Multiple Pattern Matching for Network Security Applications: Acceleration through Vectorization

Downloaded from: <https://research.chalmers.se>, 2023-04-26 15:16 UTC

Citation for the original published paper (version of record):

Stylianopoulos, C., Almgren, M., Landsiedel, O. et al (2017). Multiple Pattern Matching for Network Security Applications: Acceleration through Vectorization. 46th International Conference on Parallel Processing, ICPP 2017: 472-482. <http://dx.doi.org/10.1109/ICPP.2017.56>

N.B. When citing this work, cite the original published paper.

Multiple Pattern Matching for Network Security Applications: Acceleration through Vectorization

Charalampos Stylianopoulos, Magnus Almgren, Olaf Landsiedel, Marina Papatriantafyllou
Chalmers University of Technology
Gothenburg, Sweden
Email: {chasty, magnus.almgren, olafl, ptrianta}@chalmers.se

Abstract—Pattern matching is a key building block of Intrusion Detection Systems and firewalls, which are deployed nowadays on commodity systems from laptops to massive web servers in the cloud. In fact, pattern matching is one of their most computationally intensive parts and a bottleneck to their performance. In Network Intrusion Detection, for example, pattern matching algorithms handle thousands of patterns and contribute to more than 70% of the total running time of the system.

In this paper, we introduce efficient algorithmic designs for multiple pattern matching which (a) ensure cache locality and (b) utilize modern SIMD instructions. We first identify properties of pattern matching that make it fit for vectorization and show how to use them in the algorithmic design. Second, we build on an earlier, cache-aware algorithmic design and we show how cache-locality combined with SIMD gather instructions, introduced in 2013 to Intel’s family of processors, can be applied to pattern matching. We evaluate our algorithmic design with open data sets of real-world network traffic: Our results on two different platforms, Haswell and Xeon-Phi, show a speedup of 1.8x and 3.6x, respectively, over Direct Filter Classification (DFC), a recently proposed algorithm by Choi et al. for pattern matching exploiting cache locality, and a speedup of more than 2.3x over Aho-Corasick, a widely used algorithm in today’s Intrusion Detection Systems.

Keywords—pattern matching; SIMD vectorization; gather;

I. INTRODUCTION

Security mechanisms, such as Network Intrusion Detection Systems and firewalls, are part of every networked system and are analyzing network traffic to protect from attacks. An essential building block of many such systems is *pattern matching*, i.e., to discover if any of many predefined patterns exist in an input stream (*multiple pattern matching*), for whitelisting or blacklisting. In the context of Network Intrusion Detection, the data stream is the reassembled protocol stream of the *packets* on the monitored network and the set of patterns (usually in the order of thousands) represents *signatures* of malicious attacks that the system aims to detect.

Motivation and Challenges. Pattern matching represents a major performance bottleneck in many security mechanisms, especially when there is a need to employ analysis on the full packet’s payload (Deep Packet Inspection). In

intrusion detection, for example, more than 70% of the total running time is spent on pattern matching [1, 2]. Moreover, with the increasing interest in Network Function Virtualization (NFV) [3, 4], applications like firewalls and Network Intrusion Detection are now moved into the cloud, where they need to rely on commodity hardware features for performance, like multi-core parallelism and vector processing pipelines.

In this paper, we introduce a vectorizable design of an exact pattern matching algorithm which nearly doubles the performance when compared to the state of the art on modern, SIMD capable commodity hardware, such as Intel’s Haswell processors or Xeon Phi [5]. *Vectorization* as a technique to increase throughput is gradually taking a more central role [6]. For example, architectures with SIMD instruction-sets now provide wider vector registers (256 bits with AVX) and introduce new instructions, such as gathers, that make vectorization applicable to a wider range of applications. Moreover, modern processor designs are shifting towards new architectures, like Intel’s Xeon Phi [5], that, for example, supports 512 bit vector registers. On those platforms, vectorization is not just an option but a must, in order to achieve high performance [7]. In this work we introduce algorithmic designs to utilize these capabilities.

Approach and Contributions. The introduction of *gatherers* and other advanced SIMD instructions (cf. section III) allows even applications with irregular data patterns to gain performance from data parallelism. For example, SIMD can speed up regular expression matching [8, 9, 10]. Here, the input is matched against a single regular expression at a time, represented by a finite state machine that can fit in L1 or L2 cache. Working close to the CPU is crucial for these approaches, otherwise the long latency of memory accesses would hide any computation speedup through vectorization.

The domain of multiple pattern matching for Network Intrusion Detection has challenging constraints that limit the effectiveness of these approaches: applications need to simultaneously evaluate thousands of patterns and traditional state-machine-based algorithms, such as Aho-Corasick [11], use big data structures that by far exceed the size of the cache of today’s CPUs. The size of the patterns varies

greatly (from 1-byte to several hundred byte patterns) and can appear anywhere in the input. That is why SIMD techniques have not been previously considered for exact multiple pattern matching – with a few exceptions discussed in Section VI – for Network Intrusion Detection.

Building upon recent work [12, 13] that take steps in addressing the cache-locality issues for this problem, our approach fills this gap: we propose algorithmic designs for multiple pattern matching that bring together cache locality and modern SIMD instructions, to achieve significant speedups when compared to the state of the art. Combining cache locality and vectorization introduces new trade-offs on existing algorithms. Compared to traditional approaches that perform the minimum required number of instructions, but on data that is away from the processor, our approach, instead, performs more instructions, but these instructions find data close to the processor and can process them in parallel using vectorization.

