPF/XDP on NFP currently does not support programs that modify the program-readable permanent state. Our program analysis can be easily extended to support reading and writing the same DRAM location but may result in an underestimation which is far off from the actual worst-case. Due to the parallelized packet processing, memory content can change in between two reads or a write followed by a read from the same memory location. Additionally, in case a single program path reads and writes the same memory location, this path may not be triggerable multiple times in direct succession. Therefore, the actual worst-case throughput may not be the result of always triggering a single worst-case program path, but rather a sequence of packets triggering different paths on multiple processing cores. Establishing an underestimation for the single worst program path still results in a valid throughput guarantee, but further work is necessary to improve this lower bound.

Beyond BPF/XDP. This paper focuses on programs using the BPF/XDP toolchain executed on a Netronome SmartNIC. For example, a VPN endpoint cannot be reasonably implemented with XDP as the SmartNICs crypto co-processor is currently not accessible from BPF. Extending our approach to analyze programs written in Micro-C [43] or with the Netronome P4 SDK [47] is in principle possible but requires further work. Our approach relies on program constraints such as bounded loops and a clear division between a program which processes a single packet and a main loop which iterates over the received packets. To identify the SSP, a program, therefore, needs to be split into those two parts and loop bounds need to be calculated. When accessing additional co-processor, e.g., for fast crypto operations, the cost vector needs to be extended to handle the additional potential bottlenecks. Executing different code on multiple processing cores may further complicate the analysis.

Other SmartNICs. When analyzing programs for other processor-based SmartNICs such as Mellanox Bluefield [39] or Marvel LiquidIO [38], our approach needs to be adapted to their throughput characteristics [33]. Ideally, the manufacturer of each SmartNIC would provide a throughput model which is then used by our approach to enumerate paths ordered by throughput capacity.

To highlight our approach, we chose a SmartNIC which is easier to predict. As some SmartNICs have memory caches and branch prediction instead of cooperative hyper-threading and many simple processing cores, the throughput underestimation will likely be less accurate. When analyzing the worst-case for such SmartNICs, one must assume that in most cases, the accessed memory locations are not cached and therefore cause worst-case memory access latency and worst-case memory hit rate. Similarly, to provide worst-case guarantees, one must assume incorrect branch predictions. Our approach still provides throughput guarantees, but the determined guarantee will be more off from the typically experienced throughput. Replaying the example packets from our approach will likely cause a much higher throughput than the predicted worst-case since cache misses and branch mispredictions can often be caused only by systematic variations of parts of the packets. Incorporating the memory models and branch predictions models from CASTAN [52] and SymPerf [54] can help in generating packet traces that are closer to the actual worst-case.

In some cases cache misses or branch mispredictions are impossible to trigger. E.g., when all bad programs paths access the same memory location or take the same branch direction, the lower throughput bound can be improved. Future work could focus on proving whether memory access will always hit the cache and analyze the interaction between different program paths.

Network Analysis. SmartNIC programs are not running in isolation but are part of a network of non-programmable and programmable devices and applications and often execute only parts of an application. When automatically splitting programs [62] or NF chains [65] our approach can reason about the performance of program parts. The actual worst-case throughput capacity depends on the behavior and interaction of all devices in the network and can be higher than the minimum over the individual devices. A next step could be the performance analysis of a network of SmartNICs [34].

8 RELATED WORK

Packet Processing Performance. Packet rate and bit rate estimates have been a concern ever since packets were processed on processors [14] in the beginning of the Internet. An important step towards predictable throughput on general-purpose processors is the packet processing system by Dobrescu et al. [16] for which they can extrapolate the throughput when the number of flows changes. Today, the conventional wisdom to achieve predictable throughput is dominated by fixed or programmable match-action pipelines. We show the possibility of predictable throughput on processors by proposing a methodology to analyze SmartNIC programs for their throughput capacity.

