Novel bloom filter algorithms and architectures for ultra-high-speed network security applications

Sateesan, A.; Vliegen, J.; Daemen, J.; Mentens, N.; Trost, A.; Zemva, A.; Skavhaug, A.

Citation

Sateesan, A., Vliegen, J., Daemen, J., & Mentens, N. (2020). Novel bloom filter algorithms and architectures for ultra-high-speed network security applications. *2020 23Rd Euromicro Conference On Digital System Design (Dsd)*, 262-269. doi:10.1109/DSD51259.2020.00050

Version:	Publisher's Version
License:	Licensed under Article 25fa Copyright Act/Law (Amendment Taverne)
Downloaded from:	https://hdl.handle.net/1887/4092804

Note: To cite this publication please use the final published version (if applicable).

Novel Bloom filter algorithms and architectures for ultra-high-speed network security applications

Arish Sateesan
imec-COSIC/ES&S, ESAT
KU Leuven, Belgium
arish.sateesan@kuleuven.be

Jo Vliegen
imec-COSIC/ES&S, ESAT
KU Leuven, Belgium
jo.vliegen@kuleuven.be

Joan Daemen
Digital Security Group
Radboud University
The Netherlands
joan@cs.ru.nl

Nele Mentens
imec-COSIC/ES&S, ESAT
KU Leuven, Belgium
LIACS, Leiden University
The Netherlands
nele.mentens@kuleuven.be

***Abstract*—This paper proposes novel Bloom filter algorithms and FPGA architectures for high-speed searching applications. A Bloom filter is a memory structure that is used to test whether input search data are present in a table of stored data. Bloom filters are extensively used in network security solutions that apply traffic flow monitoring or deep packet inspection. Improving the speed of Bloom filters can therefore have a significant impact on the speed of many network applications. The most important components determining the speed of Bloom filters are hash functions. While hash functions in Bloom filters do not require strong cryptographic properties, they do need a minimized computational delay. We take on the challenge of developing ultra-high-speed Bloom filters on FPGAs by proposing a new non-cryptographic hash function, called Xoodoo-NC, derived from the cryptographic permutation Xoodoo. Xoodoo-NC is a reduced-round, reduced-state version of Xoodoo, inheriting Xoodoo's desired avalanche properties and low logical depth, resulting in an ultra-low-latency non-cryptographic hash function. We evaluate the performance of Bloom filter architectures based on Xoodoo-NC on a Xilinx UltraScale+ FPGA and we compare the performance and resource occupation to existing Bloom filter implementations. We additionally compare our results to memories that use the built-in CAM cores in Xilinx UltraScale+ FPGAs. Our proposed algorithmic and architectural advances lead to Bloom filters that, to the best of our knowledge, outperform all other FPGA-based solutions.**

I. Introduction

Network applications require efficient lookup algorithms for, e.g., packet forwarding, traffic flow monitoring and security, and deep packet inspection. The most popular data structures for high-speed lookups are Content-addressable Memories (CAMs) and Bloom filters. Whereas a CAM searches for a match between input data and stored data, and returns the address of the matching data, a Bloom filter can only detect a match without returning the address. Another difference is that false positives are possible in a Bloom filter, while a CAM has no false positives nor false negatives. Nevertheless, CAMs perform worse than Bloom filters when it comes to chip area, energy consumption and operating speed, as presented by George Varghese in [1]. Since these are critical implementation properties for network routers [2], Bloom filters are the preferred lookup mechanism when a match address is not needed, and when (a limited number of) false positives are acceptable. This is the case in various network (security) applications, such as deep packet inspection [3], network intrusion detection [4], distributed caching, resource routing, and network measurement infrastructures [5], [6]. Additionally, Bloom filters also provide the benefit of constant-time queries, access refinement, and content anonymization [6]. Applications like flow measurement or traffic mapping [7]–[9] check whether or not the identifier (ID) of a network flow has been recorded earlier. The ID could, for example, be a combination of the host and destination addresses and ports.

In order to adhere to the strong bandwidth and energy requirements in Terabit networks (which are defined as networks with a bandwidth higher than 100 Gbps), hardware implementation platforms, such as FPGAs, are increasingly used in network applications [10]. This motivates our work which studies the design and implementation of high-speed Bloom filters on FPGA. Since the speed of a Bloom filter heavily depends on the speed of the hash functions that process the incoming data, we first concentrate on the design and implementation of a high-speed, hardware-friendly hash function. The desired properties of the non-cryptographic hash functions in Bloom filters are more relaxed than the properties of cryptographic hash functions. Both cryptographic and non-cryptographic hash functions map an input of arbitrary length to an output of fixed length. Good hash functions, cryptographic as well as non-cryptographic: 1) map the inputs as uniformly as possible over the entire set of outputs, 2) are easy to compute, and 3) exhibit the avalanche effect, which says that many output bits change, even when there is only a small change in the input. Cryptographic hash functions additionally require preimage resistance, second-preimage resistance and collision resistance [11]. We propose to start from a cryptographic hash function that possesses all the aforementioned properties, and to simplify it such that the computational delay is minimized, while the non-cryptographic properties are maintained. This leads to the Xoodoo-NC (Xoodoo-non-cryptographic) algorithm, derived from the cryptographic permutation Xoodoo [12]. We use Xoodoo-NC in the Bloom-1 architecture, proposed by Qiao et al. in [13]. We show that our Bloom filter implementation is more efficient with respect to occupied FPGA resources and computational delay in comparison to other Bloom filters and CAM-based lookup architectures on FPGA.

This paper is organized as follows: Sect. II consists of background information on lookup algorithms and architectures in network applications and explains the Bloom-1 architecture

978-1-7281-9535-3/20/$31.00 ©2020 IEEE
DOI 10.1109/DSD51259.2020.00050

Authorized licensed use limited to: Universiteit Leiden. Downloaded on September 04,2024 at 14:31:23 UTC from IEEE Xplore. Restrictions apply.

that we use. Sect. III gives an overview of related work. Sect. IV presents our novel algorithms and architectures as well as the CAM architectures that we compare with. The hardware architectures of the implemented Bloom filters are presented in Sect. V. The results are discussed in Sect. VI. Finally, conclusions are drawn in Sect. VII.

