

Approximate Reduction of Finite Automata for High-Speed Network Intrusion Detection

Milan Češka, Vojtěch Havlena, Lukáš Holík, Ondřej Lengál, and Tomáš Vojnar

the date of receipt and acceptance should be inserted later

Abstract We consider the problem of *approximate reduction of non-deterministic automata* that appear in hardware-accelerated network intrusion detection systems (NIDSes). We define an error *distance* of a reduced automaton from the original one as the probability of packets being incorrectly classified by the reduced automaton (wrt the probabilistic distribution of packets in the network traffic). We use this notion to design an *approximate reduction procedure* that achieves a great size reduction (much beyond the state-of-the-art language preserving techniques) with a controlled and small error. We have implemented our approach and evaluated it on use cases from SNORT, a popular NIDS. Our results provide experimental evidence that the method can be highly efficient in practice, allowing NIDSes to follow the rapid growth in the speed of networks.

1 Introduction

The recent years have seen a boom in the number of security incidents in computer networks. In order to alleviate the impact of network attacks and intrusions, Internet service providers want to detect malicious traffic at their network's entry points and on the backbones between sub-networks. Software-based network intrusion detection systems (NIDSes), such as the popular open-source system SNORT [48], are capable of detecting suspicious network traffic by testing (among others) whether a packet payload matches a regular expression (regex) describing known patterns of malicious traffic. NIDSes collect and maintain vast databases of such regexes that are typically divided into groups according to types of the attacks and target protocols.

Regex matching is the most computationally demanding task of a NIDS as its cost grows with the speed of

the network traffic as well as with the number and complexity of the regexes being matched. The current software-based NIDSes cannot perform the regex matching on networks beyond 1 Gbps [5,26], so they cannot handle the current speed of backbone networks ranging between tens and hundreds of Gbps. A promising approach to speed up NIDSes is to (partially) offload regex matching into hardware [26,25,34]. The hardware then serves as a pre-filter of the network traffic, discarding the majority of the packets from further processing. Such pre-filtering can easily reduce the traffic the NIDS needs to handle by two or three orders of magnitude [26].

Field-programmable gate arrays (FPGAs) are the leading technology in high-throughput regex matching. Due to their inherent parallelism, FPGAs provide an efficient way of implementing *nondeterministic finite automata* (NFAs), which naturally arise from the input regexes. Although the amount of available resources in FPGAs is continually increasing, the speed of networks grows even faster. Working with multi-gigabit networks requires the hardware to use many parallel packet processing branches in a single FPGA [34]; each of them implementing a separate copy of the concerned NFA, and so reducing the size of the NFAs is of the utmost importance. Various language-preserving automata reduction approaches exist, mainly based on computing (bi)simulation relations on automata states (cf. the related work). The reductions they offer, however, do not satisfy the needs of high-speed hardware-accelerated NIDSes.

Our answer to the problem is *approximate reduction* of NFAs, allowing for a trade-off between the achieved reduction and the precision of the regex matching. To formalise the intuitive notion of precision, we propose a novel *probabilistic distance* of automata. It captures the probability that a packet of the input network traffic is incorrectly accepted or rejected by the approximated NFA.

The distance assumes a *probabilistic model* of the network traffic (we show later how such a model can be obtained).

Having formalised the notion of precision, we specify the target of our reductions as two variants of an optimization problem: (1) minimizing the NFA size given the maximum allowed error (distance from the original), or (2) minimizing the error given the maximum allowed NFA size. Finding such optimal approximations is, however, computationally hard (**PSPACE**-complete, the same as precise NFA minimization).

Consequently, we sacrifice the optimality and, motivated by the typical structure of NFAs that emerge from a set of regexes used by NIDSes (a union of many long “tentacles” with occasional small strongly-connected components), we limit the space of possible reductions by restricting the set of operations they can apply to the original automaton. Namely, we consider two reduction operations: (i) collapsing the future of a state into a *self-loop* (this reduction over-approximates the language), or (ii) *removing states* (such a reduction is under-approximating).

The problem of identifying the optimal sets of states on which these operations should be applied is still **PSPACE**-complete. The restricted problem is, however, more amenable to an approximation by a *greedy algorithm*. The algorithm applies the reductions state-by-state in an order determined by a precomputed *error labelling* of the states. The process is stopped once the given optimization goal in terms of the size or error is reached. The labelling is based on the probability of packets that may be accepted through a given state and hence over-approximates the error that may be caused by applying the reduction at a given state. As our experiments show, this approach can give us high-quality reductions while ensuring formal error bounds.

Finally, it turns out that even the pre-computation of the error labelling of the states is costly (again **PSPACE**-complete). Therefore, we propose several ways to cheaply over-approximate it such that the strong error bound guarantees are still preserved. In particular, we are able to exploit the typical structure of the “union of tentacles” of the hardware NFA in an algorithm that is exponential in the size of the largest “tentacle” only, which gives us a method that is indeed much faster in practice.

We have implemented our approach and evaluated it on regexes used to classify malicious traffic in SNORT. We obtain quite encouraging experimental results demonstrating that our approach provides a much better reduction than language-preserving techniques with an almost negligible error. In particular, our experiments, going down to the level of an actual implementation of NFAs in FPGAs, confirm that we can squeeze into an up-to-date FPGA chip real-life regexes encoding malicious traffic, allowing them to be used with a negligible error for filtering at speeds of 100 Gbps (and even 400 Gbps). This is far

beyond what one can achieve with current exact reduction approaches.

This paper is an extended version of the paper that appeared in the proceedings of TACAS’18 [51], containing complete proofs of the presented lemmas and theorems.

Related Work Hardware acceleration for regex matching at the line rate is an intensively studied technology that uses general-purpose hardware [28,47,29,3,4,27,53,30,31] as well as FPGAs [37,8,12,23,45,43,26,25,34]. Most of the works focus on DFA implementation and optimization techniques. NFAs can be exponentially smaller than DFAs but need, in the worst case, $\mathcal{O}(n)$ memory accesses to process each byte of the payload where n is the number of states. In most cases, this incurs an unacceptable slowdown. Several works alleviate this disadvantage of NFAs by exploiting reconfigurability and fine-grained parallelism of FPGAs, allowing one to process one character per clock cycle (e.g. [37,8,45,43,26,25,34]).