SmartNIC Performance. The packet processing performance of SmartNICs is an established topic and we use the same SmartNIC as several previous publications. The current paper is a continuation of our previous work [27] where we showed that throughput and latency of BPF/XDP programs on Netronome SmartNICs can vary greatly. Hasanin et al. [25] presented similar results for P4 programs on the same SmartNIC whereas Katsikas et al. [33] showed similar results for Mellanox SmartNICs. George et al. [23], Dai et al. [12], Wu et al. [64], and Chen et al. [7] optimize SmartNIC programs, but give no guarantee on the resulting performance. Qiu et al. [53] applies a performance model to unported programs to estimate the performance of a potential ported program. Since all these works use traffic traces to estimate the performance, they cannot estimate the throughput for unknown traffic. We instead determine throughput guarantees by analyzing programs for their worst-case.

Mitigating Performance Problems. Without a throughput guarantee, it is unknown how much overprovisioning is needed to prevent throughput bottlenecks. There is a line of work which mitigates performance problems when they occur. iPipe [37] dynamically adapts the offloaded portion of a program, but can only react once it observes an overload. FairNIC [24] partitions the SmartNIC resources among multiple programs, giving each program exclusive access to a subset of processing cores and caches, thereby limiting an overload to a single program. Unlike these approaches, we can

tell beforehand whether an overload may happen and how much SmartNIC resources are needed to prevent throughput bottlenecks.

Non-Performance Program Analysis. Formal methods such as symbolic execution have successfully been applied to packet processing programs to analyze non-performance properties such as finding bugs [15, 51, 57, 60], verifying reachability [15, 61] and proving correctness [18, 20, 42, 66]. These approaches rely on similar program properties as our approach, e.g., no unbounded loops, and their success shows that it is easier to analyze packet processing programs in comparison to many other programs.

Performance Analysis. We analyze SmartNIC programs for their worst throughput capacity. This is similar to worst-case execution time analysis which is a well-established research field [63] and is hard for arbitrary programs on general-purpose processors. Our problem is easier because we analyze throughput instead of latency, because packet processing programs are sufficiently restricted, and because the targeted SmartNICs are comparably simple.

Our approach has similarities and is inspired by using symbolic execution for execution time analysis of packet processing on general-purpose processors. Chipounov et al. [8] proposed this idea in S2E by exemplarily analyzing the longest path through the Apache HTTP Server's URL parser. We presented a very basic approach (Rath et al., SymPerf) [54] on analyzing BPF performance on Intel processors, followed by Pedrosa et al. (CASTAN) [52] and Rishabh et al. (BOLT) [30], which analyze the performance of DPDK programs. All of these works analyze the processing latency of a single-threaded program on a general-purpose Intel processor, whereas this paper analyzes throughput on a highly parallelized SmartNIC, which results in some unique challenges. S2E, SymPerf, and BOLT enumerate all satisfiable program paths, thereby incurring very long analysis times and the inability to analyze programs with path explosion. In contrast, our path enumeration approach provides short analysis times by only analyzing the slowest paths and provides valid throughput guarantees even in the case of path explosion. The authors of BOLT suggest restricting the search space by adding additional constraints on the input packet. Performance results for a constrained input can however not be generalized for packets beyond these constraints. Our approach often runs faster with an unconstrained input, since we stop on the first satisfiable path. S2E and BOLT only provide coarse metrics, such as the number of executed instructions and memory accesses which cannot easily be mapped to throughput. Both, SymPerf and CASTAN use a memory model to infer the memory latency for a single-threaded program, but do not analyze the worst-case with respect to their memory model. We additionally analyze memory throughput for parallelized and multi-threaded execution to underestimate the worst-case throughput.

CASTAN [52] is closest to our work, as it performs directed symbolic execution to find bad packet sequences without analyzing all program paths. CASTAN is a tool to debug bad performance without giving any performance guarantees. Its strength is the ability to find short packet sequences which deterministically result in a similar number of cache misses as long random packet sequences. The authors acknowledge the difficulties of analyzing Intel processors and use several heuristics to guide the directed symbolic execution towards bad performance, thereby often finding local maxima instead of a global maximum. We not only find bad performance but establish throughput guarantees by tightly underestimating the worst-case throughput capacity. To our knowledge, we are the first to incorporate packet size requirements to analyze not only packet-rate but also bit rate performance.