II. Lookup algorithms and architectures

A. Content-addressable Memory

Content-addressable memory (CAM) compares input data with data that are stored in the memory. If the same data are present in the memory, the CAM returns the address of the matching data. In a conventional memory, such as a Random-Access Memory (RAM), the address is used as an input and the data that are stored at that address are returned by the RAM. CAMs are, for example, used in network switches, which do a lookup of the destination MAC address of incoming network traffic to determine the port on which the traffic needs to be sent out. A CAM stores each bit explicitly in a memory cell, just like a RAM does. CAMs contain additional hardware circuits that facilitate the look-up. In recent FPGAs, IP cores are available to build fast and resource-efficient CAM structures based on the embedded RAM in the FPGA [14].

B. Bloom filters

In 1970, Howard Bloom presented space/time trade-offs for lookup algorithms in data structures based on hash coding with allowable errors [2]. The author's name has been linked to the presented algorithms ever since. The Bloom filter is a data structure that is used to assist in determining whether an element is in a set of elements. The filter has two possible outcomes: 1) the element might be in the set, or 2) the element is not in the set. Although the first option can trigger false positives, the second option rules out false negatives.

A standard Bloom filter (SBF) that looks for an input x in a set S, consists of a bit vector of length m, which is a compacted representation of the elements that are present in S, and k hash functions, which are used to map inputs to the bit vector. An example is shown in Fig. 1, where m is 14 and k is 3. Before the SBF can be used, the bit vector is first loaded with zeroes. Then, all k hash functions are applied to each element s in the set S, i.e. $H_{k_i}(s)$ is calculated, with $i \in 1, \ldots, k$. For each hash function, the resulting hash value is used as an index in the bit vector and the corresponding bit is set to 1. This process is repeated for each hash function, and subsequently for each possible $s \in S$. In Fig. 1, this means that 3 bits out of 14 are set to 1 for each loaded element s. In the figure, this process is illustrated for two elements s_1 and s_2 After these steps, the set S has been been successfully loaded into the Bloom filter.

To query the SBF using input x, all hash digests $H_{k_i}(x)$ on x are calculated. Analogous to the initialisation phase, the hash values are used as indices in the bit vector. If all the corresponding bits in the bit vector are set to 1, the SBF returns $x \in S$. Otherwise, $x \notin S$ is returned. As stated above, a membership query can have a false positive result. This happens when, during the initialization of the Bloom filter, the hash values of different elements in the set lead to a 1 being written to the same position in the bit vector.

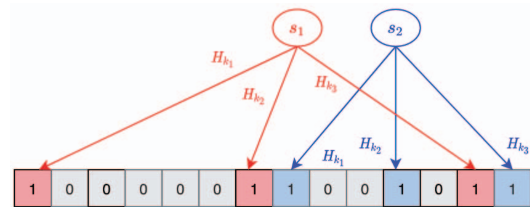

Fig. 1. Standard Bloom filter (SBF) representation

Unlike CAMs, Bloom filters do not explicitly store each element in a memory cell, but rather store the bit vector that represents all elements that are present in the data structure. This leads to more economical memory resource utilisation.

C. Bloom-1 architecture

When the SBF structure in Fig. 1 is used, k locations in the bit vector need to be read out. In order to reduce the number of read-outs, a more efficient structure was proposed by Qiao et al. [13]. They propose Bloom-1, a fast Bloom filter architecture that only requires one read-out for each query. Bloom-1 uses hash functions to map an input element s to $log_2(l) + k \times log_2(w)$ bits. The first $log_2(l)$ bits are the output of the hash function $H_l(s)$ and are used as an address in a data structure that has l memory locations, as shown in Fig. 2. We call each word in the data structure a membership word. The next k values (which are $log_2(w)$ bits each) are the outputs of the hash functions $H_{k_i}(s)$, with $i \in 1, \ldots, k$. They are used as addresses that point to bit locations in the membership word. The size of a membership word is w and a 1 is written at the bit locations that correspond to the addresses $H_{k_i}(s)$. To query an input x, a similar procedure is followed to read out k bits (of which the locations are determined by $H_{k_i}(x)$, with $i \in 1, \ldots, k$) from the membership word at the address determined by $H_l(x)$. If all k bits are 1, the Bloom-1 filter returns $x \in S$. Otherwise, $x \notin S$ is returned.

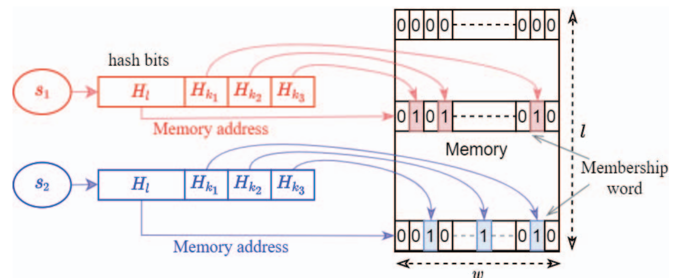

Fig. 2. Bloom-1 representation

Typically, w corresponds to the data bus width in a computer system, i.e., $w = 32$ or 64. The size of the entire Bloom-1 filter is denoted as $m = l \times w$. The advantage of Bloom-1 is that the number of required hash bits is significantly lower when k is

Authorized licensed use limited to: Universiteit Leiden. Downloaded on September 04,2024 at 14:31:23 UTC from IEEE Xplore. Restrictions apply.

large, compared to other Bloom filters [13]. The total number of hash bits for Bloom-1 is $log_2(l) + k \times log_2(w)$, compared to $k \times log_2(m)$ for SBF. Additionally, the w-bit membership word can be read out at once, while the k 1-bit values in SBF require k read-outs. This significantly increases the speed of Bloom-1 compared to SBF.

