

Fast Practical Compression of Deterministic Finite Automata^{*}

Philip Bille^{1[0000-0002-1120-5154]}, Inge Li Gørtz^{1[0000-0002-8322-4952]}, and Max Rishøj Pedersen^[0000-0002-8850-6422]

Technical University of Denmark {phbi, inge}@dtu.dk

Abstract. We revisit the popular *delayed deterministic finite automaton* (D²FA) compression algorithm introduced by Kumar et al. [SIGCOMM 2006] for compressing deterministic finite automata (DFAs) used in intrusion detection systems. This compression scheme exploits similarities in the outgoing sets of transitions among states to achieve strong compression while maintaining high throughput for matching.

Unfortunately, the D²FA algorithm and later variants of it require at least quadratic compression time since they compare all pairs of states to compute an optimal compression. This is too slow and, in some cases, even infeasible for collections of regular expression in modern intrusion detection systems that produce DFAs of millions of states.

Our main result is a simple, general framework for constructing D²FA based on locality-sensitive hashing that constructs an approximation of the optimal D²FA in near-linear time. We apply our approach to the original D²FA compression algorithm and two important variants, and we experimentally evaluate our algorithms on DFAs from widely used modern intrusion detection systems. Overall, our new algorithms compress up to an order of magnitude faster than existing solutions with either no or little loss of compression size. Consequently, our algorithms are significantly more scalable and can handle larger collections of regular expressions than previous solutions.

Keywords: Finite Automata · Algorithms · Data Compression.

1 Introduction

Signature-based deep packet inspection is a key component of modern intrusion detection and prevention systems. The basic idea is to maintain a collection of regular expressions, called *signatures*, that correspond to malicious content, violations of security policies, etc., and then match the collection of signatures against the input. In the typical scenario of high-throughput network traffic, the matching must be fast enough to process the input at the network's speed.

A natural approach to solve this is to construct a *deterministic finite automaton* (DFA) of the collection of signatures and then simulate it on the input.

^{*} Partially supported by the Independent Research Fund Denmark (DFF-9131-00069B and 10.46540/3105-00302B).

The DFA matches each character of the input with a single constant time state transition and requires only a single memory access. Unfortunately, DFAs for the collections of signatures used in modern intrusion detection systems are prohibitively large and not feasible for practical implementation [6, 51]. To overcome this space issue while still maintaining fast matching, significant work has been done on compressing DFAs [3–5, 7–9, 15, 20–22, 25, 29–32, 34, 35, 37, 42, 45, 48].

In this paper, we revisit the elegant and powerful *delayed DFA* compression technique introduced by Kumar et al. [29] and applied in many subsequent solutions [7, 9, 22, 30–35, 37]. The basic observation is that many states in the DFAs for real-world regular expression collections have similar sets of outgoing transitions, and we can take advantage of this to reduce the space significantly. Specifically, if two states s and s' share many such transitions, we can replace these in s with a special *default transition* to the other state s' .

Kumar et al. [29] proposed an algorithm to compute an optimal set of default transitions to compress any DFA. The key idea is to compute the similarity of all pairs of states, i.e., the number of shared outgoing transitions. We store these in a complete graph, called the *space reduction graph* (SRG), on the states with edges weighted by similarity. Finally, we compute a maximum spanning tree on the SRG and use each edge in the tree as a default transition leading to a compressed version of the original DFA called the *delayed deterministic finite automaton* (D²FA).

To implement matching, we traverse the D²FA similar to the Aho-Corasick algorithm for multi-string matching [2]. When we are at a state s and want to match a character α , we first inspect the outgoing transitions at s for a match of α . If we find a match, we continue to that state and process the next character in the input, and if not, we follow the default transition to state s' and repeat the process from s' with character α .

Compared to a standard DFA solution, Kumar et al. [29] showed that the D²FA dramatically reduces space by more than 90% on real-world collections of regular expressions while still achieving fast matching performance.

Processing a character D²FA during matching may require following multiple default transitions, thus incurring (as the name suggests) a *delay*. Minimizing the delay is essential in high-throughput applications, and two important variants of D²FAs that address this problem have been proposed. Kumar et al. [29] gave a modified D²FA construction that, given an integer parameter L , limits the maximum path of default transitions to L . We refer to this as the *longest delay* variant of the problem. Alternatively, Becchi and Crowley [7, 9] gave a modified D²FA construction that limits the maximum total number of default transitions traversed on any input string S by $|S|$. We refer to this as the *matching delay* variant.

The main bottleneck in the above algorithms is computing the similarity of all pairs of states to construct the SRG. If the input DFA contains n states, this requires at least $\Omega(n^2)$ time and space. This is too slow and, in some cases, even infeasible for collections of regular expression in modern intrusion detection systems that produce DFAs of millions of states.

Contributions We present a simple, general framework for fast compression of DFAs with default transitions based on locality-sensitive hashing. We apply our approach to general D²FA compression and the longest delay and matching delay variant and experimentally evaluate our algorithms on collections of regular expressions used in the popular Snort [43], Zeek [38], and Suricata [1] intrusion detection systems (see also www.snort.org, zeek.org, and suricata.io). Overall, we obtain new algorithms that compress up to an order of magnitude faster than existing solutions with either no or little loss of compression size. Consequently, our algorithms are significantly more scalable and can handle larger collections of regular expressions than previous solutions.