In [31], which is probably the closest work to ours, the authors consider a set of regexes describing network attacks. They replace a potentially prohibitively large DFA by a tree of smaller DFAs, an alternative to using NFAs that minimizes the latency occurring in a non-FPGA-based implementation. The language of every DFA-node in the tree over-approximates the languages of its children. Packets are filtered through the tree from the root downwards until they belong to the language of the encountered nodes, and may be finally accepted at the leaves, or are rejected otherwise. The over-approximating DFAs are constructed using a similar notion of probability of an occurrence of a state as in our approach. The main differences from our work are that (1) the approach targets approximation of DFAs (not NFAs), (2) the over-approximation is based on a given traffic sample only (it cannot benefit from a probabilistic model), and (3) no probabilistic guarantees on the approximation error are provided.

Approximation of DFAs was considered in various other contexts. Hyper-minimization is an approach that is allowed to alter language membership of a finite set of words [33, 19]. A DFA with a given maximum number of states is constructed in [18], minimizing the error defined either by (i) counting prefixes of misjudged words up to some length, or (ii) the sum of the probabilities of the misjudged words wrt the Poisson distribution over Σ^* . Neither of these approaches considers reduction of NFAs nor allows to control the expected error with respect to the real traffic.

In addition to the metrics mentioned above when discussing the works [18,33,19], the following metrics should also be mentioned. The Cesaro-Jaccard distance studied in [42] is, in spirit, similar to [18] and does also not reflect the probability of individual words. The edit distance of weighted automata from [38] depends on the

minimum edit distance between pairs of words from the two compared languages, again regardless of their statistical significance. One might also consider using the error metric on a pair of automata introduced by Angluin in the setting of PAC (probably approximately correct) learning of DFAs [1], where n words are sampled from a given distribution and their (non-)acceptance tested in the two automata. If the outputs of both automata agree on all n words, one can say that with confidence δ the distance between the two automata is at most ϵ , where δ and ϵ can be determined from n . None of these notions is suitable for our needs.

Language-preserving minimization of a given NFA is a **PSPACE**-complete problem [24,32]. More feasible (polynomial-time) is language-preserving size reduction of NFAs based on (bi)simulations [22,40,9,11], which does not aim for a truly minimal NFA. A number of advanced variants exist, based on multi-pebble or look-ahead simulations, or on combinations of forward and backward simulations [35,16,13]. The practical efficiency of these techniques is, however, often insufficient to allow them to handle the large NFAs that occur in practice and/or they do not manage to reduce the NFAs enough. Finally, even a minimal NFA for the given set of regexes is often too big to be implemented in the given FPGA operating on the required speed (as shown even in our experiments). Our approach is capable of a much better reduction for the price of a small change of the accepted language.

2 Preliminaries

We use $\langle a, b \rangle$ to denote the set $\{x \in \mathbb{R} \mid a \leq x \leq b\}$ and \mathbb{N} to denote the set $\{0, 1, 2, \dots\}$. Given a pair of sets X_1 and X_2 , we use $X_1 \triangle X_2$ to denote their *symmetric difference*, i.e., the set $\{x \mid \exists! i \in \{1, 2\} : x \in X_i\}$. We use the notation $[v_1, \dots, v_n]$ to denote a vector of n elements, $\mathbf{1}$ to denote the all 1's vector $[1, \dots, 1]$ (the dimension of $\mathbf{1}$ is always clear from the context), \mathbf{A} to denote a matrix, and \mathbf{A}^\top for its transpose, and \mathbf{I} for the identity matrix.

In the following, we fix a finite non-empty alphabet Σ . A *nondeterministic finite automaton* (NFA) is a quadruple $\mathcal{A} = (Q, \delta, I, F)$ where Q is a finite set of states, $\delta : Q \times \Sigma \rightarrow 2^Q$ is a transition function, $I \subseteq Q$ is a set of initial states, and $F \subseteq Q$ is a set of accepting states. We use $Q[\mathcal{A}]$, $\delta[\mathcal{A}]$, $I[\mathcal{A}]$, and $F[\mathcal{A}]$ to denote Q , δ , I , and F , respectively, and $q \xrightarrow{a} q'$ to denote that $q' \in \delta(q, a)$. Often, we abuse notation and treat δ as a subset of $Q \times \Sigma \times 2^Q$. A sequence of states $\rho = q_0 \cdots q_n$ is a *run* of \mathcal{A} over a word $w = a_1 \cdots a_n \in \Sigma^*$ from a state q to a state q' , denoted as $q \xrightarrow{w, \rho} q'$, if $\forall 1 \leq i \leq n : q_{i-1} \xrightarrow{a_i} q_i$, $q_0 = q$, and $q_n = q'$. Sometimes, we use ρ in set operations where it behaves as the set of states it contains. We also use $q \xrightarrow{w} q'$ to denote that $\exists \rho \in Q^* : q \xrightarrow{w, \rho} q'$ and $q \rightsquigarrow q'$ to denote that $\exists w : q \xrightarrow{w} q'$. The *language* of a state q is defined as

$L_{\mathcal{A}}(q) = \{w \mid \exists q_F \in F : q \xrightarrow{w} q_F\}$ and its *language* (back-language) is defined as $L_{\mathcal{A}}^b(q) = \{w \mid \exists q_I \in I : q_I \xrightarrow{w} q\}$. Both notions can be naturally extended to a set $S \subseteq Q$: $L_{\mathcal{A}}(S) = \bigcup_{q \in S} L_{\mathcal{A}}(q)$ and $L_{\mathcal{A}}^b(S) = \bigcup_{q \in S} L_{\mathcal{A}}^b(q)$. We drop the subscript \mathcal{A} when the context is obvious. \mathcal{A} *accepts* the language $L(\mathcal{A})$ defined as $L(\mathcal{A}) = L_{\mathcal{A}}(I)$. \mathcal{A} is called *deterministic* (DFA) if $|I| = 1$ and $\forall q \in Q$ and $\forall a \in \Sigma : |\delta(q, a)| \leq 1$, and *unambiguous* (UFA) if $\forall w \in L(\mathcal{A}) : \exists! q_I \in I, \rho \in Q^*, q_F \in F : q_I \xrightarrow{w, \rho} q_F$.