Qiao et al. gave an approximation of the false positive rate (fpr) for Bloom-1 in [13]. A comment published by Reviriego et al. in [15] refined the approximation, resulting in Eq. 1. Besides the Bloom-1 parameters, l, w and k, the number of entries, n, has an influence on the false positive rate.

$$fpr = \sum_{x=0}^{n} \left(\binom{n}{k} \cdot \left(\frac{1}{l}\right)^{x} \cdot \left(1 - \frac{1}{l}\right)^{n-x} \cdot \left(\frac{w!}{w^{k(x+1)}} \sum_{i=1}^{w} \sum_{j=1}^{i} (-1)^{i-j} \frac{j^{kx} i^{k}}{(w-i)!j!(i-j)!} \right) \right) \quad (1)$$

III. Related work

A wide range of related work is available, elaborating on the design and optimization of Bloom filters for different applications. Various aspects are to be considered when designing hardware architectures of Bloom filters, such as the lookup delay, the false positive rate and the operating frequency.

One of the main factors affecting the lookup delay of a Bloom filter is the execution delay of the hash functions. A number of fast non-cryptographic hash functions with good avalanche properties have been proposed, such as Murmur 3 and FNV-1a. Murmur3 is the result of a general study of non-cryptographic hash functions by Estebanez et al. [16]. FNV-1a was taken from an idea sent as reviewer comments to the IEEE POSIX P1003.2 committee by Fowler and Vo and later improved by Landon Curt Noll [17]. It is considered to have the smallest computational delay. In this work, we propose Xoodoo-NC, a fast hardware-oriented non-cryptographic hash function with good avalanche properties, which outperforms all other solutions in terms of execution delay in hardware. Xoodoo-NC is derived from the cryptographic permutation Xoodoo [12]. It uses a reduced number of rounds and a reduced size of the internal state, which leads to a lower logical depth than cryptographic hash functions based on Xoodoo, while maintaining the desired avalanche properties.

Besides optimizing the hash functions, there are other methods that contribute to the fast generation of the lookup addresses in Bloom filters. Examples are the one-hashing Bloom filter (OHBF) [18], the less hashing Bloom filter (LHBF) [19], and the ultra-fast Bloom filter (UFBF) [20]. OHBF and UHBF follow similar techniques, where a single base hash function is used, and k modulo operations are performed to obtain k addresses. LHBF employs two base hash functions, and a linear equation to generate additional addresses from the base hash functions. Although these techniques can reduce the generation delay of the lookup addresses, the overall delay of a query will still be dominated by the delay of the memory accesses when a standard Bloom filter (SBF) array [2] is used as depicted in Fig. 1. This is because each read/write operation will require one clock cycle when the embedded RAM in an FPGA is used for the storage of the Bloom filter. Hence, it is important to not only reduce the delay for the generation of the lookup addresses, but to also consider Bloom filter architectures that reduce the number of memory accesses.

In an SBF, the number of memory accesses is directly proportional to the number of hash values (k). In order to reduce the false positive rate, the number of hash functions and the size of the data structure (m) must be increased, but this has a negative effect on the execution speed. Parallel Bloom filters, as proposed by Dharmapurikar et al. [3], are a classical solution to overcome this issue. There are other approaches like multi-partitioning counting Bloom filters (MPCBF) [21], OMASS [22] and One-memory-access Bloom filters (Bloom-1) [13], where it takes only one memory access per query. MPCBF, OMASS and Bloom-1 employ similar approaches based on the partitioning of memory blocks to limit the number of memory accesses to one, as explained for Bloom-1 in Sect. II. In this work, we do not focus on counting Bloom filters which store a count value for each memory cell in order to enable the deletion of elements. That is why MPCBF is not considered for our implementation. In OMASS, each element is represented in multiple sets with different hash mappings, which decreases the false positive rate. In this work, we concentrate on Bloom-1 [13], because it requires less memory storage than OMASS. For applications in which the false positive rate of Bloom-1 is not sufficiently low, the possibility of moving to an OMASS architecture can be considered.

A number of works are available on the implementation of Bloom filters in hardware. Dharmapurikar et al. [4] present an implementation of parallel Bloom filters on FPGA for string matching. The implementation is fast in terms of number of memory accesses with acceptable resource utilization, but has a very low operating frequency of 73.5 MHz, which is not suitable for high speed networks. In [23], a rolling hash function based Bloom filter is presented and implemented on a high-end FPGA for fast streaming data. However, the architecture assumes a byte-based interface which is only suitable for Gigabit Ethernet networks. When moving towards Terabit Ethernet, the incoming data need to be processed in larger parallel blocks. The work of Kaya et al. [24] and of Lyons et al. [25] focus mainly on the power reduction of Bloom filter implementations. A counting Bloom filter implementation proposed by [26] for network intrusion detection seems to provide operating frequencies above 200 MHz for an input data size up to 8 bytes. Although all these publications present very good implementations, an even higher parallelization level and/or operating frequency is needed in Terabit Ethernet networks.

In this work, we propose to use the Bloom-1 architecture, which requires only one memory access per query. We also propose Xoodoo-NC, an extremely fast hash function that outputs a hash value large enough to be split into chunks that can directly be used in the Bloom-1 architecture. We implement a speed-optimized hardware architecture that outperforms all other FPGA-based Bloom filter implementations as well as architectures based on the CAM IP cores in Xilinx FPGAs.

Authorized licensed use limited to: Universiteit Leiden. Downloaded on September 04,2024 at 14:31:23 UTC from IEEE Xplore. Restrictions apply.

IV. Novel algorithms and architectures

A. Experimental setting and basis of comparison

We concentrate on Terabit Ethernet (defined as Ethernet with speeds above 100 Gbps). An example of an FPGA platform that supports Terabit Ethernet is the 200Gbps NFB-200G2QL board [27], which operates at a clock frequency of 200 MHz and processes network packets at a line rate of 1024 bits per clock cycle. It is clear that implementing network security applications on Terabit Ethernet FPGA platforms requires the processing of many data bits in parallel. Parallelization is also suggested by the authors of the large flow detection algorithm EARDET [7]. An example of a network encryption implementation that adheres to the strong operating frequency and parallel processing requirements of Terabit FPGA platforms is given in [28].