Technically, our main idea is to use locality-sensitive hashing to identify approximately similar states quickly. We then add edges between these states weighted by their similarity, producing the *sparse space reduction graph* (SSRG). We then compute the maximum spanning tree on the SSRG and then the D²FA. The SSRG contains significantly fewer edges, leading to a significant improvement in compression time. While the SSRG approximates the SRG by discarding edges, which may lead to worse compression size, we observe that this loss is negligible experimentally. For the longest delay variant, the previous solution by Kumar et al. [29] also uses a costly heuristic to maintain a bounded diameter maximum spanning forest. If we directly apply our sparsification technique, this heuristic dominates the running time, and we do not experimentally observe a significant speed-up in compression time. Instead, we develop an efficient alternative heuristic that first constructs a maximum spanning tree and then cuts edges to achieve the desired maximum spanning forest with bounded diameter. We show that combining our sparsification with the new heuristic leads to improvements in compression time similar to our other variants with little or no loss of compression size.

Related Work Substantial work has been done on compression DFAs. For an overview, see surveys [41, 50]. A popular approach is compressing the set of transitions [3–5, 7, 9, 20, 21, 29–32, 34, 35, 37, 42, 45, 48]. This approach includes the popular D²FA algorithm we focus on in this paper. Another approach is to compress the alphabet to reduce the size of the transition table [8, 9, 15, 25, 47]. The main idea is that, if some sets of characters (almost) always cause the same transitions throughout the DFA, they can be replaced by a single character [8, 9, 15, 25]. Alternatively, we can also compress the alphabet by replacing infrequent characters with sequences of frequent characters [47]. Finally, we can also compress the set of states as proposed by Becchi and Cambadi [5]. They showed how states with similar sets of outgoing transitions could be merged into one, thus compressing the set of states.

Using locality-sensitive hashing for fast compression of collections of sets has been used widely in many other contexts [11, 17, 18, 26, 28, 36, 39, 46, 49]. Our work naturally extends this work to fast DFA compression.

2 Preliminaries

Deterministic Finite Automata A *deterministic finite automaton* (DFA) is a 5-tuple $D = (Q, \Sigma, \delta, q_0, A)$ where Q is a set of states, Σ is an alphabet, $\delta : Q \times \Sigma \rightarrow Q$ is a transition function, $q_0 \in Q$ is the initial state and $A \subseteq Q$ is a set of accepting states. We let $n = |Q|$ denote the number of states. A DFA can be thought of as a *labeled directed graph* where Q is the set of nodes and each transition $\delta(u, c) = v$ is a labeled, directed edge, denoted $(u, v)_c$. See Figure 1. For simplicity, we assume every state has exactly one labeled transition for each character in the alphabet, i.e., δ is *total*, as in previous work.

Given a string S and a path p in D we say that p matches S if the concatenation of the labels of p equals S . A path that starts in q_0 and ends in A is *accepting* and D *accepts* a string S if there exists an accepting path that matches S . The *language* of D is the set of strings it accepts.

Locality-Sensitive Hashing A family of hash functions is *locality-sensitive*, for some similarity measure, if the probability of two objects hashing to the same value is *high* (lower-bounded for some parameter) when they are *similar* (similarity above some threshold) and, conversely, *low* when they are *dissimilar* (see e.g. [23] for formal details). There are different families of locality-sensitive hash functions for different distance or similarity measures, with some of the most popular being *simhash* [16], *MinHash* [14] and *sdhash* [44]. For example, the MinHash of a set is the minimum element according to a uniformly random permutation. The probability that two sets A and B hash to the same value is precisely their Jaccard similarity $(|A \cap B|)/(|A \cup B|)$.

3 Delayed Deterministic Finite Automata

A *delayed deterministic finite automaton* (D²FA) [29] is a deterministic finite automaton that is augmented with unlabeled *default transitions*. Formally a D²FA is a 6-tuple $D^2 = (Q, \Sigma, \delta, q_0, A, F)$. As for DFAs, Q is the set of states, Σ is the alphabet, δ is the transition functions, $q_0 \in Q$ is the initial state, and $A \subseteq Q$ is the set of accepting states. The final component $F : Q \rightarrow Q$ is the *default transition function*. Viewed as a graph, default transitions are ϵ -labeled directed edges, where ϵ is the empty string, and each state has at most one outgoing default transition. See Figure 1.

To transition from a state u according to a character c , we follow a c -labeled transition if it exists or otherwise follow the default transition:

$$\delta(u, c) = \begin{cases} v & \text{if } (u, v)_c \text{ is a edge.} \\ \delta(F(u), c) & \text{otherwise} \end{cases}$$

Note that for δ to be well-defined, reaching a state from u with a c -labeled transition must always be possible. This implies that any cycle of default transitions must have an outgoing c -labeled transition for any character c . To transition from

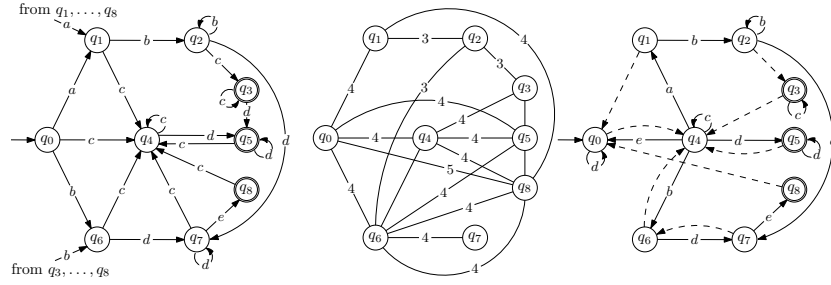


Fig. 1. Example from [29]. DFA D for regular expression $*(ab+c+)(cd+)(bd+e)$. Edges to q_0 are omitted (left). Space reduction graph for D with edges annotated with similarity. Edges with similarity less than 4 omitted, except those connecting q_2 to avoid disconnecting the graph (middle). D^2FA equivalent to D . All transitions are shown, default transitions are dashed (right).