9 CONCLUSION

The achievable packet and bit rate of a SmartNIC program is not obvious and varies between different packets and triggered program paths. SmartNICs are easier to program, whereas programmable match-action pipelines and FPGAs can provide a guaranteed packet rate. We want to provide similar guarantees to SmartNICs by analyzing programs for their guaranteed packet rate and guaranteed bit rate. With our approach, a program developer or network operator can determine whether a SmartNIC program will always achieve the needed throughput. In case the program does not yet achieve this throughput, the program can be further optimized or be parallelized onto the right number of SmartNICs.

Different packets trigger different paths through a program. We analyze the guaranteed throughput by identifying or underestimating the slowest program path. We only consider satisfiable paths, since a program may have slow paths which contain contradicting branch conditions and therefore cannot be triggered by any packet. An underestimation for the throughput capacity of the slowest satisfiable path therefore gives a throughput guarantee for the complete program. Programs may have huge numbers of paths, such that it is unfeasible to check all paths through a program for satisfiability and their throughput capacity. Instead, we incrementally enumerate paths from slowest to fastest and stop analyzing on the first satisfiable path. Our prototype determines throughput guarantees for real programs with an error of at most 1.7% and provides tight lower bounds for the processor- and memory-bottlenecked programs with only up to 8.5% and 18.2% underestimation.

We enable developers and network operators to determine if a program meets the throughput requirements. When integrated into the development toolchains for SmartNIC programs, the developer can get rapid feedback on the throughput capabilities and can iterate on optimizations until the requirements are met. When used for automatic regression tests, changes which lead to undesirable throughput are caught without impacting the production network.

With our throughput guarantees, SmartNICs can be used with the same determinism as programmable match-action pipelines and FPGAs. This enables a step towards more freely programmable switches based on processors without sacrificing throughput guarantees. A network operator does not need to fear throughput problems if our approach assures that the used program has an adequate throughput guarantee. Typical match-action pipelines have a fixed packet rate and allow only few processing steps, even on large packets. However, large packets take longer to transmit and therefore allow for more processing time before the next packet arrives. Our approach can assure high bit rate guarantees, even when a program iterates over the complete payload. With our approach, a program on a processor-based switch can perform many operations on large payloads and still meet the required bit rates.

ACKNOWLEDGMENTS

This work has been funded by the German Research Foundation (DFG) within the Collaborative Research Center (CRC) 1053 "MAKI – Multi-Mechanism-Adaption for the Future Internet".

REFERENCES

[1] Gilberto Bertin. 2017. XDP in practice: integrating XDP into our DDoS mitigation pipeline. In *NETDEV 2.1*.

[2] Borzacchiello, Luca and Coppa, Emilio and Cono D'Elia, Daniele and Demetrescu, Camil. 2019. Memory models in symbolic execution: key ideas and new thoughts. *Software Testing, Verification and Reliability* 29, 8 (Dec. 2019), 35 pages. https://doi.org/10.1002/stvr.1722

[3] Pat Bosshart, Glen Gibb, Hun-Seok Kim, George Varghese, Nick McKeown, Martin Izzard, Fernando Mujica, and Mark Horowitz. 2013. Forwarding Metamorphosis: Fast Programmable Match-action Processing in Hardware for SDN. In *Proceedings of the 2013 ACM SIGCOMM Conference (SIGCOMM '13)*. ACM. https://doi.org/10.1145/2486001.2486011

[4] Marco Spaziani Brunella, Giacomo Belocchi, Marco Bonola, Salvatore Pontarelli, Giuseppe Siracusano, Giuseppe Bianchi, Aniello Cammarano, Alessandro Palumbo, Luca Petrucci, and Roberto Bifulco. 2020. hXDP: Efficient Software Packet Processing on FPGA NICs. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*. USENIX Association. https://www.usenix.org/conference/osdi20/presentation/brunella

[5] Mihai Budiu and William Tu. 2020. Backend for the P4 compiler targeting XDP. https://github.com/vmware/p4c-xdp.