Typically in network applications flow identifiers need to be queried in a data structure. Examples of such applications are given in Sect. I. The considered flow ID in this work contains the source IP address, destination IP address, and both the source and the destination ports. The size of this flow ID is 96 bits, which is the case in an IPv4 network. To adhere to the requirements of Terabit Ethernet FPGA platforms, these 96 bits need to be processed in parallel at an operating frequency of at least 200 MHz. We compare four different lookup data structures in terms of cycle count, clock frequency, occupied FPGA resources and false positive rate:

1) CAM block generated by the design tools of Xilinx,
2) speed-optimized custom CAM consisting of manually combined smaller CAM blocks,
3) Bloom filter based on the Bloom-1 architecture and the FNV-1a hash function,
4) Bloom filter based on the Bloom-1 architecture and the newly proposed Xoodoo-NC hash function.

All four architectures store 1024 data words of 96 bits each. For the first two architectures, which are CAM-based, this means that exactly 1024 x 96 memory cells are occupied. For the last two architectures, which are based on Bloom filters, we explore different options for the size of the Bloom filter and the number and output size of the hash functions. The number of stored entries in the Bloom filters is equal to 1024 to make a fair comparison with the CAM-based implementations.

The first CAM architecture is directly generated by the design tools of Xilinx, resulting in the utilization of embedded RAM (Block RAM or BRAM) with size 1024 x 96 bits, in combination with a limited amount of logic in the form of FPGA LUTs (lookup tables) and FFs (flip-flops). A lookup in the directly generated CAM takes 2 clock cycles. The second CAM architecture consists of 128 CAM blocks of size 24 x 32 bits. These smaller CAM blocks are again generated by the Xilinx design tools, and the LUTs and FFs for combining and addressing these blocks are added manually. This leads to an architecture with a higher operating frequency, but with a larger resource occupation compared to the directly generated CAM architecture. A lookup in the manually modified CAM takes 3 clock cycles. The third and fourth architecture that we evaluate are Bloom filters, which differ in the utilized hash function. We compare FNV-1a, the fastest hash function in hardware according to our knowledge, to Xoodoo-NC, the proposed hash function in this paper.

For all experiments, the Vivado 2017.4 design tool of Xilinx is used and the UltraScale+ XCVU7P-FLVC2104-1-E FPGA is targeted. Only for the generation of the CAM blocks, the ISE 14.7 design tool of Xilinx is used, because CAM generation is no longer supported in Vivado. The resulting code is imported in Vivado to be used in the UltraScale+ FPGA (which is not supported in ISE 14.7). The operating frequency and resource utilization of the CAM-based architectures that are implemented as a reference for comparison, are given in Fig. 3.

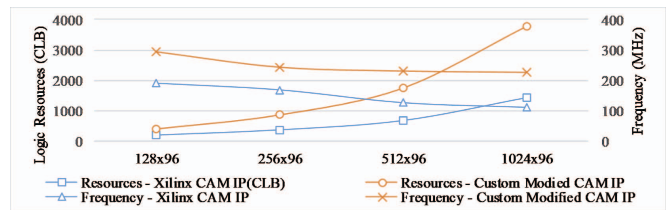

Fig. 3. Operating frequency and resource utilization of the CAM directly implemented with the Xilinx tools (referred to as Xilinx CAM IP) and the speed-optimized CAM architecture (Custom Modified CAM IP)

B. Design of a high-speed non-cryptographic hash function

The optimization goals in the design of a non-cryptographic hash function for Bloom filters are low logical depth or critical path (i.e., high operating frequency), low cycle count, high avalanche score and acceptable resource utilization. FNV-1a [17] and Murmur3 [16] are non-cryptographic hash functions with excellent avalanche properties. They both use multiplication operations in which the width of the multiplier is proportional to the output hash size. FNV-1a and Murmur3 process 8 bits and 32 bits per cycle, respectively, which means the execution time is proportional to the number of 8-bit or 32-bit input blocks. That is why FNV-1a and Murmur3 perform very well on an 8-bit or 32-bit microprocessor. When implemented in hardware, however, the high number of flip-flops to store the intermediate values and the large multipliers have a negative effect on the resource utilization, the operating frequency and the execution delay, especially when the required output hash size is large.

In our effort to find a non-cryptographic hash function that has a low logical depth and a low resource utilization in hardware, we start from the hardware-friendly cryptographic permutation Xoodoo, presented by Daemen et al. in [12]. A Xoodoo round employs only shift, AND and XOR operations, and does not need flip-flops inside the round. While Xoodoo uses a 384-bit state, Xoodoo-NC operates on a 96-bit state, which perfectly fits our needs to process incoming network flow IDs of 96 bits. Output values can be taken as multiples of 96-bits. A more detailed explanation of Xoodoo-NC is given in the following paragraphs.

The original Xoodoo permutation is parameterized by the number of rounds n_r, and it iteratively applies the round

Authorized licensed use limited to: Universiteit Leiden. Downloaded on September 04,2024 at 14:31:23 UTC from IEEE Xplore. Restrictions apply.

function R_i to the state A. The state A is depicted in Fig. 4 and has a size of 48 bytes, divided into three 2-dimensional planes, which are indexed by y. Each plane has four 1-dimensional lanes which are indexed by x. Each lane consists of 32 bits in size, indexed by z. A collection of three parallel stacked lanes in A is called a sheet, and any three parallel stacked bits in A are called a column. Each byte in the state can thus be referred to with the coordinates (x, y, z).

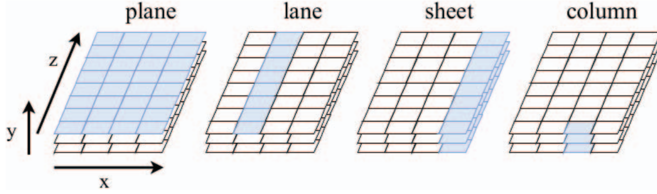

Fig. 4. Graphical representation of the Xoodoo terminology

The round function R_i of Xoodoo consists of 5 sequential steps: a mixing step θ, a plane shifting step ρ_{west}, a step for the addition of round constants ι, a non-linear layer χ, and a second plane shifting step ρ_{east}. For more in-depth details on these steps, we refer to [12].

In this work, we need a hash function with an input size of 96 bits or 12 bytes. Therefore, we reduce the state A of the Xoodoo permutation to 12 bytes, corresponding to one sheet. This fixes the x coordinate and allows the y coordinate to vary from 0 to 2 and the z coordinate from 0 to 31. The original definitions of lanes and columns still hold. Because of the reduction of the size of the state, the Xoodoo round function is slightly adapted. The cyclic shift operations involving the x and z coordinates are replaced by operations involving only z. The other operations and the round constant C_i remain the same as in the original algorithm. The resulting Xoodoo-NC permutation is specified in Alg. 1.