In particular, our works build on a family of recent methods [12, 13] that propose filtering of the input streams using small, cache efficient data structures. We argue that, as a result, memory latencies are no longer the dominant bottleneck for this family of algorithms while their computational part becomes more significant. In this work, we follow a two-step approach. First, we propose a refined and extended method, which is able to benefit from vectorization while ensuring cache locality. Second, we design its vectorized version by utilizing SIMD hardware *gather* operations. To evaluate our approach, we apply our techniques to the DFC algorithm [12], as a representative example that outperforms existing techniques in Network Intrusion Detection applications, including [13], on which our proposed approach can be applied as well. In particular, we target the computational part of pattern matching for performance optimization and make the following contributions:

- We propose algorithmic designs for multiple pattern matching which (a) ensure cache locality and (b) utilize modern SIMD instructions.
- We devise a new pattern matching algorithm, based on these designs, that utilizes SIMD instructions to outperform the state of the art, while staying flexible with respect to pattern sizes.
- We (implement the algorithm and) thoroughly evaluate it under both real-world traces and synthetic data sets. We outperform the state of the art by up to 1.8x on commodity hardware and up to 3.6x on the Xeon-Phi platform.

The remainder of the paper is organized as follows: Section II gives an overview of important pattern matching algorithms and background on vectorization. Section III describes our system model. In Section IV, we present our approach leading to a new, vectorized design. Section V presents our experimental evaluation. In Section VI, we

give an overview of other related work and we conclude in Section VII.

II. BACKGROUND

In this section we present traditional approaches to pattern matching, followed by a brief description of the DFC algorithm (Choi et al. [12]) to which we apply our approach. Next, we introduce the required background on vectorization techniques.

A. Traditional Approach to Multiple-Pattern Matching

The most commonly used pattern matching algorithm for network-based intrusion detection is by Aho-Corasick [11]. It creates a finite-state automaton from the set of patterns and reads the input byte by byte to traverse the automaton and match multiple patterns. Even though it performs a small number of operations for every input byte, it implies– in practice and on commodity hardware – a low instruction throughput due to frequent memory accesses with poor cache locality [12]: As the number of patterns increases, the size of the state automaton increases exponentially and does not fit in the cache. Nevertheless, the method is heavily used in practice; e.g., both Snort [14], one of the best known intrusion detection systems, as well as CloudFlare’s web application firewall [15], use it for string matching.

B. Filtering Approaches and Cache Locality

Besides state-machine based approaches, there is a family of algorithms that rely on *filtering* to separate the innocuous input from the matches. Recent work focuses on alleviating the problem of long latency lookups on large data structures. Choi et al. [12] present a novel algorithmic design called DFC (Direct Filter Classification), that replaces the state machine approach of Aho-Corasick with a series of small, succinct summaries called *filters*. Such a filter is a bit-array that summarizes only a specific part of each pattern, e.g. its first two bytes, having one bit for every possible combination of two characters that can be found in the patterns. The algorithm is structured in two phases, the *filtering* and *verification*:

- In the *filtering* phase, a sliding window of two bytes over the input goes through an initial filter, as described above, to quickly evaluate whether the current position is a possible starting point of a match. The two-byte windows that passed the initial filter are fed to other, similar filters, each specializing on a family of patterns depending on their length. Since the filters are small (8KB each), they usually fit in L1 cache. Thus, the main part of the algorithm differs from Aho-Corasick and uses only cache-resident data structures, resulting in up to 3.8 times less cache misses [12].
- If a window of two characters passed all filters, there is a strong indication that it is a starting point of a match. For this reason, in the next *verification* phase, the DFC

algorithm performs lookups on specially designed hash tables, containing the actual patterns and performs exact matching on the input and the pattern, to verify the match.

Other algorithms in this family, like [13] as well as this work, operate on the same idea: the input is filtered using cache resident data structures, and only the “interesting” parts of the input is forwarded for further evaluation.

C. Vectorization

Single Instruction Multiple Data (SIMD) is an execution model for data parallel applications, which utilizes processing units that operate on a vector of elements simultaneously, instead of separate elements at a time. SIMD vectorization is a desirable goal in computationally intensive, number-crunching applications, where computation is performed on independent data, *sequentially* stored in memory.

Vector instruction sets have evolved over time, introducing bigger registers and support for more complex instructions. Recently, vector instruction sets have been enriched with the *gather* instruction [16] that enables accessing data from *non-contiguous memory locations* (described in detail in Section III). Polychroniou et al. [17] study the effect of vectorization with the *gather* instruction on Bloom filters, hash tables joins and selection scans among others. We are building on these works with SIMD instructions and extend their design to pattern matching with the applications we focus on.

III. SYSTEM MODEL

In this section we introduce the assumptions and requirements that our approach makes on the hardware. We focus on mainstream CPUs, with vector processing units (VPUs) that support *gather* instructions. The latter make it possible to fetch memory from non-contiguous locations using only SIMD instructions¹

The semantics of *gather* are as follows: let W be the vector length, which is the maximum number of elements that each vector register can hold. The parameters to the instruction are a vector register (I) that holds W indexes and an array pointer (A). As output, *gather* returns a vector register (O) with the W values of the array at the respective indexes. It is important to note that *gather* does not parallelize the memory accesses; the memory system can only serve a few requests at a time. Instead, its usefulness lies in the fact that it can be used to obtain values from non-contiguous memory locations using only SIMD code. This increases the flexibility of the SIMD model and allows to efficiently employ it for workloads previously not considered, i.e., where the memory access patterns are irregular. The alternative is to load the values using scalar code, then transfer them one by one from the scalar registers into vector

registers. Generally, switching between scalar and vector code is not efficient [18, 17].

Apart from *gather*, the rest of the instructions we use can be found across almost all the vector instruction sets available. Worth mentioning is the *shuffle* instruction, that makes it possible to permute individual elements within the vector register in any desired order. For example, we employ it for handling the input and output of the algorithm (cf. Section IV-B).