a state u according a string $S = c_1 \dots c_m$ we recursively transition according to each character:

$$\delta(u, c_1 c_2 \dots c_m) = \delta(\delta(u, c_1), c_2 \dots c_m).$$

Given a character c and a path $p = u_1, \dots, u_k$ we say p matches c if $\delta(u_1, c) = u_k$ and all but the last transition is default, i.e., $F(u_i) = u_{i+1}$ for $1 \leq i < k$. Note that the concatenation of the labels of p equals c . Given a string $S = c_1 \dots c_m$ and a path p we say p matches S if p is the concatenation of paths matching the individual characters c_1, \dots, c_m . Note that the concatenation of the labels of p is S . We define acceptance as before. Two D^2FA s are equivalent if they have the same language, and two transitions are equivalent if they have the same destination and label.

Given a DFA, we can compress it by replacing sets of equivalent transitions with single default transitions to obtain an equivalent D^2FA with fewer total transitions. We define the *similarity* of two states u and v , denoted $sim(u, v)$, to be their number of equivalent transitions, that is, $sim(u, v) = |\{c \in \Sigma \mid \delta(u, c) = \delta(v, c)\}|$. See Figure 1. Inserting a default transition (u, v) and removing the equivalent transitions from u does not affect the language but saves $sim(u, v) - 1$ transitions. Each transition we can remove without affecting the language we say is *redundant*.

Following a default transition does not consume an input character, which introduces a *delay* when matching. We define the *longest delay* of D^2 to be maximum number of default transitions in any path matching a single character, i.e., the longest delay is the maximum number of default transitions in D^2 to match any single character. Given a string S , we define the *matching delay* of S in D^2 to be the number of default transitions in the path starting in q_0 and matching S .

4 Compressing DFAs with Default Transitions

We now review the algorithm of Kumar et al. [29] that compress a DFA D into an equivalent D^2 FA. Let D^2 be an initially empty D^2 FA with the same set of states as D . We proceed as follows.

- Step 1: Space Reduction Graph** Construct a complete, undirected graph on the states of D , and to each edge (u, v) assign weight $\text{sim}(u, v)$. This is the *space reduction graph* (SRG). See Figure 1.
- Step 2: Maximum Spanning Tree** Build a maximum spanning tree over the SRG. Root the spanning tree in a central node, i.e., a node of minimal radius, and direct all edges towards the root to obtain a directed spanning tree.
- Step 3: Transitions** For each edge (u, v) in the tree insert the default transition (u, v) into D^2 . Copy every labeled transition from D into D^2 that is not redundant in D^2 .

Since the SRG weighs the edges by similarity, computing the maximum spanning tree maximizes the overall compression.

Step 1 uses $O(n^2|\Sigma|)$ to compare all states and construct the SRG. Step 2 constructs the maximum spanning tree with Kruskal’s algorithm [27] in $O(n^2 \log n)$ time, and step 3 takes $O(n|\Sigma|)$ time. In total, the running time is $O(n^2 \log n + n^2|\Sigma|)$ ¹.

5 Fast Compression

We now show how to speed up the DFA compression algorithm using locality-sensitive hashing to sparsify the SRG construction in step 1.

Let D be the input DFA and let r and k be two constant, positive integer parameters. We initialize an undirected graph $G = (Q, E)$ where Q is the set of states in D and $E = \{(q_0, u) \mid u \neq q_0\}$, i.e., we have an edge between the initial state and all other states. We then add edges to the graph in r rounds, where each round proceeds as follows:

First, pick k unique random characters $c_1, \dots, c_k \in \Sigma$. Then, for each state $v \in Q$ we construct the sequence of k states $V = \delta(v, c_1), \dots, \delta(v, c_k)$. We hash V into a single hash value $h(v)$ using a standard hashing scheme of Black et al. [12]. We insert v into a table with key $h(v)$. For each unique hash value h_i , consider the set of states C_i that hash to h_i . For each state $u \in C_i$, we pick another state $v \in C_i$ uniformly at random and insert (u, v) into E if it does not already exist.

After r rounds, the algorithm terminates, and we assign weights to each edge of G equal to the similarity of the endpoint states. The resulting graph G is the *sparse space reduction graph* (SSRG).

The above scheme is inspired by Har-Peled et al. [23] hashing bitstrings w.r.t. Hamming distance. Their solution samples positions from the input bitstrings.

¹ The running time is not explicitly stated in the paper, but follows from the description.

In our solution, the sampled positions correspond to the sampled characters from Σ .

Hashing a state takes $O(k)$ time, and sampling an edge takes constant time. Thus, a round takes $O(kn)$ time and we use $O(rkn)$ time for all rounds. Each round inserts at most n edges and hence the SSRG G has $O(rn)$ edges at the end. We compute the similarity between two states in $O(|\Sigma|)$ time, and hence the final algorithm uses $O(rkn + rn|\Sigma|) = O(n|\Sigma|)$ time.

We plug in the modified Step 1 in the algorithm from Section 4. Step 1 takes $O(n|\Sigma|)$ time and since G has $O(n)$ edges, Step 2 takes $O(n \log n)$ time. Step 3 takes $O(n|\Sigma|)$ time as before. In total, we use $O(n \log n + n|\Sigma|)$ time.