The *restriction* of \mathcal{A} to $S \subseteq Q$ is an NFA $\mathcal{A}|_S$ given as $\mathcal{A}|_S = (S, \delta \cap (S \times \Sigma \times 2^S), I \cap S, F \cap S)$. We define the *trim* operation as $\text{trim}(\mathcal{A}) = \mathcal{A}|_C$ where $C = \{q \mid \exists q_I \in I, q_F \in F : q_I \rightsquigarrow q \rightsquigarrow q_F\}$. For a set of states $R \subseteq Q$, we use $\text{reach}(R)$ to denote the set of states reachable from R , $\text{reach}(R) = \{r' \mid \exists r \in R : r \rightsquigarrow r'\}$. We use the number of states of \mathcal{A} as a measure of its size, i.e., $|\mathcal{A}| = |Q|$.

A (discrete probability) *distribution* over a countable set X is a mapping $\text{Pr} : X \rightarrow \langle 0, 1 \rangle$ such that $\sum_{x \in X} \text{Pr}(x) = 1$. An *n*-state *probabilistic automaton* (PA) over Σ is a triple $\mathcal{P} = (\alpha, \gamma, \{\Delta_a\}_{a \in \Sigma})$ where $\alpha \in \langle 0, 1 \rangle^n$ is a vector of *initial weights*, $\gamma \in \langle 0, 1 \rangle^n$ is a vector of *final weights*, and for every $a \in \Sigma$, $\Delta_a \in \langle 0, 1 \rangle^{n \times n}$ is a *transition matrix* for symbol a . We abuse notation and use $Q[\mathcal{P}]$ to denote the set of states $Q[\mathcal{P}] = \{1, \dots, n\}$. Moreover, the following two properties need to hold: (i) $\sum \{\alpha[i] \mid i \in Q[\mathcal{P}]\} = 1$ (the initial probability is 1) and (ii) for every state $i \in Q[\mathcal{P}]$ it holds that $\sum \{\Delta_a[i, j] \mid j \in Q[\mathcal{P}], a \in \Sigma\} + \gamma[i] = 1$ (the probability of accepting or leaving a state is 1). We define the *support* of \mathcal{P} as the NFA $\text{supp}(\mathcal{P}) = (Q[\mathcal{P}], \delta[\mathcal{P}], I[\mathcal{P}], F[\mathcal{P}])$ s.t.

$$\begin{aligned} \delta[\mathcal{P}] &= \{(i, a, j) \mid \Delta_a[i, j] > 0\}, \\ I[\mathcal{P}] &= \{i \mid \alpha[i] > 0\}, \\ F[\mathcal{P}] &= \{i \mid \gamma[i] > 0\}. \end{aligned}$$

Let us assume that every PA \mathcal{P} is such that $\text{supp}(\mathcal{P}) = \text{trim}(\text{supp}(\mathcal{P}))$. For a word $w = a_1 \dots a_k \in \Sigma^*$, we use Δ_w to denote the matrix $\Delta_{a_1} \cdots \Delta_{a_k}$. For the empty word ε , we define $\Delta_\varepsilon = \mathbf{I}$. It can be easily shown that \mathcal{P} represents a distribution over words $w \in \Sigma^*$ defined as $\text{Pr}_{\mathcal{P}}(w) = \alpha^\top \cdot \Delta_w \cdot \gamma$. We call $\text{Pr}_{\mathcal{P}}(w)$ the *probability* of w in \mathcal{P} . Given a language $L \subseteq \Sigma^*$, we define the probability of L in \mathcal{P} as $\text{Pr}_{\mathcal{P}}(L) = \sum_{w \in L} \text{Pr}_{\mathcal{P}}(w)$.

In some of the proofs later, we use the PA \mathcal{P}_{Exp} defined as $\mathcal{P}_{Exp} = (\mathbf{1}, [\mu], \{[\mu]_a\}_{a \in \Sigma})$ where $\mu = \frac{1}{|\Sigma|+1}$. \mathcal{P}_{Exp} models a distribution over the words from Σ^* using a combination of an exponential distribution (for selecting the length l of a word) and the uniform distribution (for selecting symbols in a word of the length l). In particular, the purpose of \mathcal{P}_{Exp} is to assign every word $w \in \Sigma^*$ the (non-zero) probability $\text{Pr}_{\mathcal{P}_{Exp}}(w) = \mu^{|w|+1}$; any other PA assigning non-zero probabilities to all words would work as well.

If Conditions (i) and (ii) from the definition of PAs are dropped, we speak about a *pseudo-probabilistic automaton (PPA)*, which may assign a word from its support a quantity that is not necessarily in the range $\langle 0, 1 \rangle$, denoted as the *significance* of the word below. PPAs may arise during some of our operations performed on PAs. Note that PPAs can be seen as instantiations of multiplicity or weighted automata [44].

3 Approximate Reduction of NFAs

In this section, we first introduce the key notion of our approach: a *probabilistic distance* of a pair of finite automata wrt a given probabilistic automaton that, intuitively, represents the significance of particular words. We discuss the complexity of computing the probabilistic distance. Finally, we formulate two problems of *approximate automata reduction via probabilistic distance*.

3.1 Probabilistic Distance

We start by defining our notion of a probabilistic distance of two NFAs. Assume NFAs \mathcal{A}_1 and \mathcal{A}_2 and a probabilistic automaton \mathcal{P} specifying the distribution $\Pr_{\mathcal{P}} : \Sigma^* \rightarrow \langle 0, 1 \rangle$. The *probabilistic distance* $d_{\mathcal{P}}(\mathcal{A}_1, \mathcal{A}_2)$ between \mathcal{A}_1 and \mathcal{A}_2 wrt $\Pr_{\mathcal{P}}$ is defined as

$$d_{\mathcal{P}}(\mathcal{A}_1, \mathcal{A}_2) = \Pr_{\mathcal{P}}(L(\mathcal{A}_1) \triangle L(\mathcal{A}_2)).$$

Intuitively, the distance captures the significance of the words accepted by one of the automata only. We use the distance to drive the reduction process towards automata with small errors and to assess the quality of the result. (The distance is sometimes called the *symmetric difference semi-metric* [15].)