Algorithm 1: Xoodoo-NC[n_r]

Parameters: Number of rounds n_r
for *Round index i from* $1 - n_r$ **to** 0 **do**
  $A = R_i(A)$
end
  Round R_i:
   θ:
    $P \leftarrow A_0 \bigoplus A_1 \bigoplus A_2, E \leftarrow P \lll 5 \bigoplus P \lll 14$
    $A_y \leftarrow A_y \bigoplus E$ for $y \in \{0, 1, 2\}$
   ρ_{west}:
    $A_2 \leftarrow A_2 \lll 11$
   ι:
    $A_0 \leftarrow A_0 \bigoplus C_i$
   χ:
    $B_0 \leftarrow A_1.A_2, B_1 \leftarrow A_2.A_0, B_2 \leftarrow A_0.A_1$
    $A_y \leftarrow A_y \bigoplus B_y$ for $y \in \{0, 1, 2\}$
   ρ_{east}:
    $A_1 \leftarrow A_1 \lll 1, A_2 \leftarrow A_2 \lll 8$

For the Xoodoo permutation, Table 8 in [12] summarizes the avalanche properties. The avalanche metrics mentioned in [12] are used here also to determine the avalanche behaviour of Xoodoo-Nc. The metrics, namely avalanche dependence (D_{av}), avalanche weight ($\overline{w}_{av}$), and avalanche entropy (H_{av}), are calculated as in [12] for single-bit differences at the input. Only the worst-case values are reported, which are the minimum values taken over all individual input differences of a given type. The results are presented in Table I. Xoodoo-NC provides full bit-dependence and quasi-strict avalanche weight after 2.5 rounds, and $\approx$ 90% dependence after round 2. We assume that it is sufficient for our intended application and generate a 96-bit Xoodoo-NC hash output after 2 rounds. To obtain an output of 192 bits, one additional round is executed and the output of round 2 is concatenated with the output of round 3. In a similar way, an output size of any multiple of 96 bits can be obtained by increasing the number of rounds and concatenating the results.

Rounds	D_{av}	$\overline{w}_{av}$	H_{av}
2	84	35.408	80.332
2.5	96	47.324	95.864
3	96	47.309	95.867
3.5	96	47.922	95.996

TABLE I
Avalanche Scores

V. Hardware implementation

We discuss the hardware implementation of the Bloom-1 architectures that are evaluated in this work. We employ a single non-cryptographic hash function of which the output is split into $k + 1$ parts, i.e., one part of size $log_2(l)$ and k parts of size $log_2(w)$, as explained in Sect. II-C. Two non-cryptographic hash functions with an input size of 96 bits are evaluated, namely FNV-1a and Xoodoo-NC. In the following paragraphs, we elaborate on the FNV-1a implementation, the Xoodoo-NC implementation and the Bloom-1 implementation that uses one of these hash functions.

A. Hash functions

1) FNV-1a: FNV-1a can provide output hash sizes ranging from 32 to 1024 bits. It uses only two parameters - a non-zero FNV Offset basis, and an FNV Prime. Both parameters depend on the hash output size. FNV-1a processes one byte at a time, and the algorithm consists of iterations of two sequential operations - XOR and multiplication.

The input can be of any size, but only 8 bits are processed per cycle. The received input is stored in an octet shift register, which shifts a single byte per cycle and has a total size equal to the size of the incoming message, namely q bytes. In this work, the size of the flow ID is 96 bits, which corresponds to $q = 12$. Every clock cycle, a byte is passed to the processing element (PE), which performs the XOR operation with the intermediate hash value, or with the FNV Offset basis for the first byte. Subsequently, a multiplication with the FNV Prime is calculated. The hash_temp register stores the intermediate hash value. The output is obtained after processing the final byte of the input. Before the hash value is valid, zeroes are sent to the output.

The drawback of a straightforward FNV-1a architecture is that it processes only one byte per clock cycle. The input needs to be buffered until the processing is completed. This is a major concern for high-speed networks like Terabit Ethernet networks, in which multiple bytes need to be processed in each

Authorized licensed use limited to: Universiteit Leiden. Downloaded on September 04,2024 at 14:31:23 UTC from IEEE Xplore. Restrictions apply.

clock cycle. For the applications that we consider in this paper, 12 bytes need to be processed every clock cycle, but the FNV-1a architecture requires 12 clock cycles to process 12 bytes. Another concern regarding the hardware implementation is that one of the core operations in the PE is a multiplication. If we want to increase the input size and the output hash size, the multiplier width also needs to be increased, which further increases the hardware complexity and the critical path. Pipelining somewhat helps to mitigate this issue and to improve the overall throughput (after an initial delay equal to q clock cycles. The hardware architecture of the pipelined version is shown in Fig. 5.

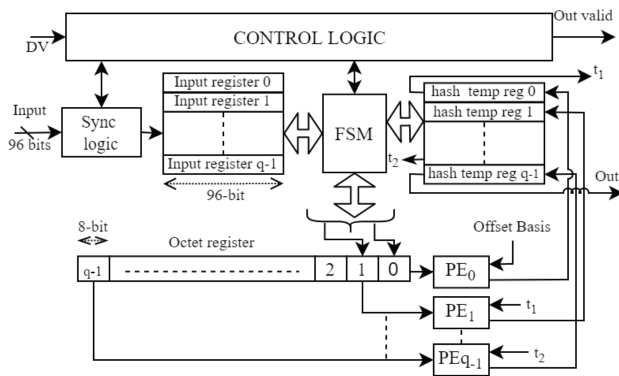

Fig. 5. FNV-1a - Pipelined architecture

The core of the design is a free-running FSM, which processes the n-byte inputs every clock cycle based on the data valid (DV) signal. The input buffers store the incoming q-byte inputs. In total, q input buffers are implemented that each store an q-bit input word. In each clock cycle, the bytes from the input buffers are shifted to the octet register of size n bytes. For every input, the first byte is stored in the 0^{th} register, the second byte in the 1^{st} register and so on. Each 8-bit register in the octet register is connected to the input of a PE. Each PE takes one byte from the octet register and the output from the previous PE (except for PE_0), as inputs. PE_0 takes the FNV Offset basis as the second input. All PEs operate simultaneously. Once the processing is complete, the inputs are removed from the input buffers. The output of each PE is connected to an output buffer that stores the intermediate hash value and is connected to the input of the next PE in the sequence. The validity of the output is determined by the Out_valid signal. The pipelined version of FNV-1a can deal with a new 96-bit input every clock cycle and starts outputting the 8-bit hash values of the 96-bit inputs after an initial delay of 12 clock cycles.