6 Compression with Bounded Longest Delay

We now consider the bounded longest delay variant of D^2FA , that is, given a DFA D and an integer parameter L , construct a D^2FA D^2 equivalent to D , such that the longest delay of D^2 is at most L . Recall that the longest delay of a D^2FA is the length of the longest path of default transitions in D^2 .

Bounded Longest Delay by Constructing Small Trees We first review the algorithm by Kumar et al. [29]. The algorithm is based on a simple modification to the maximum spanning tree construction in Step 2 of the algorithm from Section 4.

Let L be a parameter. We modify Step 2 by constructing a maximum spanning forest with the constraint that each tree in the forest has a diameter of at most $\Delta = 2L$. To do so, we run Kruskal's algorithm but simply ignore any edges that would cause a tree diameter to exceed Δ . Also, among the edges with maximum similarity, we select one that causes a minimum increase to any tree diameter. After constructing the forest, we root each tree in a central node and direct edges toward each root.

Since each tree has diameter at most Δ and is rooted in a central node, the final D^2FA has a longest delay of at most $\lceil \Delta/2 \rceil = L$. To implement the modified Step 2, Kumar et al. [29] maintains the radius of the tree for each node during the maximum spanning forest construction. When we add an edge, we need to merge two trees and potentially update the radius for each node in the resulting tree. Hence, we may need to update $\Omega(n^2)$ radii in total during the maximum spanning forest construction.

We note that new Step 2 uses $\Omega(n^2)$ time, whether or not we run it on a sparse or a dense space reduction graph. Hence, we cannot directly apply our sparsification technique from Section 5 to obtain an efficient algorithm. We present a new algorithm in the next section that efficiently combines with sparsification.

Fast Compression with Bounded Longest Delay We now present a new algorithm that uses sparsification to efficiently construct D^2FAs with bounded longest delay. The idea is to construct a single large maximum spanning tree and then cut edges until each tree in the resulting forest has small diameter.

Given a DFA D and an integer parameter L , we proceed as follows.

- Step 1: Construct SSRG** Construct a sparse space reduction graph G for D as in Section 5.
- Step 2: Construct MST** Construct a maximum spanning tree T_0 over G using Kruskal’s algorithm. Pick a central node v_0 in T_0 and then discard T_0 . Construct a new maximum spanning tree T using Prim’s algorithm [40] with v_0 as the initial node. When we queue a new edge (u, v) , where u is the node already in T , assign weight $w'_{u,v} = \text{sim}(u, v) - 2^{d_v}$ where d_v is the distance from v_0 to v in T .
- Step 3: Cut Edges** Cut a minimum number of edges in T to obtain a forest with each tree diameter at most $\Delta = 2L$. We do so using the algorithm of Farley et al. [19] that cuts the necessary edges in a bottom-up traversal of T . Then, direct the edges of each tree towards the root.
- Step 4: Construct Transitions** As in Step 3 in Section 4, create default transitions along edges in the trees and then copy in every labeled transition from D that is not redundant in D^2 .

After cutting, each tree has a diameter of at most $\Delta = 2L$. Since we direct edges toward each root, the longest delay is at most $\lceil \Delta/2 \rceil = L$. Step 1 and 4 uses $O(n|\Sigma|)$ time as before. Step 2 uses $O(n \log n)$ time, and Step 3 uses $O(n)$ time. In total, we use $O(n \log n + n|\Sigma|)$ time.

Each edge (u, v) we cut results in $\text{sim}(u, v) - 1$ more labeled transitions in D^2 , as that default transition is then not constructed. Intuitively, the lower the diameter of T , the fewer edges we cut to get each tree below the bound. Therefore, we use an edge weight that trades similarity for lower diameter, as we observed that the fewer cuts outweighed the lost similarity.

We found the simple heuristic of the modified weight w' performed well in practice. A similar idea was used by Kumar et al [29] in their solution.

Because T is not a maximum spanning tree w.r.t. similarity, the choice of initial node v_0 affects the total similarity of the final tree. We found that picking v_0 to be a central node in a MST w.r.t. similarity (T_0) yielded the best compression in practice.

Note that we cut the minimum *number* of edges to uphold the diameter constraint. Alternatively, we could cut edges of minimum *total similarity*, which might result in better compression. However, our approach is simple and fast in practice, and because each edge in the SRG has near-maximum weight, the difference in compression is negligible.

7 Compression with Bounded Matching Delay

We now consider the bounded matching delay variant of D^2 FA, that is, given a DFA D , construct a D^2 FA D^2 equivalent to D , such that the matching delay of D^2 is at most $|S|$ on any input string S .

Bounded Matching Delay by the A-DFA Algorithm We first review A-DFA algorithm by Becchi and Crowley [7, 9]. Let D be an input DFA. For a state $v \in Q$, define the *depth* of v , denoted $d(v)$, to be the length of the shortest path from

the initial state q_0 to v . The key idea is only to add default transitions from state v to state u if $d(u) < d(v)$. This implies that the matching delay is at most $|S|$ on any input string S (see, e.g., Aho and Corasick [2]).

Initialize a D²FA D^2 with no default transitions. We proceed as follows.

Step 1: Calculate Depth Calculate the depth $d(v)$ of each state $v \in Q$ by a breadth-first traversal of D .

Step 2: Construct Default Transitions For each state $u \in Q$ add default transition (u, v) to D^2 , where v is the state such that $\text{sim}(u, v)$ is maximum and $d(v) < d(u)$.

Step 3: Construct Labeled Transitions Copy every labeled transition from D that is not redundant in D^2 .