[6] Cristian Cadar, Daniel Dunbar, and Dawson Engler. 2008. KLEE: Unassisted and Automatic Generation of High-Coverage Tests for Complex Systems Programs. In *8th USENIX Symposium on Operating Systems Design and Implementation (OSDI 08) (OSDI 08)*. USENIX Association. https://www.usenix.org/conference/osdi-08/klee-unassisted-and-automatic-generation-high-coverage-tests-complex-systems

[7] Michael K. Chen, Xiao Feng Li, Ruiqi Lian, Jason H. Lin, Lixia Liu, Tao Liu, and Roy Ju. 2005. Shangri-La: Achieving High Performance from Compiled Network Applications While Enabling Ease of Programming. In *Proceedings of the 2005 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '05)*. ACM. https://doi.org/10.1145/1065010.1065038

[8] Vitaly Chipounov, Volodymyr Kuznetsov, and George Candea. 2011. S2E: A Platform for in-Vivo Multi-Path Analysis of Software Systems. In *Proceedings of the Sixteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS XVI)*. ACM. https://doi.org/10.1145/1950365.1950396

[9] Cloudflare. 2014. BPF Tools. https://github.com/cloudflare/bpftools.

[10] P4 Language Consortium. 2020. P4_16 reference compiler. https://github.com/p4lang/p4c.

[11] Graham Cormode and S. Muthukrishnan. 2005. An improved data stream summary: the count-min sketch and its applications. *Journal of Algorithms* 55, 1 (2005), 58–75. https://doi.org/10.1016/j.jalgor.2003.12.001

[12] Jinquan Dai, Bo Huang, Long Li, and Luddy Harrison. 2005. Automatically Partitioning Packet Processing Applications for Pipelined Architectures. In *Proceedings of the 2005 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '05)*. ACM. https://doi.org/10.1145/1065010.1065039

[13] Huynh Tu Dang, Daniele Sciascia, Marco Canini, Fernando Pedone, and Robert Soulé. 2015. NetPaxos: Consensus at Network Speed. In *Proceedings of the 1st ACM SIGCOMM Symposium on Software Defined Networking Research (SOSR '15)*. ACM. https://doi.org/10.1145/2774993.2774999

[14] Donald Watts Davies, K. A. Bartlett, R. A. Scantlebury, and P. T. Wilkinson. 1967. A Digital Communication Network for Computers Giving Rapid Response at Remote Terminals. In *Proceedings of the First ACM Symposium on Operating System Principles (SOSP '67)*. ACM. https://doi.org/10.1145/800001.811669

[15] Mihai Dobrescu and Katerina Argyraki. 2014. Software Dataplane Verification. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*. USENIX Association. https://www.usenix.org/conference/nsdi14/technical-sessions/presentation/dobrescu

[16] Mihai Dobrescu, Katerina Argyraki, and Sylvia Ratnasamy. 2012. Toward Predictable Performance in Software Packet-Processing Platforms. In *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12) (NSDI 12)*. USENIX Association. https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/dobrescu

[17] M. Duke and N. Banks. 2020. QUIC-LB: Generating Routable QUIC Connection IDs. https://tools.ietf.org/html/draft-ietf-quic-load-balancers-03.

[18] Dragos Dumitrescu, Radu Stoenescu, Matei Popovici, Lorina Negreanu, and Costin Raiciu. 2019. Dataplane equivalence and its applications. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*. USENIX Association. https://www.usenix.org/conference/nsdi19/presentation/dumitrescu

[19] Daniel Firestone, Andrew Putnam, Sambhrama Mundkur, Derek Chiou, Alireza Dabagh, Mike Andrewartha, Hari Angepat, Vivek Bhanu, Adrian Caulfield, Eric Chung, Harish Kumar Chandrappa, Somesh Chaturmohta, Matt Humphrey, Jack Lavier, Norman Lam, Fengfen Liu, Kalin Ovtcharov, Jitu Padhye, Gautham Popuri, Shachar Raindel, Tejas Sapre, Mark Shaw, Gabriel Silva, Madhan Sivakumar, Nisheeth Srivastava, Anshuman Verma, Qasim Zuhair, Deepak Bansal, Doug Burger, Kushagra Vaid, David A. Maltz, and Albert Greenberg. 2018. Azure Accelerated Networking: SmartNICs in the Public Cloud. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI '18)*. USENIX Association. https://www.usenix.org/conference/nsdi18/presentation/firestone

[20] Lucas Freire, Miguel Neves, Lucas Leal, Kirill Levchenko, Alberto Schaeffer-Filho, and Marinho Barcellos. 2018. Uncovering Bugs in P4 Programs with Assertion-Based Verification. In *Proceedings of the Symposium on SDN Research (SOSR '18)*. ACM. https://doi.org/10.1145/3185467.3185499

[21] G.711 1988. *Pulse Code Modulation (PCM) of Voice Frequencies*. ITU-T Recommendation.