The size of the cache, especially the L1 and L2, is very important for the algorithmic design, as we describe later in Section IV. Common sizes in modern architectures is 32 KB of L1 data cache with 256 KB of L2 cache and we will use this as a running example. Our design is applicable to other cache sizes as well.

IV. ALGORITHMIC DESIGN

In this section, we begin by introducing S-PATCH, an efficient algorithmic design for multiple pattern matching. It is designed with both cache locality and vectorizability in mind. Next, we propose our vectorization approach V-PATCH, Vectorized PATtern matCHing.

A. S-PATCH: a vectorizable version of DFC

To enable efficient vectorization, we introduce significant modifications to the original DFC design. The key insight for the modifications, explained later in detail, is that small patterns will be found frequently in real traffic, so they should be identified quickly without adding too much overhead. On the other hand, long patterns are found less frequently, but detecting them takes longer and requires more characters from the input to pinpoint them accurately.

As the original DFC, our approach has two parts, organized as two separate rounds. In the **filtering** round, we examine the whole input and feed it through a series of filters that bear some similarities to DFC, but adapted to consider properties of realistic traffic, as motivated above. The **verification** round is as in DFC and performs exact matching on the full patterns that are stored in hash tables. Compared with DFC, S-PATCH focuses on efficient filtering in the first round, because this is the computationally intensive part of the algorithm that, as we show, can be efficiently vectorized. Splitting the two parts in separate rounds improves cache locality, since the data structures used in each round do not evict each other and, as shown in Section IV-B, makes vectorization more practical.

1) *Filtering*: In this first phase the goals are to (i) quickly eliminate the parts of the input that cannot generate a match and (ii) store the input positions where there is indication for a match. In general, key properties of the filtering phase include:

- Good filtering rate. A big fraction of the input is filtered out at this stage.

¹In Intel processors, the *gather* instruction was introduced with the AVX2 instruction set and is included in the latest family of mainstream processors (Haswell and Broadwell); *gather* also exists in other architectures, such as the Xeon Phi co-processor [5].

Figure 1. Filter Design of S-PATCH. HT stands for the *Hash Tables* that contain the full patterns.

- Low overhead. Every filter introduces additional computations and memory accesses, so there needs to be a balance between its overhead and the amount of input that is filtered out.
- Size-efficiency. All the filters need to fit in L1 or L2 cache, while also leaving room for the input and the array for the intermediate results in cache. This is very important, because it ensures that the lookups on the filters will be fast and, as explained later, vectorization using the gather instruction will be feasible.

Our proposed filter design (cf. Figure 1) consists of three filters, each with a specific purpose. The first one stores information about the short patterns (less than 4 characters). It has one bit for every possible combination of two characters, and if a particular combination is the beginning of a pattern, the corresponding bit is set. Similarly, the second filter uses the same indexing and accounts for the longer patterns together with the third filter. In more detail, (cf. also Algorithm 1):

First filter. In the first part of the filtering, we examine two bytes of the input at a time and use them to calculate an index for filters 1 and 2. If the corresponding bit in the first filter is set, we directly store the current input position in an array for further processing (lines 5-7).

Second filter. We also perform a lookup on the second filter using the same index, at line 8. A hit may indicate that we have a match with a longer pattern, but it may also be a false positive (e.g. compare the strings “attribute” and “attack”). Thus, before storing the current input position after a match with the second filter, the algorithm uses more bytes (in our case four) from the input stream with a third filter to gain stronger indications whether there is actually a match. Only when the match in the second filter is corroborated with a match from the third filter is the current position in the input stream stored for further processing (line 11).

Third filter. For the third filter, the index is calculated differently; we cannot have a filter with all combinations of four bytes, due to cache-size limitations. Instead, we use a multiplicative hash function for the four bytes of input to compute the index in the filter, at line 9. There is a trade-off between having a large enough filter to avoid collisions (thus providing a good filtering rate) and having it small enough to fit in cache. The reason why we choose four bytes as input will become clear in the next section (4 bytes fit in

each one of the 32-bit vector register values).

Note that the performance of the filtering phase is intrinsically tied to the filter designs and the type of input. The reason why our proposed design is more effective is twofold. Short patterns, although few,² are likely to generate many matches. As an example, if strings like GET and HTTP are part of the pattern set, they will frequently be found in real network traffic. Treating them separately in a dedicated filter allows us to focus on the longer patterns in other filters. Long patterns, found more rarely, require more information to be distinguished from innocuous traffic.

2) *Verification:* After the filtering, all the possible match positions in the input have been stored in a temporary array. At this point, we need to compare the input at these positions with the actual patterns, before we can safely report a match. As mentioned before, the verification phase is as described by Choi et al. [12], except that it is now done in a separate round, after the current chunk of input has been processed by the filtering phase. For ease of reference we paraphrase here.

Among several optimizations, Choi et al. [12] use specially designed *compact hash tables* that are different for different pattern lengths. Translated to our improved filtering design, if the input at some position i passed the filtering, in the verification phase the algorithm will perform a match on the compact hash table that stores references to all the patterns of appropriate size. For example, if i passed the third filter that stores information on patterns that are four bytes or longer, in the verification phase, the algorithm performs a match on the compact hash table that stores patterns of four bytes or longer (lines 18-20). Each hash table is indexed with as many bytes as the shortest pattern that the hash table contains (in this case, four bytes of the input will be used as an index to the hash table). Each bucket in the hash table contains references to the full patterns and the algorithm has to compare each one of them individually with the input, before reporting a match. Eventually, the algorithm identifies all the occurrences of all the patterns, producing the same output as Aho-Corasick.

In general, the compact hash tables as we use them in this phase, do not fit L1 or L2 cache (but they might fit L3 cache) and accessing them incurs high latency misses. However, the success of the approach lies in the fact that the filtering phase will reject most of the input, so the algorithm resorts to verification only when it is needed (when there is a high probability for a match). That is why our efforts focus on the filtering part, where the data structures are close to the processor and can benefit from vectorization.