The value of $\Pr_{\mathcal{P}}(L(\mathcal{A}_1) \triangle L(\mathcal{A}_2))$ can be computed as follows. Using the fact that (1) $L_1 \triangle L_2 = (L_1 \setminus L_2) \uplus (L_2 \setminus L_1)$ and (2) $L_1 \setminus L_2 = L_1 \setminus (L_1 \cap L_2)$, we get

$$\begin{aligned} d_{\mathcal{P}}(\mathcal{A}_1, \mathcal{A}_2) &= \Pr_{\mathcal{P}}(L(\mathcal{A}_1) \setminus L(\mathcal{A}_2)) + \Pr_{\mathcal{P}}(L(\mathcal{A}_2) \setminus L(\mathcal{A}_1)) \\ &= \Pr_{\mathcal{P}}(L(\mathcal{A}_1) \setminus (L(\mathcal{A}_1) \cap L(\mathcal{A}_2))) + \\ &\quad \Pr_{\mathcal{P}}(L(\mathcal{A}_2) \setminus (L(\mathcal{A}_2) \cap L(\mathcal{A}_1))) \\ &= \Pr_{\mathcal{P}}(L(\mathcal{A}_1)) + \Pr_{\mathcal{P}}(L(\mathcal{A}_2)) - 2 \cdot \Pr_{\mathcal{P}}(L(\mathcal{A}_1) \cap L(\mathcal{A}_2)). \end{aligned}$$

Hence, the key step is to compute $\Pr_{\mathcal{P}}(L(\mathcal{A}))$ for an NFA \mathcal{A} and a PA \mathcal{P} . Problems similar to computing such a probability have been extensively studied in several contexts including verification of probabilistic systems [50, 2, 6].

In our approach, we apply the method of [6] and compute $\Pr_{\mathcal{P}}(L(\mathcal{A}))$ in the following way. We first check whether the NFA \mathcal{A} is unambiguous. This can be done by

using the standard product construction (denoted as \cap) for computing the intersection of the NFA \mathcal{A} with itself and trimming the result, formally $\mathcal{B} = \text{trim}(\mathcal{A} \cap \mathcal{A})$, followed by a check whether there is some state $(p, q) \in Q[\mathcal{B}]$ s.t. $p \neq q$ [39]. If \mathcal{A} is ambiguous, we either determinise it or disambiguate it [39], leading to a DFA/UFA \mathcal{A}' , respectively.¹ Then, we construct the trimmed product of \mathcal{A}' and \mathcal{P} (this can be seen as computing $\mathcal{A}' \cap \text{supp}(\mathcal{P})$ while keeping the probabilities from \mathcal{P} on the edges of the result), yielding a PPA $\mathcal{R} = (\alpha_{\mathcal{R}}, \gamma_{\mathcal{R}}, \{\Delta_a^{\mathcal{R}}\}_{a \in \Sigma})$.² Intuitively, \mathcal{R} represents not only the words of $L(\mathcal{A})$ but also their probability in \mathcal{P} (we give the formal definition of \mathcal{R} inside the proof of Lemma 2). Now, let $\Delta = \sum_{a \in \Sigma} \Delta_a$ be the matrix that expresses, for any $p, q \in Q[\mathcal{R}]$, the significance of getting from p to q via any $a \in \Sigma$. Further, it can be shown (cf. the proof of Lemma 1) that the matrix Δ^* , representing the significance of going from p to q via any $w \in \Sigma^*$, can be computed as $(\mathbf{I} - \Delta)^{-1}$. Then, to get $\Pr_{\mathcal{P}}(L(\mathcal{A}))$, it suffices to take $\alpha^{\top} \cdot \Delta^* \cdot \gamma$. Note that, due to the determinisation/disambiguation step, the obtained value indeed is $\Pr_{\mathcal{P}}(L(\mathcal{A}))$ despite \mathcal{R} being a PPA. The two lemmas below summarise the complexity of this step for NFAs and UFAs respectively.

Lemma 1 *Let \mathcal{P} be a PA and \mathcal{A} an NFA. The problem of computing $\Pr_{\mathcal{P}}(L(\mathcal{A}))$ is **PSPACE**-complete.*

Proof The membership in **PSPACE** can be shown as follows. The computation described above corresponds to solving a linear equation system. The system has an exponential size because of the blowup caused by the determinisation/disambiguation of \mathcal{A} required by the product construction. The equation system can, however, be constructed by a **PSPACE** transducer \mathcal{M}_{eq} . Moreover, as solving linear equation systems can be done using a polylogarithmic-space transducer \mathcal{M}_{SysLin} , one can combine these two transducers to obtain a **PSPACE** algorithm. Details of the construction follow:

First, we construct a transducer \mathcal{M}_{eq} that, given an NFA $\mathcal{A} = (Q_{\mathcal{A}}, \delta_{\mathcal{A}}, I_{\mathcal{A}}, F_{\mathcal{A}})$ and a PA $\mathcal{P} = (\alpha, \gamma, \{\Delta_a\}_{a \in \Sigma})$ on its input, constructs a system of $m = 2^{|Q_{\mathcal{A}}|} \cdot |Q[\mathcal{P}]|$ linear equations $\mathcal{S}(\mathcal{A}, \mathcal{P})$ of m unknowns $\xi_{[R,p]}$ for $R \subseteq Q_{\mathcal{A}}$ and $p \in Q[\mathcal{P}]$ representing the product of \mathcal{A}' and \mathcal{P} , where \mathcal{A}' is a deterministic automaton obtained from \mathcal{A} using the standard subset construction. The system of

¹ In theory, disambiguation can produce smaller automata, but, in our experiments, determinisation proved to work better.

² \mathcal{R} is not necessarily a PA since there might be transitions in \mathcal{P} that are either removed or copied several times in the product construction.

equations $\mathcal{S}(\mathcal{A}, \mathcal{P})$ is defined as follows (cf. [6]):

$$\xi_{[R,p]} = \begin{cases} 0 & \text{if } L_{\mathcal{A}}(R) \cap L_{\mathcal{P}'}(p) = \emptyset, \\ \sum_{a \in \Sigma} \sum_{p' \in Q[\mathcal{P}]} (\Delta_a[p, p'] \cdot \xi_{[\delta_{\mathcal{A}}(R,a), p']}) + \gamma[p] & \text{if } R \cap F_{\mathcal{A}} \neq \emptyset, \\ \sum_{a \in \Sigma} \sum_{p' \in Q[\mathcal{P}]} \Delta_a[p, p'] \cdot \xi_{[\delta_{\mathcal{A}}(R,a), p']} & \text{otherwise,} \end{cases}$$

such that $\mathcal{P}' = \text{supp}(\mathcal{P})$ and $\delta_{\mathcal{A}}(R, a) = \bigcup_{r \in R} \delta(r, a)$. The test $L_{\mathcal{A}}(R) \cap L_{\mathcal{P}'}(p) = \emptyset$ can be performed by checking $\exists r \in R : L_{\mathcal{A}}(r) \cap L_{\mathcal{P}'}(p) = \emptyset$, which can be done in polynomial time.