[22] Nadeen Gebara, Alberto Lerner, Mingran Yang, Minlan Yu, Paolo Costa, and Manya Ghobadi. 2020. Challenging the Stateless Quo of Programmable Switches. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks (HotNets '20)*. ACM. https://doi.org/10.1145/3422604.3425928

[23] Lal George and Matthias Blume. 2003. Taming the IXP Network Processor. In *Proceedings of the ACM SIGPLAN 2003 Conference on Programming Language Design and Implementation (PLDI '03)*. ACM. https://doi.org/10.1145/781131.781135

[24] Stewart Grant, Anil Yelam, Maxwell Bland, and Alex C. Snoeren. 2020. SmartNIC Performance Isolation with FairNIC: Programmable Networking for the Cloud. 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)*. ACM. https://doi.org/10.1145/3387514.3405895

[25] Hasanin Harkous, Michael Jarschel, Mu He, Rastien Pries, and Wolfgang Kellerer. 2019. Towards Understanding the Performance of P4 Programmable Hardware. In *2019 ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS '19)*. IEEE. https://doi.org/10.1109/ANCS.2019.8901881

[26] Frederik Hauser, Marco Häberle, Daniel Merling, Steffen Lindner, Vladimir Gurevich, Florian Zeiger, Reinhard Frank, and Michael Menth. 2021. A Survey on Data Plane Programming with P4: Fundamentals, Advances, and Applied Research. https://arxiv.org/abs/2101.10632.

[27] Oliver Hohlfeld, Johannes Krude, Jens Helge Reelfs, Jan Rüth, and Klaus Wehrle. 2019. Demystifying the Performance of XDP BPF. In *2019 IEEE Conference on Network Softwarization (NetSoft) (NetSoft 2019)*. IEEE. https://doi.org/10.1109/NETSOFT.2019.8806651

[28] Toke Høiland-Jørgensen, Jesper Dangaard Brouer, Daniel Borkmann, John Fastabend, Tom Herbert, David Ahern, and David Miller. 2018. 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)*. ACM. https://doi.org/10.1145/3281411.3281443

[29] Jack Tigar Humphries, Kostis Kaffes, David Mazières, and Christos Kozyrakis. 2019. Mind the Gap: A Case for Informed Request Scheduling at the NIC. In *Proceedings of the 18th ACM Workshop on Hot Topics in Networks (HotNets '19)*. ACM. https://doi.org/10.1145/3365609.3365856

[30] Rishabh Iyer, Luis Pedrosa, Arseniy Zaostrovnykh, Solal Pirelli, Katerina Argyraki, and George Candea. 2019. Performance Contracts for Software Network Functions. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*. USENIX Association. https://www.usenix.org/conference/nsdi19/presentation/iyer

[31] Nic Viljoen Jakub Kicinski. 2016. eBPF Hardware Offload to SmartNICs: cls bpf and XDP. In *Netdev 1.2*.

[32] Xin Jin, Xiaozhou li, Haoyu Zhang, Robert Soulé, Jeongkeun Lee, Nate Foster, Changhoon Kim, and Ion Stoica. 2017. NetCache: Balancing Key-Value Stores with Fast In-Network Caching. In *Proceedings of the 26th Symposium on Operating Systems Principles (SOSP '17)*. ACM. https://doi.org/10.1145/3132747.3132764

[33] Georgios P. Katsikas, Tom Barbette, Marco Chiesa, Dejan Kostić, and Gerald Q. Maguire. 2021. What You Need to Know About (Smart) Network Interface Cards. In *Passive and Active Measurement*. Springer International Publishing. https://doi.org/10.1007/978-3-030-72582-2_19

[34] Johannes Krude, Matthias Eichholz, Maximilian Winck, Klaus Wehrle, and Mira Mezini. 2019. Optimizing Data Plane Programs for the Network. In *Proceedings of the ACM SIGCOMM 2019 Workshop on Networking and Programming Languages (NetPL '19)*. ACM. https://doi.org/10.1145/3341561.3349590