B. V-PATCH: Vectorized algorithmic design

A basic issue when vectorizing S-PATCH is its non-contiguous memory accesses. The sequential version accesses the filters at nonadjacent locations for every window

²21% of Snort’s v2.9.7 patterns are 1-4 bytes long [12].

Data: D: data to inspect

```

1 # A_short : temporary array for short patterns
2 # A_long : temporary array for long patterns
3 for i=0, i < D.length, i++ do
4   index = Read two bytes from pos i in D
5   if (Filter1[index] is set) then
6     Store i in A_short
7   end
8   if (Filter2[index] is set) then
9     new_index = hash 4 bytes from input
10    if Filter3[new_index] is set then
11      Store i in A_long
12    end
13  end
14 end
15 for i=0, i < A_short.length, i++ do
16   Verification for small patterns
17 end
18 for i=0, i < A_long.length, i++ do
19   Verification for big patterns
20 end

```

Algorithm 1: Pseudocode for S-PATCH.

Figure 2. Input Transformation from consecutive characters to sliding windows of two characters.

of two characters, whereas in a vectorized design W indexes are stored in a vector register (of length W), each pointing to a separate part of the data structure. For this reason, we use the SIMD *gather* instruction that allows us to fetch values from W separate places in memory and pack them in a vector register.

Algorithm 2 gives a high level summary of the filtering phase of V-PATCH. The first step towards vectorizing the algorithm is loading the consecutive input characters from memory and storing them in the appropriate vector registers. Figure 2 shows the initial layout of the input and the desired transformation to W elements, each holding a sliding window of two characters. The transformation is efficiently achieved with the use of the *shuffle* instruction, allowing to manually reposition bytes in the vector registers (Algorithm 2, line 8).

Once the vector registers are filled, the next step is to calculate the set of indexes for the filters. Note that every 2-byte input value maps to a specific *bit* in the filter, but the memory locations in the filter are addressable in *bytes*. A standard technique used in the literature [19, 12] is to perform a bit-wise right shift of the input value to the corresponding index in the filter. The remainder of the shift indicates which bit to choose from the ones returned.

Data: D: input data to inspect

```

1 # W : the vector register length
2 # A_short : temporary array for short patterns
3 # A_long : temporary array for long patterns
4 #  $\vec{M1}$  : constant mask used to convert the input to 2 byte
   sliding window format
5 #  $\vec{M2}$  : constant mask used to convert the input to 4 byte
   sliding window format
6 for i=0, i < D.length, i += W do
7    $\vec{R}$  = Fill register with raw input from D
8    $\vec{Indexes}$  = shuffle( $\vec{R}, \vec{M1}$ )
9    $\vec{V1}$  = gather(filter1_address,  $\vec{Indexes}$ )
10  if at least one element in  $\vec{V1}$  is set then
11    Store positions of matches in A_short
12  end
13   $\vec{V2}$  = gather(filter2_address,  $\vec{Indexes}$ )
14  if at least one element in  $\vec{V2}$  is set then
15     $\vec{NewIndexes}$  = shuffle( $\vec{R}, \vec{M2}$ )
16     $\vec{Keys}$  = hash( $\vec{NewIndexes}$ )
17     $\vec{V3}$  = gather(filter3_address,  $\vec{Keys}$ )
18    if at least one element in  $\vec{V3}$  is set then
19      Store positions of matches in A_long
20    end
21  end
22 end

```

Algorithm 2: Pseudocode for the V-PATCH filtering phase.

Having computed the indexes, we use them as arguments to the *gather* instruction that fetches the filter values at those locations (Algorithm 2, lines 9 and 13).

Regarding the number of *gather* instructions used, to optimize in latency, note that the first two filters (lines 9 and 13) are specifically designed to use the same indexes for a given input value in *gather* but different base addresses for the filters. Thus, with the **filter merging** optimization where the filters are interleaved in memory (at the same base address), we can merge lines 9 and 13 into a single *gather*, to bring the information from both filters from memory simultaneously. This optimization is not shown in the pseudo-code but depicted in Figure 3, giving an example in which a single *gather* instruction fetches information from both filters. Using bit-wise operations we can choose one filter or the other, once the data is in the vector register.

If at least one of the W values has passed the second filter, they need to be further processed through the third filter. Remember that the third filter uses a window of four input characters as an index. Thus, we load a sliding window of four input characters in each vector element in the register (line 15) and create the hash values that we use as indexes in the third filter (lines 16-17).

Not all of the values in the vector register are useful; only the ones that passed the second filter need to be processed further by the third filter. This is a common challenge when

Figure 3. Figure describing the **filter merging** optimization. In the upper half, lookups on two filters require two gather invocations. Once the filters are merged in memory in the lower half, one gather brings information from both filters to the registers.

vectorizing algorithms with conditional statements, since for different input we need to run different instructions. There are approaches [19] that manipulate the elements in the vector registers, so that they only operate on useful elements. For this particular algorithm, experiments with preliminary implementations showed that the cost of moving the elements in the registers out-weighted the benefits. Thus, we choose to speculatively perform the filtering on all the values and then mask out the ones that do not pass the second filter. In our evaluation (Section V), we observe that operating speculatively on all the elements is actually not a wasteful approach, especially with a large number of patterns to match.

As with the scalar algorithm, after a hit in the first or third filter we need to store the position of the input where a potential match occurred. We store the positions of the input that passed the filter from the set of W values in the register (lines 11 and 19). Here, we postpone the actual verification to avoid a potential costly mix of vectorized and scalar code, where the values from the vector registers need to be written to the stack and from there read into the scalar registers. Such a conversion can be costly and can negate any benefits we gain from vectorization [18].