Step 1 and 3 uses $O(n|\Sigma|)$ time to traverse D . Step 2 uses $O(n^2|\Sigma|)$ to compute the similarity of each pair of states. In total, we use $O(n^2|\Sigma|)$ time.

Fast Compression with Bounded Matching Delay We now speed up the A-DFA algorithm using sparsification. Let D be the input DFA and let r and k be two constant, positive integer parameters. We initialize a D²FA D^2 with no default transitions, i.e., we set $F(u) = u$ for each $u \in Q$. The algorithm runs in r rounds, where each round proceeds as follows.

First, pick k unique random characters $c_1, \dots, c_k \in \Sigma$. Then, for each state $v \in Q$ we construct the sequence $V = \delta(v, c_1), \dots, \delta(v, c_k)$, and hash V into a single hash value $h(v)$. We insert v into a table with key $h(v)$. For each unique hash value h_i , we consider the set of states C_i that hash to that value. For each state $u \in C_i$ we pick another state $v \in C_i$ uniformly at random. If v has lower depth and the default transition (u, v) compresses better than the current default transition of u , i.e., $d(v) < d(u)$ and $\text{sim}(u, v) > \text{sim}(u, F(u))$, we update the default transition of u to point to v in D^2 , otherwise, we continue. After r rounds, the algorithm terminates and returns the resulting D²FA D^2 .

Hashing a state and computing the similarity of the potential new default transition uses $O(k|\Sigma|)$ time. Hence, the full algorithm uses $O(rkn + rn|\Sigma|) = O(n|\Sigma|)$ time.

8 Experimental Evaluation

We implemented our methods described in the previous section and measured their performance on regular expressions extracted from widely used intrusion detection systems. The implementation is available at <https://github.com/MaxRishoj/fcomp-dfa>.

Datasets We extracted our datasets from regular expressions used in the popular Snort [43], Zeek (formerly Bro) [38], and Suricata [1] intrusion detection systems (see current homepages for these systems at www.snort.org, zeek.org, and suricata.io). The Snort and Zeek datasets are extracted from current versions of the datasets used in most of the previous work.

For each dataset, we extracted prefixes of the rules to generate DFAs of different sizes to explore the scalability of our algorithms on DFAs that have between roughly 1k to 1M states. The details are in the full version of the paper [10]. As in previous work, we have filtered out some rules that used advanced features. All regular expressions use an ASCII alphabet size of 256.

Algorithms Tested We evaluate the following algorithms.

- D²FA. The algorithm of [29], described in Section 4.
- D²FA-LD. The algorithm of [29] for bounded longest delay, described in Section 6 with parameter $L = 2$.
- D²FA-LD-CUT. The algorithm in Section 6 for bounded longest delay using the SRG instead of the SSRG.
- D²FA-MD. The algorithm of [7, 9] for bounded matching delay, described in Section 7.
- SPARSE-D²FA. The algorithm in Section 5.
- SPARSE-D²FA-LD. As D²FA-LD using the SSRG instead of the SRG.
- SPARSE-D²FA-LD-CUT. The algorithm in Section 6 for bounded longest delay.
- SPARSE-D²FA-MD. The algorithm in Section 7 for bounded matching delay.

We use the locality-sensitive hashing scheme from Section 5 without replacement and parameters $k = 8$ and $r = 512$. We evaluated several locality-sensitive hashing schemes, including the one in Section 5 with replacement and min-hash [13] over the set of outgoing transitions with one or k random permutations of the universe. We also evaluated several combinations of r and k and found that $k = 8$ gave the best compression size. Increasing r results in better compression size but increases the compression time linearly. Our chosen variant achieved the best combination of compression size and compression time.

In our longest delay variant, we report results for parameter $L = 2$. We have experimented with other values of L but did not observe significant differences in the relative performances of the algorithm for this variant. We note that the bounded matching delay variant also appears in a more general version, where we can tune the overhead of the default transitions (see the journal version of the result [9]). The version tested here is the simplest version and leads to the best compression size.

Setup Experiments were run on a machine with an Intel Xeon Gold 6226R 2.9GHz processor and 128GB of memory. The operating system was Scientific Linux 7.9 kernel version 3.10.0-1160.80.1.el7.x86_64. Source code was compiled with g++ version 9.4 with options `-Wall -O4`. The input to each algorithm is a DFA constructed from a set of regular expressions. We measured the time for constructing an equivalent D²FA for the input DFA, using the `clock` function of the C standard library.

Results We compare the algorithms across the datasets and measure compression time and compression size (number of states in D²FA as a percent of the