It holds that $\Pr_{\mathcal{P}}(L(\mathcal{A})) = \sum_{p \in Q[\mathcal{P}]} \alpha[p] \cdot \xi_{[I_{\mathcal{A}}, p]}$. Although the size of $\mathcal{S}(\mathcal{A}, \mathcal{P})$ (which is the output of \mathcal{M}_{eq}) is exponential in the size of the input of \mathcal{M}_{eq} , the internal configuration of \mathcal{M}_{eq} only needs to be of polynomial size, i.e., \mathcal{M}_{eq} works in **PSPACE**. Note that the size of each equation is at most polynomial.

Given a system \mathcal{S} of m linear equations with m unknowns, solving \mathcal{S} can be done in the time $\mathcal{O}(\log^2 m)$ using $\mathcal{O}(m^k)$ processors for a fixed k [14, Corollary 2] (i.e., it is in the class **NC**).³ According to [17, Lemma 1b], an $\mathcal{O}(\log^2 m)$ time-bounded parallel machine can be simulated by an $\mathcal{O}(\log^4 m)$ space-bounded Turing machine. Therefore, there exists an $\mathcal{O}(\log^4 m)$ space-bounded Turing machine \mathcal{M}_{SysLin} that solves a system of m linear equations with m unknowns. As a consequence, \mathcal{M}_{SysLin} can solve $\mathcal{S}(\mathcal{A}, \mathcal{P})$ using the space

$$\begin{aligned} \mathcal{O}(\log^4(2^{|Q_{\mathcal{A}}|} \cdot |Q[\mathcal{P}]|)) &= \mathcal{O}(\log^4 2^{|Q_{\mathcal{A}}|} + \log^4 |Q[\mathcal{P}]|) \\ &= \mathcal{O}(|Q_{\mathcal{A}}|^4 + \log^4 |Q[\mathcal{P}]|). \end{aligned}$$

The missing part is how to combine \mathcal{M}_{eq} and \mathcal{M}_{SysLin} to avoid using the exponential-size output tape of \mathcal{M}_{eq} . For this, we use the following standard technique for combining reductions [41, Proposition 8.2].

We take turns in simulating \mathcal{M}_{SysLin} and \mathcal{M}_{eq} . We start with simulating \mathcal{M}_{SysLin} . When \mathcal{M}_{SysLin} moves its head right, we pause it and simulate \mathcal{M}_{eq} until it outputs the corresponding bit, which is fed into the input of \mathcal{M}_{SysLin} . Then we pause \mathcal{M}_{eq} and resume the run of \mathcal{M}_{SysLin} . On the other hand, when \mathcal{M}_{SysLin} moves its head left (from the k -th position on the tape), we pause it, restart \mathcal{M}_{eq} from its initial state, and simulate it until it outputs the $(k-1)$ -st bit of its output tape, and then pause \mathcal{M}_{eq} and return the control to \mathcal{M}_{SysLin} . In order to keep track of the position k of the head of \mathcal{M}_{SysLin} on its tape, we use a binary counter.

The internal configuration of both \mathcal{M}_{eq} and \mathcal{M}_{SysLin} is of a polynomial size and the overhead of keeping track of the position of the head of \mathcal{M}_{SysLin} also requires only

polynomial space. Therefore, the whole transducer runs in a polynomially-bounded space.

The **PSPACE**-hardness is obtained by a reduction from the (**PSPACE**-complete) universality of NFAs: using the PA \mathcal{P}_{Exp} defined in §2, which assigns every word a non-zero probability, it holds that

$$L(\mathcal{A}) = \Sigma^* \quad \text{iff} \quad \Pr_{\mathcal{P}_{Exp}}(L(\mathcal{A})) = 1. \quad \square$$

Lemma 2 *Let \mathcal{P} be a PA and \mathcal{A} a UFA. The problem of computing $\Pr_{\mathcal{P}}(L(\mathcal{A}))$ is in **PTIME**.*

Proof We modify the proof from [6] into our setting. First, we give a formal definition of the *product* of a PA $\mathcal{P} = (\alpha, \gamma, \{\Delta_a\}_{a \in \Sigma})$ and an NFA $\mathcal{A} = (Q, \delta, I, F)$ as the $(|Q[\mathcal{P}]| \cdot |Q|)$ -state PPA $\mathcal{R} = (\alpha_{\mathcal{R}}, \gamma_{\mathcal{R}}, \{\Delta_a^{\mathcal{R}}\}_{a \in \Sigma})$ where⁴

$$\begin{aligned} \alpha_{\mathcal{R}}[(q_{\mathcal{P}}, q_{\mathcal{A}})] &= \alpha_{\mathcal{P}}[q_{\mathcal{P}}] \cdot |\{q_{\mathcal{A}}\} \cap I|, \\ \gamma_{\mathcal{R}}[(q_{\mathcal{P}}, q_{\mathcal{A}})] &= \gamma_{\mathcal{P}}[q_{\mathcal{P}}] \cdot |\{q_{\mathcal{A}}\} \cap F|, \\ \Delta_a^{\mathcal{R}}[(q_{\mathcal{P}}, q_{\mathcal{A}}), (q'_{\mathcal{P}}, q'_{\mathcal{A}})] &= \Delta_a[q_{\mathcal{P}}, q'_{\mathcal{P}}] \cdot |\{q'_{\mathcal{A}}\} \cap \delta(q_{\mathcal{A}}, a)|. \end{aligned}$$

Note that \mathcal{R} is not necessarily a PA any more because for $w \in \Sigma^*$ such that $\Pr_{\mathcal{P}}(w) > 0$, (i) if $w \notin L(\mathcal{A})$, then $\Pr_{\mathcal{R}}(w) = 0$ and (ii) if $w \in L(\mathcal{A})$ and \mathcal{A} can accept w using n different runs, then $\Pr_{\mathcal{R}}(w) = n \cdot \Pr_{\mathcal{P}}(w)$. As a consequence, the probabilities of all words from Σ^* are no longer guaranteed to add up to 1. If \mathcal{A} is unambiguous, the second issue is avoided and \mathcal{R} preserves the probabilities of words from $L(\mathcal{A})$, i.e., $\Pr_{\mathcal{R}}(w) = \Pr_{\mathcal{P}}(w)$ for all $w \in L(\mathcal{A})$, so \mathcal{R} can be seen as the restriction of $\Pr_{\mathcal{P}}$ to $L(\mathcal{A})$. In the following, we assume \mathcal{R} is trimmed.