Furthermore, to fully exploit the available instruction-level parallelism, we manually unroll the main loop of the algorithm by operating on two vectors (R_j) of W values instead of one, a technique that has proven to be efficient especially for SIMD code [19]. This has the benefit that, while the results of a *gather* on one set of W values are fetched from memory (line 9), the pipeline can execute computations on the other set of values in parallel.

V. EVALUATION

In this section, we evaluate the benefits that our vectorization techniques bring to pattern matching algorithms.

Our evaluation criteria are the processing throughput and the performance under varying number of patterns. We show the improvements of V-PATCH with both realistic and synthetic datasets, as well as with changing number of patterns. For a comprehensive evaluation, we compare the results from five different algorithms: the original Aho-Corasick ([11]; implementation directly taken from the Snort source code [14]), DFC (Choi et al. [12], summarized in Section II-B), Vector-DFC (a direct vectorization of DFC done by us), S-PATCH (the scalar version of our algorithm, described in Section IV-A, that facilitates vectorization and addresses properties of realistic traffic that were not addressed before), and V-PATCH (the final vectorized algorithm described in Section IV-B).

A. Experimental setup

Systems. For the evaluation we use both Intel Haswell and Xeon-Phi. More specifically, the first system is an Intel Xeon E5-2695 (Haswell) CPU with 32KB of L1 data cache, 256KB of L2 cache and 35MB of L3 cache. We use the ICC compiler (version 16.0.3) with -O3 optimization under the operating system CentOS. Unless otherwise noted, the experiments in this section are run on this platform. The second system is the Intel Xeon-Phi 3120 co-processor platform. Xeon-Phi has 57 simple, in-order cores at 1.1 GHz each, with 512-bit vector processing units. The memory subsystem includes a L1 data cache and a L2 cache (32KB and 512KB respectively) private to each core, as well as a 6GB GDDR5 memory, but no L3 cache. We compile with ICC -O3 (version 16.0.3) under embedded Linux 2.6. We are only using Xeon-Phi in native mode as a co-processor. The next versions of Xeon-Phi are standalone processors, so the problem of processor-to-co-processor communication is alleviated. Since different hardware threads can operate independently on different parts of the stream, in our experiments with both platforms, we focus on the speedup achieved by a single hardware thread, through vectorization.

Patterns. We use two sets of patterns: a smaller one, named *S1*, consisting of approximately 2,500 patterns that comes with the standard distribution of Snort³ [20] – the de-facto standard for network intrusion detection systems – and a larger one, named *S2*, with approximately 20,000 patterns, that is distributed by emergingthreats.net. The patterns affect the performance of the algorithm and this is analyzed in detail in Section V-C.

Data sets. In our evaluation, we use both real-world traces and synthetic data-sets. The real-world traces are the ICSX dataset [21, 22] (created to evaluate intrusion detection systems) and the DARPA intrusion detection dataset [23]. From ICSX, we randomly take 1GB of data from each of days 2 and 6 (thereafter named ICSX day 2 and ICSX day 6, respectively) and we also use 300MB of data from

³We used version 2.9.7 for our experiments.

the DARPA 2000 capture. We are aware of the artifacts in the latter set, and the discussions in the community about its suitability for measuring the *detection capability* of intrusion detection systems [24]. In our experiments, we use it only for the purpose of comparing throughput between algorithms, allowing for future comparisons on a known dataset. The synthetic data set consists of 1GB of randomly generated characters.

An important point, considering the evaluation validity, is that, typically, not all the patterns are evaluated at the same time. In a Network Intrusion Detection System such as Snort, patterns are organized in groups, depending on the type of traffic they refer to. When traffic arrives in the system, the reassembled payload is matched only against patterns that are relevant (e.g. if the stream has HTTP traffic, it is checked against HTTP related patterns, as well as more general patterns that do not refer to a specific protocol or service). To evaluate our algorithm in a realistic setting, we also pair traffic with relevant patterns. Since, in our datasets, most of the traffic is HTTP [21], we focus on HTTP traffic and match it against the patterns that are applicable based on the rule definitions. A similar approach can be used for other protocols (e.g. DNS, FTP), but we focus on HTTP traffic as it typically dominates the traffic mix and many attacks use HTTP as a vector of infection.

B. Overall Throughput

In this section we compare the overall performance between the different algorithms. Using the HTTP-related patterns of each set gives us 2K patterns from pattern set *S1* and 9K patterns from pattern set *S2*. All algorithms count the number of matches. We use 10 independent runs of each experiment. We report the average throughput values, as well as standard deviation as error bars.

Figure 4a shows the throughput of all algorithms under realistic traffic traces and synthetic traces, when matched against the small pattern set (*S1*). In Figure 4b we use the bigger pattern set (*S2*). The numbers above the bars indicate the relative speedup compared to the original DFC algorithm.

We first discuss the results by only considering each pattern set and each traffic set separately. For realistic traffic traces, our vectorized implementation consistently outperforms the DFC algorithm by up to 1.86x (left parts of Figure 4), due to the parallelization we introduce in the filtering phase. The direct vectorization of the original DFC algorithm (Vector-DFC) has limited performance gain, because much of the running time of DFC is spent on verification and not filtering. This is the main motivation for introducing a modified version of DFC, in Section IV-A, focused on improving the filtering phase. By treating small, frequently occurring patterns separately and by examining more information in the case of long patterns, S-PATCH outperforms the original by up to 1.47x. More importantly,

Figure 4. Performance comparison between the different algorithms for public and random data sets, on the Xeon platform.

it allows for much greater vectorization potential, since the biggest portion of the algorithm’s running time is shifted to efficient filtering of the input, and verification is done much more seldom.