2) Xoodoo-NC: The heart of the Xoodoo permutation is the round function, which consists of the 5 steps shown in Alg. 1: θ, ρ_{west}, ι, χ, and ρ_{east}. The round is executed twice when a 96-bit hash output is needed and once more to generate a second 96-bit output, to be concatenated with the first output when a 192-bit hash output is needed. This way, the output of Xoodoo-NC can be any multiple of 96 bits, but we stick to two different output sizes in our evaluation: 96 and 192 bits.

The rounds are implemented in a fully unrolled architecture without registers in between the rounds, i.e., the 2 or 3 rounds of Xoodoo-NC are completely implemented in combinational logic. The detailed hardware block diagram of the round function is shown in Fig. 6. The 96-bit input A consists of 3 lanes of size 32 bits, namely A_0, A_1, and A_2 (cfr. Alg. 1. The round constant C_i used in ι is a single 32-bit lane, indexed from $1 - n_r$ to 0. The 3-lane output of the round goes to the input of the next round.

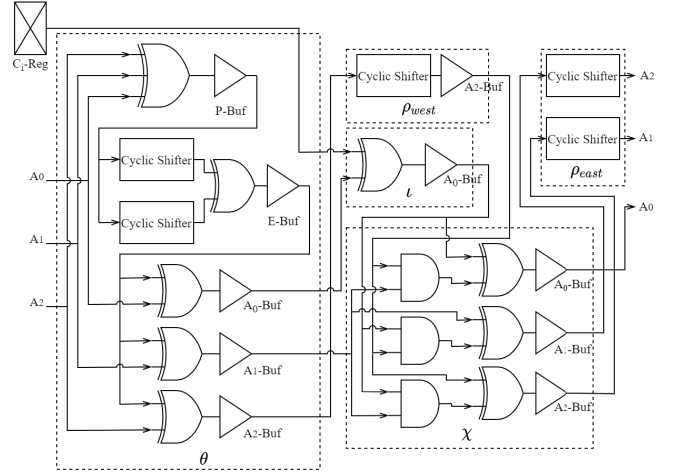

Fig. 6. Xoodoo-NC Round, which is implemented 2 or 3 times in combinational logic only for a 96-bit or 192-bit hash output, respectively.

B. Bloom-1

The hardware architecture of the Bloom-1 filter follows the concept shown in Fig. 2. The memory inside the Bloom-1 filter is implemented using BRAM. The hash bits are generated with a single hash function, of which the output is split into $k + 1$ parts. The first part, H_l, is used as an address for the BRAM. The k remaining parts are used to address the bits inside the memory words. Depending on the query/configure instruction, the memory is read/written. An FSM controls all the operations. The execution delay is determined by the BRAM read-write delay, which is 2 clock cycles, and the delay of the hash generation. Xoodoo-NC is executed in one clock cycle, while the non-pipelined and pipelined versions of FNV-1a take 14 and 2 clock cycles, respectively.

VI. Results

In this section, we present the results of the experiments conducted on a Xilinx Virtex Ultrascale+ FPGA (XCVU7P-FLVC2104-1-E) using the Vivado 2017.4 design tool. For the Bloom filter implementations, we employ FNV-1a with 64-bit and 128-bit hash outputs, and Xoodoo-NC with 96-bit and 192-bit hash outputs. The input size is always 96 bits. For FNV-1a, both the non-pipelined and the pipelined architectures are evaluated. The Bloom-1 parameters l, w, and $m = l * m$ are taken as 4096, 64 and 262144, respectively. The size of the CAMs that we compare with is 1024 x 96.

Authorized licensed use limited to: Universiteit Leiden. Downloaded on September 04,2024 at 14:31:23 UTC from IEEE Xplore. Restrictions apply.

The usage of FPGA logic and DSP units in the Bloom-1 architectures is depicted in Fig. 7. The Xoodoo-NC based Bloom filters outperform the FNV-1a based implementations, especially for larger values of k. The FNV-1a implementations make use of the multipliers inside the DSP units. No DSP units are used in Xoodoo-NC. The pipelined version of FNV-1a has a much higher consumption of logic resources and DSP units in comparison to the serial FNV-1a and the Xoodoo-NC based Bloom filters. The operating frequency and execution delay of the Bloom-1 architectures are shown in Fig. 8. The Xoodoo-NC based architectures clearly feature a much higher operating frequency and a much lower execution delay than the FNV-1a based architectures.

The results are listed in Table II. Additionally, these are shown in Fig. 7 and 8. The table also contains the comparison with the CAM based lookup architectures. The CAM block generated by the Xilinx design tools has an operating frequency of 112.75 MHz, which is not even sufficient for 1-Gbps link speeds. The custom modified CAM architecture operates at a frequency of 225 MHz, but consumes more FPGA resources. Although it takes 1 cycle longer to query the custom modified CAM, the total query time is lower compared to the directly generated CAM. The lookup delay of the Xoodoo-NC based architectures is significantly lower than the lookup delay of all other implemented architectures. This is mainly thanks to the high operating frequency, but also thanks to the cycle count of only 3 cycles. Two versions of the Bloom-1 architectures are considered in the table: those with $k = 2$ and $k = 12$. In order to calculate the false positive rate (fpr) following Eq. 1, the number of entries in the Bloom filter is assumed to be 1024, such that a fair comparison can be made with the CAMs, whose depth is 1024. Going from $k = 2$ to $k = 12$, leads to a resource occupation that is more than doubled for the Xoodoo-NC based implementation, but also drastically decreases the false positive rate. Nevertheless, the resource occupation of the Xoodoo-NC based Bloom filter is much lower than all other considered architectures, even when $k = 12$.