In order to compute $\Pr_{\mathcal{P}}(L(\mathcal{A}))$, we construct a matrix \mathbf{E} defined as $\mathbf{E} = \sum_{a \in \Sigma} \Delta_a^{\mathcal{R}}$. Because \mathcal{R} is trimmed, the *spectral radius* of \mathbf{E} , denoted as $\rho(\mathbf{E})$, is less than one, i.e., $\rho(\mathbf{E}) < 1$ (the proof of this fact can be found, e.g., in [6]). Intuitively, $\rho(\mathbf{E}) < 1$ holds because we trimmed the redundant states from the product of \mathcal{P} and \mathcal{A} . We further use the following standard result in linear algebra: if $\rho(\mathbf{E}) < 1$, then (i) the matrix $\mathbf{I} - \mathbf{E}$ is invertible and (ii) the sum of powers of \mathbf{E} , denoted as \mathbf{E}^* , can be computed as $\mathbf{E}^* = \sum_{i=0}^{\infty} \mathbf{E}^i = (\mathbf{I} - \mathbf{E})^{-1}$ [21]. Moreover, note that matrix inversion can be done in polynomial time [46].

\mathbf{E}^* represents the reachability between nodes of \mathcal{R} , i.e., $\mathbf{E}^*[r, r']$ is the sum of significances of all (possibly infinitely many) paths from r to r' in \mathcal{R} . When related to \mathcal{P} and \mathcal{A} , the matrix \mathbf{E}^* represents the reachability in \mathcal{P} wrt $L(\mathcal{A})$, i.e.,

$$\begin{aligned} \mathbf{E}^*[(q_{\mathcal{P}}, q_{\mathcal{A}}), (q'_{\mathcal{P}}, q'_{\mathcal{A}})] \\ = \sum \left\{ \Delta_w[q_{\mathcal{P}}, q'_{\mathcal{P}}] \mid q_{\mathcal{A}} \xrightarrow{w} q'_{\mathcal{A}}, w \in \Sigma^* \right\}. \quad (1) \end{aligned}$$

⁴ We assume an implicit bijection between states of the product \mathcal{R} and $\{1, \dots, |Q[\mathcal{R}]|\}$.

³ We use $\log k$ to denote the base-2 logarithm of k .

We prove Equation (1) using the following reasoning. First, we show that

$$\begin{aligned} \mathbf{E}^n[(q_{\mathcal{P}}, q_{\mathcal{A}}), (q'_{\mathcal{P}}, q'_{\mathcal{A}})] \\ = \sum \left\{ \Delta_w[q_{\mathcal{P}}, q'_{\mathcal{P}}] \mid q_{\mathcal{A}} \xrightarrow{w} q'_{\mathcal{A}}, w \in \Sigma^n \right\}, \end{aligned} \quad (2)$$

i.e., \mathbf{E}^n represents the reachability in \mathcal{P} wrt $L(\mathcal{A})$ for words of length n . We prove Equation (2) by induction on n : For $n = 0$, the equation follows from the fact that $\mathbf{E}^0 = \mathbf{I}$. For $n = 1$, the equation follows directly from the definition of \mathcal{R} and Δ . Next, suppose that Equation (2) holds for $n > 1$; we show that it holds also for $n + 1$. We start with the following reasoning:

$$\begin{aligned} \mathbf{E}^{n+1}[(q_{\mathcal{P}}, q_{\mathcal{A}}), (q'_{\mathcal{P}}, q'_{\mathcal{A}})] \\ = (\mathbf{E}^n \mathbf{E})[(q_{\mathcal{P}}, q_{\mathcal{A}}), (q'_{\mathcal{P}}, q'_{\mathcal{A}})] \\ = \sum \left\{ \mathbf{E}^n[(q_{\mathcal{P}}, q_{\mathcal{A}}), (q''_{\mathcal{P}}, q''_{\mathcal{A}})] \cdot \mathbf{E}[(q''_{\mathcal{P}}, q''_{\mathcal{A}}), (q'_{\mathcal{P}}, q'_{\mathcal{A}})] \mid \right. \\ \left. (q''_{\mathcal{P}}, q''_{\mathcal{A}}) \in Q[\mathcal{R}] \right\}. \end{aligned}$$

The last line is obtained via the definition of matrix multiplication. Further, using the induction hypothesis, we get

$$\begin{aligned} \mathbf{E}^{n+1}[(q_{\mathcal{P}}, q_{\mathcal{A}}), (q'_{\mathcal{P}}, q'_{\mathcal{A}})] \\ = \sum \left\{ \sum \left\{ \Delta_w[q_{\mathcal{P}}, q''_{\mathcal{P}}] \mid q_{\mathcal{A}} \xrightarrow{w} q''_{\mathcal{A}}, w \in \Sigma^n \right\} \cdot \right. \\ \left. \sum \left\{ \Delta_a[q''_{\mathcal{P}}, q'_{\mathcal{P}}] \mid q''_{\mathcal{A}} \xrightarrow{a} q'_{\mathcal{A}}, a \in \Sigma \right\} \mid \right. \\ \left. (q''_{\mathcal{P}}, q''_{\mathcal{A}}) \in Q[\mathcal{R}] \right\} \\ = \sum \left\{ \sum \left\{ \Delta_w[q_{\mathcal{P}}, q''_{\mathcal{P}}] \cdot \Delta_a[q''_{\mathcal{P}}, q'_{\mathcal{P}}] \mid q_{\mathcal{A}} \xrightarrow{w} q''_{\mathcal{A}}, \right. \right. \\ \left. \left. q''_{\mathcal{A}} \xrightarrow{a} q'_{\mathcal{A}}, a \in \Sigma, w \in \Sigma^n \right\} \mid (q''_{\mathcal{P}}, q''_{\mathcal{A}}) \in Q[\mathcal{R}] \right\} \\ = \sum \left\{ \Delta_{w'}[q_{\mathcal{P}}, q'_{\mathcal{P}}] \mid q_{\mathcal{A}} \xrightarrow{w'} q'_{\mathcal{A}}, w' \in \Sigma^{n+1} \right\}. \end{aligned}$$