Next, we evaluate the impact of the size of the ruleset on the overall throughput (comparing Figure 4a with Figure 4b). The overall throughput of the algorithms decreases, since the input is more likely to match and identifying every match consumes extra cycles. The performance of Aho-Corasick, in particular, decreases by more than 40%, because the extra patterns greatly increase the size of the state machine. The rest of the algorithms experience a 23-34% drop in performance.

It is important to note that the performance gain of the algorithms (DFC versus Aho-Corasick, V-PATCH versus DFC) is influenced by the input as follows: when feeding the algorithms a data set that contains random strings, DFC significantly outperforms AC (right part of Figure 4). In this case, we do not expect to find many matches in the input and the filtering phase will quickly filter out up to 95% of the input. This is also the reason why the modified versions of the algorithm (S-PATCH and V-PATCH) perform less efficiently compared to what they do in the different input scenarios; the design of the two separate filters as described in Section IV shows its benefits in more realistic traffic mixes. In turn, this poses interesting questions for the future in how to best design the filters based on the expected traffic mix. Still, the vectorized versions provides speedups over the scalar ones.

C. The effects of the number of patterns

As shown in Section V-B, it is important to account for the actual traffic mix the algorithms are expected to run upon when designing the filtering stage, as it has a large impact

Figure 5. Figure a) compares the scalar and vectorized versions of our approach, as the number of patterns increases. Figure b) shows the filtering-to-verification ratio (left axis), as well as the average number of useful elements in the vector registers after filter 2 (right axis), as the number of patterns increases. Figure c) compares the scalar and vectorized approach, as the fraction of matches in the input increases.

on the performance. As new threats emerge, more malicious patterns are introduced and the performance of the algorithm must adapt to that change.

We measure the effects of the number of patterns on the two best performing algorithms and summarize the results in Figure 5a, also including the overall speedup of V-PATCH compared to S-PATCH. In this experiment, we randomly select the number of patterns from the complete set S_2 (20,000 patterns) in order to test our algorithms with as many patterns as possible. V-PATCH consistently performs better compared to S-PATCH, regardless of the number of patterns considered. Observe that:

- As the number of patterns increases, so does the input fraction that passes the filters. This causes the verification part, which is not vectorized, to take up more of the running time, essentially reducing the parallel portion and, by Amdahl's law [25], the benefit of vectorization. The portion of the running time spent in filtering, over the total running time is shown in Figure 5b (blue line).
- As the number of patterns increases, the vectorization of the filtering becomes more efficient. Remember that V-PATCH will proceed with the third filter if at least one of the values in the vector register block passes the second filter. With a small number of patterns, we will seldom pass the second filter. When we do, it is likely we only have a single match, meaning that the rest of the values in the register are disabled and any computation performed for those values is wasteful work. Increasing the number of patterns results in more potential matches in the second filter and, as a consequence, less disabled values for the third filter and thus more useful work. In Figure 5b (red line) we measure this effect and show the average number of useful items inside the vector register every time we reach the third filter. Clearly, with an increasing number of patterns, the vectorization is performed mainly on useful data and

therefore becomes more efficient.

- The two trends essentially cancel each other out, keeping the overall performance benefit of V-PATCH compared to S-PATCH constant after a point (Figure 5a), even though the optimized filtering gradually becomes a smaller part of the total running time.
- A similar effect is observed when we keep the number of patterns constant, but increase the amount of matches in the dataset (Figure 5c). For this experiment, we created a synthetic input that contains increasingly more patterns, randomly selected from a ruleset of 2,000 patterns. As more matching strings are inserted into the input, our vectorized portion of the algorithm becomes more efficient and the relative speedup compared to the scalar version slowly increases.

D. Filtering Parallelism

In this section, in order to gain better insights about the benefits of vectorization, we measure the speedup gained in the filtering part in isolation. Figure 6 compares the filtering throughput of the scalar S-PATCH and V-PATCH, for pattern sets S_1 , S_2 , as well as the full pattern set (20K patterns). In the same figure, we also report the performance of the vectorized filtering, where we exclude the cost of storing the matches in the filtering phase in the temporary arrays. As we can see from the graph, the throughput of the filtering part is increased by up to a factor of 1.84x, on the small pattern set. Storing the matches of the filtering part in arrays comes with a cost; when it is removed, performance increases up to 2.15x for small pattern sets and up to 2.80x for the full pattern set. Even though there is a small decrease at the pattern set with 9K patterns (Figure 6b), the relative speedups of vectorized filtering increase with the number of patterns (Figure 6c).

Figure 6. Measuring the performance of the filtering part only. “V-PATCH-filtering+stores” includes the cost of storing the results of the filtering phase to temporary arrays

E. Changing the vector length: Results from Xeon-Phi

We have also evaluated the effectiveness of our approach on an architecture with a wider vector processing pipeline. The Xeon-Phi [5] co-processor from Intel supports vector instructions that operate on 512-bit registers, thus able to perform two times more operations in parallel, in the filtering phase.

Figure 7 summarizes the results from Xeon Phi, where the experiments are identical with those described in Section V-B. Note that we report the throughput of a single Xeon-Phi thread. V-PATCH takes advantage of the wider vector registers and outperforms the original scalar DFC algorithm, up to a factor of 3.6x on real data and 3.5x on synthetic random data.

As Xeon-Phi threads have much slower clock (1.1 GHz) and the pipeline is less sophisticated (e.g. there is no out-of-order execution), it is not surprising that the absolute throughput sustained by a single Phi thread is smaller than that of the single thread performance of the Xeon platform used in the previous experiments. When dealing with multiple streams in parallel, due to the higher degree of parallelism, the aggregated gain will naturally be higher.