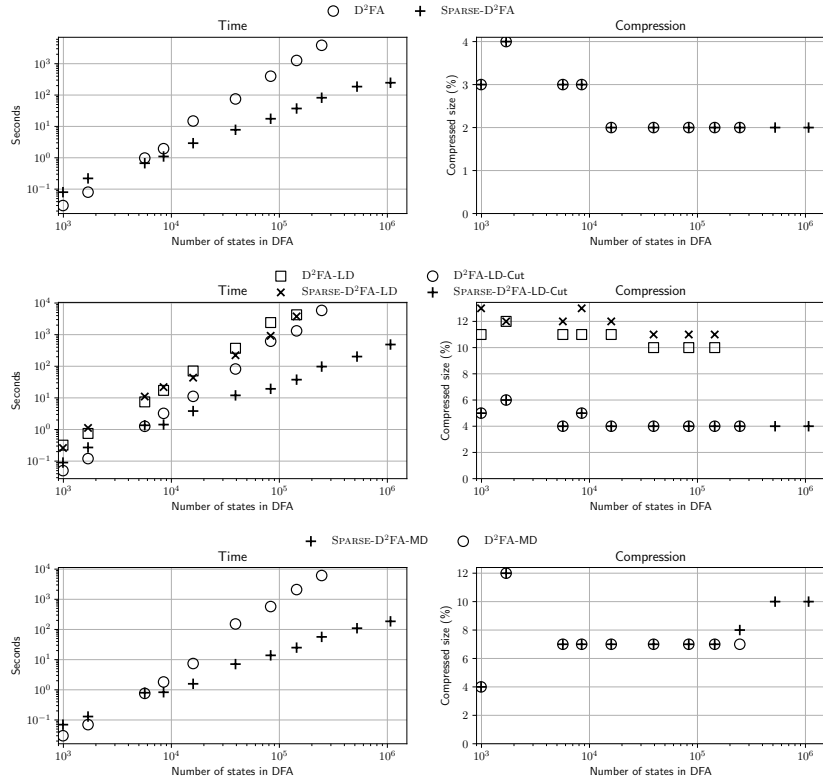


Fig. 2. Results for the Snort dataset on the algorithms for general compression (top), bounded longest delay (middle), and bounded longest matching delay (bottom). On the left, we show compression time in seconds vs. the number of states in the input DFA. On the right, we show the number of transitions in the D²FA as a percent of the number of transitions in the input DFA.

number of states in the input DFA) for each of the variants (general compression, longest delay, and bounded matching delay). The relative performance of our algorithms is similar across the datasets, so we focus on the results for the Snort dataset shown in Figure 2. The corresponding results for the Suricata and Zeek datasets are in the full version of the paper [10]. We point out whenever there are significant differences between observed results across datasets. We note that the absolute compression size varies significantly across the datasets and variants. Most instances are highly compressible to around 10% of the original DFA (many even in the low single-digit percentages). At the same time, a single one (bounded matching delay on the Zeek dataset) compresses to around 50 percent on the largest DFAs.

General Compression We compare the D²FA and SPARSE-D²FA general compression algorithms. We observe that SPARSE-D²FA compresses up to an order of magnitude faster than D²FA with either no or little loss of compression size. For DFAs with around 1k states, the compression time is comparable and increases gradually to roughly an order of magnitude for the largest DFAs.

Compression with Bounded Longest Delay We compare the D²FA-LD, SPARSE-D²FA-LD, D²FA-LD-CUT and SPARSE-D²FA-LD-CUT compression algorithms with bounded longest delay. We observe that D²FA-LD and SPARSE-D²FA-LD achieve similar compression times as expected. The compression size achieved by SPARSE-D²FA-LD is around 10-15% worse in the Snort dataset, roughly comparable in the Zeek dataset, and around 100% worse in the Suricata dataset. We believe that the bounded diameter approach in Kumar et al. [29] is highly sensitive to the greedy choice of edges at each step, leading to the observed differences in the compression size across the different datasets.

More importantly, we observe that SPARSE-D²FA-LD-CUT compresses up to an order of magnitude faster than all other algorithms. SPARSE-D²FA-LD-CUT also achieves a substantially better compression size than D²FA-LD and SPARSE-D²FA-LD (except for the Zeek dataset, where the compression size is comparable). Compared to the dense version D²FA-LD-CUT, SPARSE-D²FA-LD-CUT has either no or little loss of compression size.

Compression with Bounded Matching Delay We compare the compression algorithms D²FA-MD and SPARSE-D²FA-MD with bounded matching delay. We observe that SPARSE-D²FA-MD compresses up to an order of magnitude faster than D²FA-MD. The compression size varies depending on the dataset. For the Snort dataset, SPARSE-D²FA-MD achieves 10-15% worse compression size for DFAs with more than 10k states, for the Zeek dataset, SPARSE-D²FA-MD is comparable, and for the Suricata dataset, SPARSE-D²FA-MD achieves 100% worse compression size for DFAs with more than 10k states. We note that the absolute compression size varies significantly, which may explain this difference in relative compression size.

9 Conclusion and Acknowledgements

We presented a simple, practical framework for constructing D²FA based on locality-sensitive hashing. On DFAs from widely used modern intrusion detection systems, we achieved compression times of up to an order of magnitude faster than existing solutions with either no or little loss of compression size. An interesting open problem is to explore if our new framework can be combined with other DFA compression techniques, such as the ones mentioned in Section 1.

This paper is inspired by earlier work in an MSc. thesis [24] supervised by two of the authors. We thank the anonymous reviewers of earlier drafts of this article for many valuable comments that improved the quality of the work.