[35] Sandip Kundu. 1994. An incremental algorithm for identification of longest (shortest) paths. *INTEGRATION* 17, 1 (1994), 25–31. https://doi.org/10.1016/0167-9260(94)90018-3

[36] ChonLam Lao, Yanfang Le, Kshiteej Mahajan, Yixi Chen, Wenfei Wu, Aditya Akella, and Michael Swift. 2021. ATP: In-network Aggregation for Multi-tenant

Learning. In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21) (NSDI '21)*. USENIX Association. https://www.usenix.org/conference/nsdi21/presentation/lao
[37] Ming Liu, Tianyi Cui, Henry Schuh, Arvind Krishnamurthy, Simon Peter, and Karan Gupta. 2019. Offloading Distributed Applications onto SmartNICs Using IPipe. In *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM '19)*. ACM. https://doi.org/10.1145/3341302.3342079
[38] Marvell. 2020. Marvell® LiquidIO™ III. https://www.marvell.com/content/dam/marvell/en/public-collateral/embedded-processors/marvell-liquidio-III-solutions-brief.pdf.
[39] Mellanox. 2019. BlueField SmartNIC for Ethernet. https://www.mellanox.com/sites/default/files/related-docs/prod_adapter_cards/PB_BlueField_Smart_NIC.pdf.
[40] Microsoft Research. 2019. The Z3 Theorem Prover. https://github.com/Z3Prover/z3/tree/z3-4.8.7.
[41] Usama Naseer, Luca Niccolini, Udip Pant, Alan Frindell, Ranjeeth Dasineni, and Theophilus A. Benson. 2020. Zero Downtime Release: Disruption-Free Load Balancing of a Multi-Billion User Website. 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)*. ACM. https://doi.org/10.1145/3387514.3405885
[42] Luke Nelson, Jacob Van Geffen, Emina Torlak, and Xi Wang. 2020. Specification and verification in the field: Applying formal methods to BPF just-in-time compilers in the Linux kernel. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*. USENIX Association. https://www.usenix.org/conference/osdi20/presentation/nelson
[43] Netronome. 2014. The Joy of Micro-C. https://open-nfp.org/documents/48/the-joy-of-micro-c_fcjSfra.pdf.
[44] Netronome. 2018. Benefits of a Composable Silicon Architecture. https://www.netronome.com/media/documents/WP_Composable-Architecture.pdf.
[45] Netronome. 2019. Agilio eBPF 2.0.6 - extended Berkeley Packet Filter. https://help.netronome.com/support/solutions/articles/36000050009-agilio-ebpf-2-0-6-extended-berkeley-packet-filter.
[46] Netronome. 2019. Netronome ® Network Flow Processor 6xxx NFP SDK version 6 Flow Processor Core Programmer's Reference Manual.
[47] Netronome. 2019. P4 Data Plane Programming for Server-Based Networking Applications. https://www.netronome.com/media/documents/WP_P4_Data_Plane_Programming.pdf.
[48] Netronome. 2020. CoreNIC: a flexible SR-IOV SmartNIC firmware implementation supporting BPF and stateless offloads. https://github.com/Netronome/nic-firmware.
[49] Netronome. 2020. Netronome Flow Processor (NFP) Kernel Drivers. https://github.com/Netronome/nfp-drv-kmods.
[50] Rolf Neugebauer, Gianni Antichi, José Fernando Zazo, Yury Audzevich, Sergio López-Buedo, and Andrew W. Moore. 2018. Understanding PCIe Performance for End Host Networking. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication (SIGCOMM '18)*. ACM. https://doi.org/10.1145/3230543.3230560
[51] Andres Nötzli, Jehandad Khan, Andy Fingerhut, Clark Barrett, and Peter Athanas. 2018. P4pktgen: Automated Test Case Generation for P4 Programs. In *Proceedings of the Symposium on SDN Research (SOSR '18)*. ACM. https://doi.org/10.1145/3185467.3185497
[52] Pedrosa, Luis and Iyer, Rishabh and Zaostrovnykh, Arseniy and Fietz, Jonas and Argyraki, Katerina. 2018. 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)*. ACM. https://doi.org/10.1145/3230543.3230573