A quantitative comparison with related work is shown in shown in Table III for n=1024. It is clear that our Xoodoo-NC based Bloom-1 architecture with a 96-bit input and $m = 4096 * 64$ significantly outperforms previously proposed FPGA architectures that have a smaller input size, a smaller m and a larger false positive rate. Our design drastically improves the resource occupation and the speed. Although the total lookup delay is not reported for the considered related architectures, the table shows that the delay of the memory accesses only is longer than the delay of one query in our Xoodoo-NC based implementation. Note, however, that the considered related work uses older FPGAs with 4-input LUTs, while our FPGA has 6-input LUTs. Nevertheless, our reported number of LUTs is only around half of the LUTs reported in related work and the number of FFs is roughly one fifth to one tenth of the the FFs reported in related work. The difference in operating frequency can partially be attributed to the difference in silicon technology nodes between older and newer FPGAs. Nevertheless, the frequency of 462 MHz is extremely high thanks to the low logical depth of the Xoodoo-NC hash function.

VII. Conclusion

In this work, we targeted novel algorithms and architectures for high-speed Bloom filters on FPGA, used for fast lookups in network (security) applications. We proposed a new high-speed hardware-oriented non-cryptographic hash function called Xoodoo-NC. The hash function was integrated into a Bloom-1 architecture and evaluated on a Virtex UltraScale+ FPGA of Xilinx. The resulting resource occupation, operating frequency, lookup delay and false positive rate were compared to previously reported Bloom filter architectures and Content-addressable Memories on FPGA. To our knowledge, our work significantly outperforms all other solutions on all these comparison criteria.

Acknowledgement

This work is supported by the ESCALATE project, funded by FWO and SNSF (G0E0719N). We thank Seth Hoffert and the co-designers of Xoodoo, Gilles Van Assche and Ronny Van Keer, for their valuable assistance.

References

[1] G. Varghese. *Network Algorithmics: An Interdisciplinary Approach to Designing Fast Networked Devices*. 2004.
[2] B. H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM*, 13(7):422–426, 1970.
[3] S. Dharmapurikar, P. Krishnamurthy, T. S. Sproull, and J. W. Lockwood. Deep packet inspection using parallel Bloom filters. *IEEE Micro, vol. 24, no. 1*, pages 52–61, 2004.
[4] S. Dharmapurikar, M. Attig, and J. Lockwood. Design and implementation of a string matching system for network intrusion detection using FPGA-based Bloom filters. *12th Annual IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM '04)*, 2004.
[5] A. Z. Broder and M. Mitzenmacher. Network applications of Bloom filters: A survey. *Internet Mathematics*, 2004.
[6] L. Luo, D. Guo, Richard T.B. Ma, O. Rottenstreich, and X. Luo. Optimizing Bloom filter: Challenges, solutions, and comparisons. *arXiv:1804.04777*, 2018.
[7] C. Estan and G. Varghese. New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice. *ACM Transactions on Computer Systems*, 2003.
[8] A. Kumar, J. Xu, J. Wang, O. Spatschek, and L. Li. Space-code Bloom filter for efficient per-flow traffic measurement. *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*, pages 167–172, 2003.
[9] H. Wu, H. C. Hsiao, and Y. C. Hu. Efficient large flow detection over arbitrary windows: An algorithm exact outside an ambiguity region. *Proceedings of the ACM SIGCOMM Internet Measurement Conference, IMC*, pages 209–222, 2014.
[10] Simon Stanley. FPGAs & ASICs for telecom. Heavy Reading, http://www.heavyreading.com/details.asp?sku_id=3418&skuitem_itemid=1657, 2015.
[11] A. J. Menezes, P. C. van Oorschot, and S. A. Vanstone. *Handbook of Applied Cryptography*. CRC Press, 2001.
[12] J. Daemen, S. Hoffert, G. Van Assche, and R. Van Keer. The design of Xoodoo and Xoofff. *IACR Transactions on Symmetric Cryptology*, pages 1–38, 2018.
[13] Y. Qiao, T. Li, and S. Chen. Fast Bloom filters and their generalization. *IEEE TPDS*, 25(1):93–103, 2014.
[14] K. Locke. *Parameterizable Content-Addressable Memory*. https://www.xilinx.com/support/documentation/application_notes/xapp1151_Param_CAM.pdf.

Authorized licensed use limited to: Universiteit Leiden. Downloaded on September 04,2024 at 14:31:23 UTC from IEEE Xplore. Restrictions apply.