References

1. www.suricata.io
2. Aho, A.V., Corasick, M.J.: Efficient string matching: an aid to bibliographic search. *Commun. ACM* **18**(6), 333–340 (1975). <https://doi.org/http://doi.acm.org/10.1145/360825.360855>
3. Antonello, R., Fernandes, S.F.L., Sadok, D., Kelner, J., Szabó, G.: Deterministic finite automaton for scalable traffic identification: The power of compressing by range. In: *NOMS 2012*. pp. 155–162 (2012). <https://doi.org/10.1109/NOMS.2012.6211894>
4. Antonello, R., Fernandes, S.F.L., Sadok, D.F.H., Kelner, J., Szabó, G.: Design and optimizations for efficient regular expression matching in DPI systems. *Comput. Commun.* **61**, 103–120 (2015). <https://doi.org/10.1016/j.comcom.2014.12.011>
5. Becchi, M., Cadambi, S.: Memory-efficient regular expression search using state merging. In: *Proc. 26th INFOCOM*. pp. 1064–1072 (2007). <https://doi.org/10.1109/INFOCOM.2007.128>
6. Becchi, M., Crowley, P.: A hybrid finite automaton for practical deep packet inspection. In: *Proc. 3rd CoNEXT conference*. pp. 1–12 (2007)
7. Becchi, M., Crowley, P.: An improved algorithm to accelerate regular expression evaluation. In: *Proc. ANCS 2007*. pp. 145–154 (2007). <https://doi.org/10.1145/1323548.1323573>
8. Becchi, M., Crowley, P.: Efficient regular expression evaluation: theory to practice. In: *Proc. ANCS 2008*. pp. 50–59 (2008). <https://doi.org/10.1145/1477942.1477950>
9. Becchi, M., Crowley, P.: A-DFA: A time- and space-efficient DFA compression algorithm for fast regular expression evaluation. *ACM Trans. Archit. Code Optim.* **10**(1), 4:1–4:26 (2013). <https://doi.org/10.1145/2445572.2445576>
10. Bille, P., Gørtz, I.L., Pedersen, M.R.: Fast practical compression of deterministic finite automata. *arXiv:2306.12771* (2024)
11. Bille, P., Gørtz, I.L., Puglisi, S.J., Tarnow, S.R.: Hierarchical relative lempel-ziv compression. In: *Proc. 21st SEA* (2023)
12. Black, J., Halevi, S., Krawczyk, H., Krovetz, T., Rogaway, P.: UMAC: fast and secure message authentication. In: *Proc. 19th CRYPTO*. vol. 1666, pp. 216–233 (1999). https://doi.org/10.1007/3-540-48405-1_14
13. Broder, A.Z.: On the resemblance and containment of documents. In: *Proc. SEQUENCES*. pp. 21–29 (1997)
14. Broder, A.Z., Charikar, M., Frieze, A.M., Mitzenmacher, M.: Min-wise independent permutations. *J. Comput. Syst. Sci.* **60**(3), 630–659 (2000). <https://doi.org/10.1006/jcss.1999.1690>
15. Brodie, B.C., Taylor, D.E., Cytron, R.K.: A scalable architecture for high-throughput regular-expression pattern matching. In: *Proc. 33rd ISCA*. pp. 191–202 (2006). <https://doi.org/10.1109/ISCA.2006.7>
16. Charikar, M.: Similarity estimation techniques from rounding algorithms. In: *Proc. 34th STOC*. pp. 380–388 (2002). <https://doi.org/10.1145/509907.509965>
17. Ding, S., Attenberg, J., Suel, T.: Scalable techniques for document identifier assignment in inverted indexes. In: *Proc. 19th WWW*. pp. 311–320 (2010)
18. Douglis, F., Iyengar, A.: Application-specific delta-encoding via resemblance detection. In: *Proc. USENIX ATC, General Track 2003*. pp. 113–126 (2003), <http://www.usenix.org/events/usenix03/tech/douglis.html>

19. Farley, A.M., Hedetniemi, S.T., Proskurowski, A.: Partitioning trees: Matching, domination, and maximum diameter. *Int. J. Parallel Program.* **10**(1), 55–61 (1981). <https://doi.org/10.1007/BF00978378>
20. Ficara, D., Giordano, S., Procissi, G., Vitucci, F., Antichi, G., Pietro, A.D.: An improved DFA for fast regular expression matching. *Comput. Commun. Rev.* **38**(5), 29–40 (2008). <https://doi.org/10.1145/1452335.1452339>
21. Ficara, D., Pietro, A.D., Giordano, S., Procissi, G., Vitucci, F., Antichi, G.: Differential encoding of dfas for fast regular expression matching. *IEEE/ACM Trans. Netw.* **19**(3), 683–694 (2011). <https://doi.org/10.1109/TNET.2010.2089639>
22. Gong, L., Wang, C., Xia, H., Chen, X., Li, X., Zhou, X.: Enabling fast and memory-efficient acceleration for pattern matching workloads: The lightweight automata processing engine. *IEEE Trans. Computers* **72**(4), 1011–1025 (2023). <https://doi.org/10.1109/TC.2022.3187338>
23. Har-Peled, S., Indyk, P., Motwani, R.: Approximate nearest neighbor: Towards removing the curse of dimensionality. *Theory Comput.* **8**(1), 321–350 (2012). <https://doi.org/10.4086/toc.2012.v008a014>
24. Hemmingsen, M., Lam, B.W.: Fast Compression of DFAs for Intrusion Detection Systems. Master’s thesis, Tech. Uni. Denmark. (2021)
25. Kong, S., Smith, R., Estan, C.: Efficient signature matching with multiple alphabet compression tables. In: *Proc. 4th SECURECOMM*. p. 1 (2008). <https://doi.org/10.1145/1460877.1460879>
26. Krcál, L., Holub, J.: Incremental locality and clustering-based compression. In: *DCC 2015*. pp. 203–212 (2015). <https://doi.org/10.1109/DCC.2015.23>
27. Kruskal, J.B.: On the shortest spanning subtree of a graph and the traveling salesman problem. *Proc. Amer. Math. Soc.* **7**(1), 48–50 (1956)
28. Kulkarni, P., Douglis, F., LaVoie, J.D., Tracey, J.M.: Redundancy elimination within large collections of files. In: *Proc. USENIX ATC, General Track 2004*. pp. 59–72 (2004)
29. Kumar, S., Dharmapurikar, S., Yu, F., Crowley, P., Turner, J.S.: Algorithms to accelerate multiple regular expressions matching for deep packet inspection. In: *Proc. SIGCOMM 2006*. pp. 339–350 (2006). <https://doi.org/10.1145/1159913.1159952>
30. Kumar, S., Turner, J.S., Williams, J.: Advanced algorithms for fast and scalable deep packet inspection. In: *Proc. ANCS 2006*. pp. 81–92 (2006). <https://doi.org/10.1145/1185347.1185359>
31. Liu, A.X., Torng, E.: An overlay automata approach to regular expression matching. In: *Proc. 33rd INFOCOM*. pp. 952–960 (2014). <https://doi.org/10.1109/INFOCOM.2014.6848024>
32. Liu, S., Su, S., Liu, D., Huang, Z., Xiao, M.: Efficient compression algorithm for ternary content addressable memory-based regular expression matching. *Electronics Letters* **53**(3), 152–154 (2017)
33. Matousek, D., Kubis, J., Matousek, J., Korenek, J.: Regular expression matching with pipelined delayed input dfas for high-speed networks. In: *Proc. ANCS 2018*. pp. 104–110 (2018). <https://doi.org/10.1145/3230718.3230730>
34. Matousek, D., Matousek, J., Korenek, J.: High-speed regular expression matching with pipelined memory-based automata. In: *Proc. 26th FCCM*. p. 214 (2018). <https://doi.org/10.1109/FCCM.2018.00048>
35. Meiners, C.R., Patel, J., Norige, E., Torng, E., Liu, A.X.: Fast regular expression matching using small tcams for network intrusion detection and prevention systems. In: *19th USENIX Security*. pp. 111–126 (2010)