[53] Yiming Qiu, Qiao Kang, Ming Liu, and Ang Chen. 2020. Clara: Performance Clarity for SmartNIC Offloading. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks (HotNets '20)*. ACM. https://doi.org/10.1145/3422604,.3425929
[54] Felix Rath, Johannes Krude, Jan Rüth, Daniel Schemmel, Oliver Hohlfeld, Jó Á. Bitsch, and Klaus Wehrle. 2017. SymPerf: Predicting Network Function Performance. In *Proceedings of the SIGCOMM Posters and Demos (SIGCOMM Posters and Demos '17)*. ACM. https://doi.org/10.1145/3123878.3131977
[55] RFC1035 1987. *DOMAIN NAMES - IMPLEMENTATION AND SPECIFICATION*. Internet Requests for Comments.
[56] RFC7655 2015. *RTP Payload Format for G.711.0*. Internet Requests for Comments.
[57] Fabian Ruffy, Tao Wang, and Anirudh Sivaraman. 2020. Gauntlet: Finding Bugs in Compilers for Programmable Packet Processing. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*. USENIX Association. https://www.usenix.org/conference/osdi20/presentation/ruffy
[58] Amedeo Sapio, Ibrahim Abdelaziz, Abdulla Aldilaijan, Marco Canini, and Panos Kalnis. 2017. In-Network Computation is a Dumb Idea Whose Time Has Come. In *Proceedings of the 16th ACM Workshop on Hot Topics in Networks (HotNets 2017)*. ACM. https://doi.org/10.1145/3152434.3152461
[59] Amedeo Sapio, Marco Canini, Chen-Yu Ho, Jacob Nelson, Panos Kalnis, Changhoon Kim, Arvind Krishnamurthy, Masoud Moshref, Dan Ports, and Peter Richtarik. 2021. Scaling Distributed Machine Learning with In-Network Aggregation. In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21) (NSDI '21)*. USENIX Association. https://www.usenix.org/conference/nsdi21/presentation/sapio
[60] Radu Stoenescu, Dragos Dumitrescu, Matei Popovici, Lorina Negreanu, and Costin Raiciu. 2018. Debugging P4 Programs with Vera. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication (SIGCOMM '18)*. ACM. https://doi.org/10.1145/3230543.3230548
[61] Radu Stoenescu, Matei Popovici, Lorina Negreanu, and Costin Raiciu. 2016. SymNet: Scalable Symbolic Execution for Modern Networks. In *Proceedings of the 2016 ACM SIGCOMM Conference (SIGCOMM '16)*. ACM. https://doi.org/10.1145/2934872.2934881
[62] Nik Sultana, John Sonchack, Hans Giesen, Isaac Pedisich, Zhaoyang Han, Nishanth Shyamkumar, Shivani Burad, André DeHon, and Boon Thau Loo. 2021. Flightplan: Dataplane Disaggregation and Placement for P4 Programs. In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21) (NSDI '21)*. USENIX Association. https://www.usenix.org/conference/nsdi21/presentation/sultana
[63] Reinhard Wilhelm, Jakob Engblom, Andreas Ermedahl, Niklas Holsti, Stephan Thesing, David Whalley, Guillem Bernat, Christian Ferdinand, Reinhold Heckmann, Tulika Mitra, Frank Mueller, Isabelle Puaut, Peter Puschner, Jan Staschulat, and Per Stenström. 2008. The Worst-Case Execution-Time Problem—Overview of Methods and Survey of Tools. *ACM Transactions on Embedded Computing Systems* 7, 3, Article 36 (May 2008), 53 pages. https://doi.org/10.1145/1347375.1347389
[64] Qiang Wu and Tilman Wolf. 2009. Runtime resource allocation in multi-core packet processing systems. In *2009 International Conference on High Performance Switching and Routing (HPSR '09)*. IEEE. https://doi.org/10.1109/HPSR.2009.5307422
[65] Jane Yen, Jianfeng Wang, Sucha Supittayapornpong, Marcos A. M. Vieira, Ramesh Govindan, and Barath Raghavan. 2020. Meeting SLOs in Cross-Platform NFV. In *Proceedings of the 16th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT '20)*. ACM. https://doi.org/10.1145/3386367.3431292
[66] Arseniy Zaostrovnykh, Solal Pirelli, Rishabh Iyer, Matteo Rizzo, Luis Pedrosa, Katerina Argyraki, and George Candea. 2019. Verifying Software Network Functions with No Verification Expertise. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles (SOSP '19)*. ACM. https://doi.org/10.1145/3341301.3359647