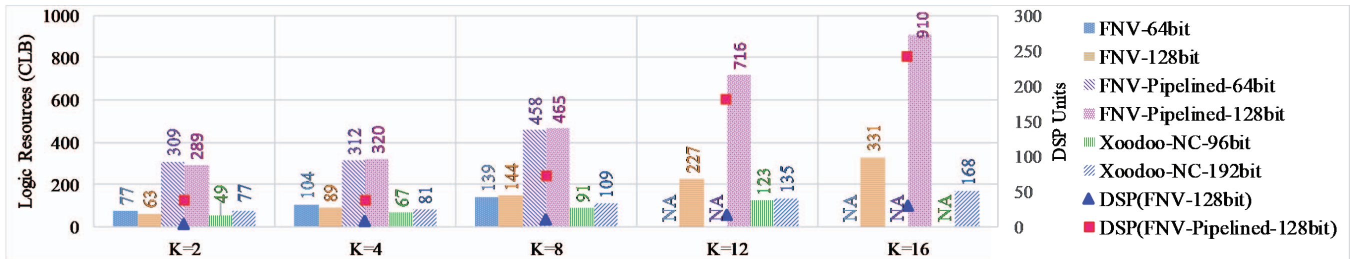

Fig. 7. Resource Utilization of Bloom-1: Logic resources and DSP units

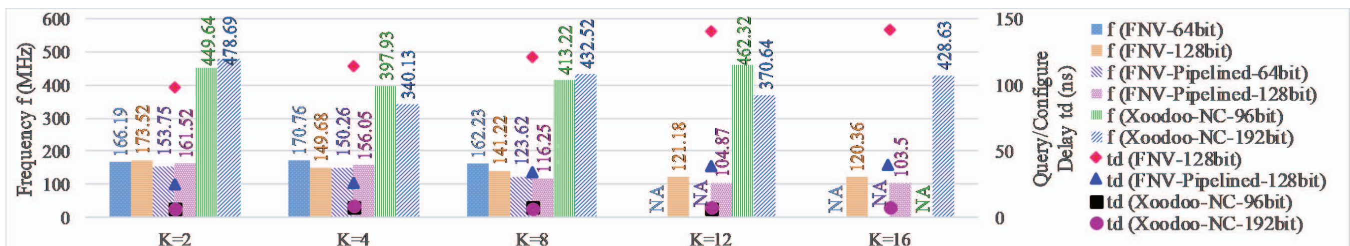

Fig. 8. Maximum Operating Frequency and Delay of Bloom-1

Design	k	fpr	LUTs	FFs	BRAM	LUTRAM	DSP	Max freq [MHz]	No. of cycles for lookup	Lookup Delay [ns]
Xoodoo-NC	**2**	**0.0002976**	**290**	**98**	**7.5**	**0**	**0**	**449.64**	**3**	**6.67**
FNV-1a	2	0.0002976	261	228	7.5	0	3	173.52	16	92.21
FNV-1a Pipelined	2	0.0002976	649	1507	7.5	6	36	151.52	4	24.76
Xoodoo-NC	**12**	**$2.61x10^{-7}$**	**675**	**158**	**7.5**	**0**	**0**	**462.32**	**3**	**6.49**
FNV-1a	12	$2.61x10^{-7}$	1366	339	7.5	0	15	121.18	16	132.03
FNV-1a Pipelined	12	$2.61x10^{-7}$	2706	1562	7.5	6	180	104.87	4	38.14
Direct CAM		0.0	6731	1145	320	1536	0	112.75	2	17.74
Custom Modified CAM		0.0	20'797	988	384	1536	0	225.07	3	13.33

TABLE II
COMPARISON OF BLOOM-1 BASED AND CAM BASED IMPLEMENTATIONS

Design	m	k	input size	LUTs	FFs	BRAMs	Frequency	fpr	No. of memory accesses	FPGA
[4]	4096	10	10 byte	1495	1297	5	73.51 MHz	$8.89x10^{-5}$	2	Virtex 2000E
[26](SBF)	16384	8	8 byte	1058	1058	4	200.6 MHz	$5.74x10^{-4}$	4	Virtex-4
[26](CBF)	16384	8	8 byte	1188	1188	4	201.6 MHz	$5.74x10^{-4}$	4	Virtex-4
Ours (xoodoo-NC)	**262144**	**12**	**12 byte**	**675**	**158**	**7.5**	**462.32 MHz**	**$2.61x10^{-7}$**	**1**	**Virtex UltraScale+**

TABLE III
COMPARISON WITH RELATED WORK

[15] P. Reviriego, K. Christensen, and J.A. Maestro. A comment on fast Bloom filters and their generalization. *IEEE Transactions on Parallel and Distributed Systems*, 27(1):303–304, 2016.

[16] C. Estébanez, Y. Sáez, G. Recio, and P. Isasi. Performance of the most common non-cryptographic hash functions. *Software: Practice and Experience*, 44, 06 2014.

[17] G. Fowler, L. C. Noll, and K.-P. Vo. The FNV non-cryptographic hash algorithm. https://tools.ietf.org/html/draft-eastlake-fnv-03, 2012.

[18] J. Lu, T. Yang, Y. Wang, H. Dai, L. Jin, H. Song, and B. Liu. One-hashing Bloom filter. *Proc. IEEE/ACM IWQoS*, 2015.

[19] A. Kirsch and M. Mitzenmacher. Less hashing, same performance: Building a better Bloom filter. *Random Struct. Algorithms*, 33(2):187–218, 2006.

[20] J. Lu, Y. Wan, Y. Li, C. Zhang, Dai H, Y. Wang, G.Zhang, and B. Liu. Ultra-fast Bloom filters using SIMD techniques. *Proc. IEEE/ACM IWQoS*, 2017.

[21] K. Huang et al. A multi-partitioning approach to building fast and accurate counting Bloom filters. *IEEE 27th International Symposium on Parallel and Distributed Processing*, pages 1159–1170, 2013.

[22] M. Mitzenmacher, P. Reviriego, and S. Pontarelli. OMASS: One memory access set separation. *IEEE Transactions on Knowledge and Data Engineering*, 28(7):1940–1943, 2016.

[23] T. Wada, N. Matsumura, K. Nakano, and Y. Ito. Efficient byte stream pattern test using Bloom filter with rolling hash functions on the FPGA. *2018 Sixth International Symposium on Computing and Networking*, pages 66–75, 2018.

[24] I. Kaya and T. Kocak. A low power lookup technique for multi-hashing network applications. *IEEE Computer Society Annual Symposium on Emerging VLSI Technologies and Architectures*, 2006.

[25] M. J. Lyons and D. Brooks. The design of a Bloom filter hardware accelerator for ultra low power systems. *Proceedings of the 2009 ACM/IEEE International Symposium on Low Power Electronics and Design*, page 371–376, 2009.

[26] J. Harwayne-Gidansky, D. Stefan, and I. Dalal. FPGA-based SoC for real-time network intrusion detection using counting Bloom filters. *IEEE Southeastcon 2009, Atlanta, GA*, pages 452–458, 2009.

[27] Netcope. *www.netcope.com/getattachment/bb2b8efa-9925-438d-b895-897d7c1e4745/NFB-200G2QL-product-brief.aspx*.

[28] Z. Martinasek, J. Hajny, D. Smekal, L. Malina, D. Matousek, M. Kekely, and N. Mentens. 200 Gbps hardware accelerated encryption system for FPGA network cards. In *Proceedings of the 2018 Workshop on Attacks and Solutions in Hardware Security*, pages 11–17, 2018.

Authorized licensed use limited to: Universiteit Leiden. Downloaded on September 04,2024 at 14:31:23 UTC from IEEE Xplore. Restrictions apply.