An interesting observation is that the DFC algorithm is slightly slower than AC on real data, where the number of matches in the input is significantly higher. In the original DFC algorithm, the filters are small and can easily fit L1 or L2 cache, and the hash tables containing the patterns are bigger, but still expected to fit L3 cache. In Xeon-Phi there is no L3 cache, so accesses to the hash tables in the verification phase are typically served by the device memory,

Figure 7. Performance comparison between the different algorithms for public and random data sets on the Xeon-Phi platform.

negating the benefits of cache locality that is part of the main idea of the algorithm. Nonetheless, our *improved filtering* design reduces the number of times we resort to verification and access the device memory, thus resulting in 1.1x-1.5x increased throughput on realistic traffic, compared to the original DFC design.

VI. OTHER RELATED WORK

A. Pattern matching algorithms

Pattern matching has been an active field of research for many years and there are numerous proposed approaches. Aho-Corasick, explained before in Section II-A is one of the fundamental algorithms in the fields. There are variants of Aho-Corasick that decrease the size of the state transition table (for example [26]) by changing the way it is mapped in memory, but they come at an increased search cost, compared to the standard version of Aho-Corasick used in our evaluation. Other approaches apply heuristics that enable the algorithm to skip some of the input bytes without examining them at all, such as Wu-Manber [27] where a table is used to store information of how many bytes one can skip in the input. The main issue with these approaches is that they perform poorly with short patterns. For the problem domain investigated here, the patterns can be of any length and the algorithm must handle all of them gracefully. Moreover, in both Aho-Corasick and Wu-Manber algorithms, there is no data parallelism because there are dependencies between different iterations of the main loop over the input.

Recent algorithms [12, 13] follow a different idea: Using small data structures that hold information from the patterns (directly addressable bitmaps in the case of [12], Bloom filters in the case of [13]), they quickly filter out the biggest parts of the input that will not match any patterns

and fallback to expensive verification when there is an indication for a match. Our work is inspired by this family of algorithms, showing how they can be modified to perform better under realistic traffic and gain significant benefit from vectorization.

B. SIMD approaches to pattern matching

Even though pattern matching algorithms are characterized by random access patterns, SIMD approaches have been used before for pattern matching, especially in the field of regular expression matching. Mytkowicz et al. [8] enumerate all the possible state transitions for a given byte of input to break data dependencies when traversing the DFA. Then they use the *shuffle* instruction to implement gathers and to compute the next set of states in the DFA. The algorithm is applied on the case where the input is matched against a single regular expression with a few hundreds of states and does not scale for the case of multiple pattern matching where we need to access thousands of states for every byte of input. Sitaridi et al. [9] use the same hardware gathers as we do, but apply them on database applications where the multiple, independent strings need to be matched against a single regular expression. There have been approaches that use other SIMD instructions for multiple exact pattern matching, but have constraints that make them impractical for the case of Network Intrusion Detection. Faro et al. [28] create fingerprints from patterns and hash them, but they require that the patterns are long, which is not always true for the typical set of patterns found e.g. in Snort.

C. Other architectures

Outside the range of approaches that target commodity hardware, there is rich literature on network intrusion detections systems that are customised for specific hardware. For example, SIMD approaches that target DFA-based algorithms have been applied on the Cell processor [29], as well as GPUs and FPGAs [30, 31, 32]. Vasiliadis et al. [30] build a GPU-based intrusion detection system that uses Aho-Corasick as the core pattern matching engine. Kouzinopoulos and Margaritis also experiment with pattern matching algorithms on GPUs and apply them on genome sequence analysis [31]. GPU parallelization has many similarities with vectorization; in fact GPUs offer more parallelism that can hide memory latencies. At the same time, it introduces additional challenges e.g. long latencies when transferring data between the host and the GPU. In this work we utilize vector pipelines that are already part of modern commodity architectures. Moreover, vectorization with CPUs requires careful algorithmic design that makes use of caches and advanced SIMD instructions. A main part of our work is showing how this problem can be tackled for the case of intrusion detection.

VII. CONCLUSION

In this paper, we introduce an efficient algorithmic design for multiple pattern matching which ensures cache locality and utilizes modern SIMD instructions. Specifically, we introduce V-PATCH: it employs carefully designed and engineered vectorization and cache locality for accelerated pattern matching and nearly doubles the performance when compared to the state of the art.

We thoroughly evaluate V-PATCH and its algorithmic design with both open data sets of real-world network traffic and synthetic ones in the context of network intrusion detection. Our results on Haswell and Xeon-Phi show a speedup of 1.8x and 3.6x, respectively, over single thread performance of Direct Filter Classification (DFC), a recently proposed algorithm by Choi et al. for pattern matching exploiting cache locality, and a speedup of more than 2.3x over Aho-Corasick, a widely used algorithm in today's Intrusion Detection Systems. Moreover, we show that the performance improvements of V-PATCH over the state of the art hold across open, realistic data sets, regardless of the number of patterns in the chosen ruleset. The experimental study also provides insights about the net effect of vectorization as well as trade-offs it implies in this family of algorithmic designs.

ACKNOWLEDGEMENTS

The research leading to these results has been partially supported by the Swedish Energy Agency under the program Energy, IT and Design, the Swedish Civil Contingencies Agency (MSB) through the project "RICS", and by the Swedish Foundation for Strategic Research (SSF) through the framework project FiC and the project LoWi, and by the Swedish Research Council (VR) through the project ChaosNet.

REFERENCES

- [1] S. Antonatos, K. G. Anagnostakis, and E. P. Markatos, "Generating realistic workloads for network intrusion detection systems," *SIGSOFT Softw. Eng. Notes*, vol. 29, 2004.
- [2] J. B. D. Cabrera, J. Gosar, W. Lee, and R. K. Mehra, "On the statistical distribution of processing times in network intrusion detection," in *2004 43rd IEEE Conf. on Decision and Control (CDC)*, vol. 1, 2004.
- [3] R. Mijumbi, J. Serrat, J.-L. Gorricho, N. Bouten, F. De Turck, and R. Boutaba, "Network function virtualization: State-of-the-art and research challenges," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, 2015.
- [4] Y. Li and M. Chen, "Software-defined network function virtualization: a survey," *IEEE Access*, vol. 3, 2015.
- [5] "Intel Xeon Phi product family," <http://www.intel.com/content/www/us/en/processors/xeon/xeon-phi-detail.html>, accessed: 2016-12-10.