Since $\mathbf{E}^* = \sum_{i=0}^{\infty} \mathbf{E}^i$, Equation (1) follows. Using the matrix \mathbf{E}^* , it remains to compute $\Pr_{\mathcal{P}}(L(\mathcal{A}))$ as

$$\Pr_{\mathcal{P}}(L(\mathcal{A})) = \alpha_{\mathcal{R}}^{\top} \cdot \mathbf{E}^* \cdot \gamma_{\mathcal{R}}. \quad \square$$

3.2 Automata Reduction using Probabilistic Distance

We now exploit the probabilistic distance introduced above to formulate the task of approximate reduction of NFAs as two optimisation problems. Given an NFA \mathcal{A} and a PA \mathcal{P} specifying the distribution $\Pr_{\mathcal{P}} : \Sigma^* \rightarrow \langle 0, 1 \rangle$, we define

- **size-driven reduction**: for $n \in \mathbb{N}$, find an NFA \mathcal{A}' such that $|\mathcal{A}'| \leq n$ and the distance $d_{\mathcal{P}}(\mathcal{A}, \mathcal{A}')$ is minimal,
- **error-driven reduction**: for $\epsilon \in \langle 0, 1 \rangle$, find an NFA \mathcal{A}' such that $d_{\mathcal{P}}(\mathcal{A}, \mathcal{A}') \leq \epsilon$ and the size $|\mathcal{A}'|$ is minimal.

The following lemma shows that the natural decision problem underlying both of the above optimization problems is **PSPACE**-complete, which matches the complexity of computing the probabilistic distance as well as that of the *exact* reduction of NFAs [24].

Lemma 3 *Consider an NFA \mathcal{A} , a PA \mathcal{P} , a bound on the number of states $n \in \mathbb{N}$, and an error bound $\epsilon \in \langle 0, 1 \rangle$. It is **PSPACE**-complete to determine whether there exists an NFA \mathcal{A}' with n states s.t. $d_{\mathcal{P}}(\mathcal{A}, \mathcal{A}') \leq \epsilon$.*

Proof Membership in PSPACE: We non-deterministically generate an automaton \mathcal{A}' with n states and test (in **PSPACE**, as shown in Lemma 1) that $d_{\mathcal{P}}(\mathcal{A}, \mathcal{A}') \leq \epsilon$. This shows the problem is in **NPSPACE** = **PSPACE**.

PSPACE-hardness: We use a reduction from the problem of checking universality of an NFA $\mathcal{A} = (Q, \delta, I, F)$ over Σ , i.e., from checking whether $L(\mathcal{A}) = \Sigma^*$, which is **PSPACE**-complete. First, for a reason that will become clear later, we test if \mathcal{A} accepts all words over Σ of length 0 and 1, which can be done in polynomial time. It holds that $L(\mathcal{A}) = \Sigma^*$ iff there is a 1-state NFA \mathcal{A}' s.t. $d_{\mathcal{P}_{Exp}}(\mathcal{A}, \mathcal{A}') \leq 0$ (\mathcal{P}_{Exp} is defined in §2). The implication from left to right is clear: \mathcal{A}' can be constructed as $\mathcal{A}' = (\{q\}, \{q \xrightarrow{a} q \mid a \in \Sigma\}, \{q\}, \{q\})$. To show the reverse implication, we note that we have tested that $\{\epsilon\} \cup \Sigma \subseteq L(\mathcal{A})$. Since the probability of any word from $\{\epsilon\} \cup \Sigma \subseteq L(\mathcal{A})$ in \mathcal{P}_{Exp} is non-zero, the only 1-state NFA that processes those words with zero error is the NFA \mathcal{A}' defined above. Because the language of \mathcal{A}' is $L(\mathcal{A}') = \Sigma^*$, it holds that $d_{\mathcal{P}_{Exp}}(\mathcal{A}, \mathcal{A}') \leq 0$ iff $L(\mathcal{A}) = \Sigma^*$. \square

The notions defined above do not distinguish between introducing a *false positive* (\mathcal{A}' accepts a word $w \notin L(\mathcal{A})$) or a *false negative* (\mathcal{A}' rejects a word $w \in L(\mathcal{A})$) answers. To this end, we define *over-approximating* and *under-approximating* reductions as reductions for which the conditions $L(\mathcal{A}) \subseteq L(\mathcal{A}')$ and $L(\mathcal{A}) \supseteq L(\mathcal{A}')$ hold.

A naïve solution to the reductions would enumerate all NFAs \mathcal{A}' of sizes from 0 up to k (resp. $|\mathcal{A}|$), for each of them compute $d_{\mathcal{P}}(\mathcal{A}, \mathcal{A}')$, and take an automaton with the smallest probabilistic distance (resp. a smallest one satisfying the restriction on $d_{\mathcal{P}}(\mathcal{A}, \mathcal{A}')$). Obviously, this approach is computationally infeasible.

4 A Heuristic Approach to Approximate Reduction

In this section, we introduce two techniques for approximate reduction of NFAs that avoid the need to iterate over all

Algorithm 1: A greedy size-driven reduction

Input : NFA $\mathcal{A} = (Q, \delta, I, F)$, PA \mathcal{P} , $n \geq 1$
Output: NFA \mathcal{A}' , $\epsilon \in \mathbb{R}$ s.t. $|\mathcal{A}'| \leq n$ and $d_{\mathcal{P}}(\mathcal{A}, \mathcal{A}') \leq \epsilon$

```

1  $V \leftarrow \emptyset$ ;
2 for  $q \in Q$  in the order  $\preceq_{\mathcal{A}, \text{label}(\mathcal{A}, \mathcal{P})}$  do
3    $V \leftarrow V \cup \{q\}$ ;  $\mathcal{A}' \leftarrow \text{reduce}(\mathcal{A}, V)$ ;
4   if  $|\mathcal{A}'| \leq n$  then break ;
5 return  $\mathcal{A}'$ ,  $\epsilon = \text{error}(\mathcal{A}, V, \text{label}(\mathcal{A}, \mathcal{P}))$ ;
```

automata of a certain size. The first approach is based on under-approximating the automata by removing states—we call it the *pruning reduction*—while the second approach is based on over-approximating the automata by adding self-loops to states and removing redundant states—we call it the *self-loop reduction*. Finding an optimal automaton using these reductions is also **PSPACE**-complete, but more amenable to heuristics like greedy algorithms. We start with introducing two high-level greedy algorithms, one for the size- and one for the error-driven reduction, and follow by showing their instantiations for the pruning and the self-loop reduction. A crucial role in the algorithms is played by a function that labels states of the automata by an estimate of the error that will be caused when some of the reductions is applied at a given state.