A ARTIFACTS

Artifacts are available at https://zenodo.org/record/5515910 or at https://github.com/johannes-krude/nfp-pred-artifacts. These can be used to **repeat** the full evaluation and can be **reused** to analyze new BPF/XDP programs. Included is the source code and documentation of the main approach, our modifications to the SmartNIC device driver and firmware, as well as the infrastructure and raw measurement data from our evaluation accompanied by documentation.

Requirements. It suffices to use Docker on Linux or a single computer with Ubuntu 20.04 to repeat the provided small evaluation example and analzye the precompiled BPF/XDP programs. To repeat the full evaluation, one needs: three computers, a Netronome Agilio CX 2x40 GbE SmartNIC, an additional 2x10 GbE NIC, and a Barefoot Tofino based EdgeCore Wedge BF100-32X switch. The proprietary compilers to build the NIC firmware and Tofino P4 program are not included and need to be obtained from Netronome and Intel.

Implementation. The main approach from the paper is implemented as a tool that determines throughput guarantees by incrementally enumerating programs paths of BPF/XDP programs compiled to Netronome Flow Processor assembly. This tool is implemented in 9600 lines of C++, heavily relies on the SMT solver Z3, and is in part inspired by the KLEE symbolic execution engine. The evaluation infrastructure consists of 4700 lines of Ruby source code with low-level tools written in C and some small additions of Bash,

Python, and P4. Our modifications to the SmartNIC device driver and firmware consist of Linux kernel level C and NFP assembly.

Measurements. The evaluation mainly consists of two different types of measurements. During the estimation phase, the approach as described in the paper is executed on example programs to estimate throughput guarantees. All discovered satisfiable program paths are recorded together with an example packet to trigger the path and the time it takes to discover that path. These results from these measurements are presented in Table 2 (columns 2 & 3), Table 3, Figure 7, and Figure 8 (estimates). For Table 4, the throughput estimation is executed again on all example programs but uses different implementation variants.

The second kind of measurements in the evaluation, are measurements of the actual throughput when executing programs on the SmartNIC. These measurements use the example packets from the estimation phase to measure the throughput capacity of each discovered program path. These throughput measurements and their comparison to the estimates are presented in Table 2 (column 4) and Figure 8 (measured throughput). Some additional throughput measurements are shown in Figure 4 and Figure 5.

Repeating the Evaluation from the Paper. The full evaluation takes approximately 10 days and requires a Netronome Agilio CX 2x40 GbE SmartNIC and a Barefoot Tofino-based EdgeCore Wedge BF100-32X programmable switch. To enable partial repetition, we structured the evaluation into smaller steps, some of which take significantly less time and do not require special hardware. All raw measurements gathered during our evaluation are included, to enable repeating each evaluation step independently of the other steps.

When having only a few minutes to spare, one can try the small evaluation example which estimates the throughput bound of the QUIC LB (IPv4) example program and compares these estimates with existing measurements. With some more available time, one can reanalyze all included measurement data and optionally repeat all throughput estimates. In case of having access to the SmartNIC and a Tofino switch, all throughput measurements can be repeated.

Reusing the Implementation for New Programs. Our implementation of the main approach, as well as the measurement infrastructure can be applied to new real BPF/XDP programs. Any BPF/XDP program which adheres to the constraints for NFP offloading and our modified NIC firmware and driver can be analyzed for its worst-case throughput. The implementation supports analyzing for bit rate or packet rate and can be configured to analyze for processing cores throughput, DRAM memory engine throughput, or both. For each enumerated path, an example packet is generated which can be used to compare the estimated throughput capacity to measured throughput. A detailed description is included which explains how to generate the data as presented in Table 2 and Table 3 for new XDP/BPf programs.

All further documentation is included in the `README.md`.