- [6] "Intel vectorization tools," <https://software.intel.com/en-us/articles/intel-vectorization-tools>, accessed: 2016-12-10.
- [7] "The importance of vectorization for Intel Many Integrated Core Architecture (Intel MIC architecture)," <https://software.intel.com/en-us/articles/the-importance-of-vectorization-for-intel-many-integrated-core-architecture-intel-mic>, accessed: 2016-12-10.
- [8] T. Mytkowicz, M. Musuvathi, and W. Schulte, "Data-parallel finite-state machines," in *Proc. of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS '14. ACM, 2014.
- [9] E. Sitaridi, O. Polychroniou, and K. A. Ross, "SIMD-accelerated regular expression matching," in *Proc. of the 12th Int. Workshop on Data Management on New Hardware*, ser. DaMoN '16. ACM, 2016.
- [10] P. Jiang and G. Agrawal, "Combining SIMD and Many/Multi-core parallelism for finite state machines with enumerative speculation," in *Proceedings of the 22nd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, ser. PPOPP '17. ACM, 2017.
- [11] A. V. Aho and M. J. Corasick, "Efficient string matching: An aid to bibliographic search," *Commun. ACM*, vol. 18, no. 6, Jun. 1975.
- [12] B. Choi, J. Chae, M. Jamshed, K. Park, and D. Han, "DFC: Accelerating string pattern matching for network applications," in *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*. USENIX Association, 2016.
- [13] I. Moraru and D. G. Andersen, "Exact pattern matching with feed-forward bloom filters," *J. Exp. Algorithmics*, vol. 17, Sep. 2012.
- [14] "Snort rules and IDS software download," <https://www.snort.org/downloads>, accessed: 2016-12-10.
- [15] "Scaling CloudFlare's massive WAF," <https://www.scalescale.com/scaling-cloudflares-massive-waf/>, accessed: 2016-12-10.
- [16] "Gather Scatter operations," <http://insidehpc.com/2015/05/gather-scatter-operations/>, accessed: 2016-12-10.
- [17] O. Polychroniou, A. Raghavan, and K. A. Ross, "Rethinking SIMD vectorization for in-memory databases," in *Proc. of the 2015 ACM SIGMOD Int. Conf. on Management of Data*, ser. SIGMOD '15. ACM, 2015.
- [18] J. Hofmann, J. Treibig, G. Hager, and G. Wellein, "Comparing the performance of different x86 SIMD instruction sets for a medical imaging application on modern multi- and manycore chips," in *Proc. of the 2014 Workshop on Programming Models for SIMD/Vector Processing*, ser. WPMVP '14. ACM, 2014.
- [19] O. Polychroniou and K. A. Ross, "Vectorized Bloom filters for advanced SIMD processors," in *Proc. of the Tenth Int. Workshop on Data Management on New Hardware*, ser. DaMoN '14. ACM, 2014.
- [20] M. Roesch, "Snort - lightweight intrusion detection for networks," in *Proc. of the 13th USENIX Conf. on System Administration*. USENIX Association, 1999.
- [21] A. Shiravi, H. Shiravi, M. Tavallaee, and A. A. Ghorbani, "Toward developing a systematic approach to generate benchmark datasets for intrusion detection," *Computers & Security*, vol. 31, no. 3, 2012.
- [22] "UNB ISCIX intrusion detection evaluation dataset," <http://www.unb.ca/research/iscx/dataset/iscx-IDS-dataset.html>, accessed: 2016-12-10.
- [23] "DARPA intrusion detection data sets," <https://www.ll.mit.edu/ideval/data/>, accessed: 2016-12-10.
- [24] M. V. Mahoney and P. K. Chan, "An analysis of the 1999 DARPA/Lincoln Laboratory evaluation data for network anomaly detection," in *Int. Workshop on Recent Advances in Intrusion Detection*. Springer, 2003.
- [25] G. M. Amdahl, "Validity of the single processor approach to achieving large scale computing capabilities," in *Proc. of the April 18-20, 1967, Spring Joint Computer Conference*, ser. AFIPS '67 (Spring). ACM, 1967.
- [26] M. Norton, "Optimizing pattern matching for intrusion detection," *Sourcefire, Inc., Columbia, MD*, 2004.
- [27] S. Wu and U. Manber, "A fast algorithm for multi-pattern searching," University of Arizona. Department of Computer Science, Tech. Rep. TR-94-17, 1994.
- [28] S. Faro and M. O. Külekci, *Fast Multiple String Matching Using Streaming SIMD Extensions Technology*. Springer, 2012.
- [29] D. P. Scarpazza, O. Villa, and F. Petrini, "Peak-performance DFA-based string matching on the Cell processor," in *2007 IEEE International Parallel and Distributed Processing Symposium*, 2007.
- [30] G. Vasiliadis, S. Antonatos, M. Polychronakis, E. P. Markatos, and S. Ioannidis, *Gnort: High Performance Network Intrusion Detection Using Graphics Processors*. Springer, 2008.
- [31] C. S. Kouzinopoulos and K. G. Margaritis, "String matching on a multicore GPU using CUDA," in *Informatics, PCI'09. 13th Panhellenic Con. on.* IEEE, 2009.
- [32] I. Sourdis and D. Pnevmatikatos, "Pre-decoded CAMs for efficient and high-speed nids pattern matching," in *Field-Programmable Custom Computing Machines, FCCM 2004. 12th Annual IEEE Symposium on*, 2004.