36. Ouyang, Z., Memon, N.D., Suel, T., Trendafilov, D.: Cluster-based delta compression of a collection of files. In: Proc. 3rd WISE. pp. 257–268 (2002). <https://doi.org/10.1109/WISE.2002.1181662>
37. Patel, J., Liu, A.X., Torng, E.: Bypassing space explosion in high-speed regular expression matching. *IEEE/ACM Trans. Netw.* **22**(6), 1701–1714 (2014). <https://doi.org/10.1109/TNET.2014.2309014>
38. Paxson, V.: Bro: a system for detecting network intruders in real-time. *Comput. Networks* **31**(23-24), 2435–2463 (1999)
39. Peel, A., Wirth, A., Zobel, J.: Collection-based compression using discovered long matching strings. In: Proc. 20th CIKM. pp. 2361–2364 (2011). <https://doi.org/10.1145/2063576.2063967>
40. Prim, R.C.: Shortest connection networks and some generalizations. *The Bell System Technical Journal* **36**(6), 1389–1401 (1957)
41. Prithi, S., Sumathi, S.: A survey on recent dfa compression techniques for deep packet inspection in network intrusion detection system. *Journal of Electrical Engineering* **17**(3), 14–14 (2017)
42. Qi, Y., Wang, K., Fong, J., Xue, Y., Li, J., Jiang, W., Prasanna, V.K.: FEACAN: front-end acceleration for content-aware network processing. In: Proc. 30th INFOCOM. pp. 2114–2122 (2011). <https://doi.org/10.1109/INFOCOM.2011.5935021>
43. Roesch, M.: Snort: Lightweight intrusion detection for networks. In: Proc. 13th LISA. pp. 229–238 (1999)
44. Roussev, V.: Data fingerprinting with similarity digests. In: IFIP Int. Conf. Digital Forensics 2010. vol. 337, pp. 207–226 (2010). https://doi.org/10.1007/978-3-642-15506-2_15
45. Shankar, S.S., Lin, P., Herkersdorf, A., Wild, T.: A divide and conquer state grouping method for bitmap based transition compression. In: Proc. 18th PDCAT. pp. 400–406 (2017). <https://doi.org/10.1109/PDCAT.2017.00071>
46. Shilane, P., Huang, M., Wallace, G., Hsu, W.: Wan-optimized replication of backup datasets using stream-informed delta compression. *ACM Trans. Storage* **8**(4), 13:1–13:26 (2012). <https://doi.org/10.1145/2385603.2385606>
47. Tang, Q., Jiang, L., Dai, Q., Su, M., Xie, H., Fang, B.: RICS-DFA: a space and time-efficient signature matching algorithm with reduced input character set. *Concurr. Comput. Pract. Exp.* **29**(20) (2017). <https://doi.org/10.1002/cpe.3940>
48. Tuck, N., Sherwood, T., Calder, B., Varghese, G.: Deterministic memory-efficient string matching algorithms for intrusion detection. In: Proc. 23rd INFOCOM. pp. 2628–2639 (2004). <https://doi.org/10.1109/INFOCOM.2004.1354682>
49. Xia, W., Jiang, H., Feng, D., Hua, Y.: Silo: A similarity-locality based near-exact deduplication scheme with low RAM overhead and high throughput. In: USENIX ATC 2011 (2011)
50. Xu, C., Chen, S., Su, J., Yiu, S., Hui, L.C.K.: A survey on regular expression matching for deep packet inspection: Applications, algorithms, and hardware platforms. *IEEE Commun. Surv. Tutorials* **18**(4), 2991–3029 (2016). <https://doi.org/10.1109/COMST.2016.2566669>
51. Yu, F., Chen, Z., Diao, Y., Lakshman, T.V., Katz, R.H.: Fast and memory-efficient regular expression matching for deep packet inspection. In: Proc. ANCS 2006. pp. 93–102 (2006). <https://doi.org/10.1145/1185347.1185360>