4.1 A General Algorithm for Size-Driven Reduction

Algorithm 1 shows a general greedy method for performing the size-driven reduction. In order to use the same high-level algorithm in both directions of reduction (over/under-approximating), it is parameterized with the functions: *label*, *reduce*, and *error*. The real intricacy of the procedure is hidden inside these three functions. Intuitively, $\text{label}(\mathcal{A}, \mathcal{P})$ assigns every state of an NFA \mathcal{A} an approximation of the error that will be caused wrt the PA \mathcal{P} when a reduction is applied at this state, while the purpose of $\text{reduce}(\mathcal{A}, V)$ is to create a new NFA \mathcal{A}' obtained from \mathcal{A} by introducing some error at states from V .⁵ Further, $\text{error}(\mathcal{A}, V, \text{label}(\mathcal{A}, \mathcal{P}))$ estimates the error introduced by the application of $\text{reduce}(\mathcal{A}, V)$, possibly in a more precise (and costly) way than by just summing the concerned error labels: Such a computation is possible outside of the main computation loop. We show instantiations of these functions later, when discussing the reductions used. Moreover, the algorithm is also parameterized with a total order $\preceq_{\mathcal{A}, \text{label}(\mathcal{A}, \mathcal{P})}$ that defines which states of \mathcal{A} are

⁵ We emphasize that this does not mean that states from V will be simply removed from \mathcal{A} —the performed operation depends on the particular reduction.

Algorithm 2: A greedy error-driven reduction.

Input : NFA $\mathcal{A} = (Q, \delta, I, F)$, PA \mathcal{P} , $\epsilon \in \langle 0, 1 \rangle$
Output: NFA \mathcal{A}' s.t. $d_{\mathcal{P}}(\mathcal{A}, \mathcal{A}') \leq \epsilon$

```

1  $\ell \leftarrow \text{label}(\mathcal{A}, \mathcal{P})$ ;
2  $V \leftarrow \emptyset$ ;
3 for  $q \in Q$  in the order  $\preceq_{\mathcal{A}, \text{label}(\mathcal{A}, \mathcal{P})}$  do
4    $e \leftarrow \text{error}(\mathcal{A}, V \cup \{q\}, \ell)$ ;
5   if  $e \leq \epsilon$  then  $V \leftarrow V \cup \{q\}$  ;
6 return  $\mathcal{A}' = \text{reduce}(\mathcal{A}, V)$ ;
```

processed first and which are processed later. The ordering may take into account the precomputed labelling. The algorithm accepts an NFA \mathcal{A} , a PA \mathcal{P} , and $n \in \mathbb{N}$ and outputs a pair consisting of an NFA \mathcal{A}' of the size $|\mathcal{A}'| \leq n$ and an error bound ϵ such that $d_{\mathcal{P}}(\mathcal{A}, \mathcal{A}') \leq \epsilon$.

The main idea of the algorithm is that it creates a set V of states where an error is to be introduced. V is constructed by starting from an empty set and adding states to it in the order given by $\preceq_{\mathcal{A}, \text{label}(\mathcal{A}, \mathcal{P})}$, until the size of the result of $\text{reduce}(\mathcal{A}, V)$ has reached the desired bound n (in our setting, *reduce* is always antitone, i.e., for $V \subseteq V'$, it holds that $|\text{reduce}(\mathcal{A}, V)| \geq |\text{reduce}(\mathcal{A}, V')|$). We now define the necessary condition for *label*, *reduce*, and *error* that makes Algorithm 1 correct.

Condition C1 holds if for every NFA \mathcal{A} , PA \mathcal{P} , and a set $V \subseteq Q[\mathcal{A}]$, we have that

- (a) $\text{error}(\mathcal{A}, V, \text{label}(\mathcal{A}, \mathcal{P})) \geq d_{\mathcal{P}}(\mathcal{A}, \text{reduce}(\mathcal{A}, V))$,
- (b) $|\text{reduce}(\mathcal{A}, Q[\mathcal{A}])| \leq 1$, and
- (c) $\text{reduce}(\mathcal{A}, \emptyset) = \mathcal{A}$.

C1(a) ensures that the error computed by the reduction algorithm indeed over-approximates the exact probabilistic distance, **C1(b)** is a boundary condition for the case when the reduction is applied at every state of \mathcal{A} , and **C1(c)** ensures that when no error is to be introduced at any state, we obtain the original automaton.

Lemma 4 Algorithm 1 is correct if **C1** holds.

Proof Follows straightforwardly from Condition **C1**. \square

4.2 A General Algorithm for Error-Driven Reduction

In Algorithm 2, we provide a high-level method of computing the error-driven reduction. The algorithm is in many ways similar to Algorithm 1; it also computes a set of states V where an error is to be introduced, but an important difference is that we compute an approximation of the error in each step and only add q to V if it does not raise the error over the threshold ϵ . Note that the *error* does not need to be

monotone, so it may be advantageous to traverse all states from Q and not terminate as soon as the threshold is reached. The correctness of Algorithm 2 also depends on **C1**.

Lemma 5 *Algorithm 2 is correct if **C1** holds.*

Proof Follows straightforwardly from Condition **C1**. \square

4.3 Pruning Reduction

The pruning reduction is based on identifying a set of states to be removed from an NFA \mathcal{A} , under-approximating the language of \mathcal{A} . In particular, for $\mathcal{A} = (Q, \delta, I, F)$, the pruning reduction finds a set $R \subseteq Q$ and restricts \mathcal{A} to $Q \setminus R$, followed by removing useless states, to construct a reduced automaton $\mathcal{A}' = \text{trim}(\mathcal{A}_{|Q \setminus R})$. Note that the natural decision problem corresponding to this reduction is also **PSPACE**-complete.

Lemma 6 *Consider an NFA \mathcal{A} , a PA \mathcal{P} , a bound on the number of states $n \in \mathbb{N}$, and an error bound $\epsilon \in (0, 1)$. It is **PSPACE**-complete to determine whether there exists a subset of states $R \subseteq Q[\mathcal{A}]$ of size $|R| = n$ such that $d_{\mathcal{P}}(\mathcal{A}, \mathcal{A}_{|R}) \leq \epsilon$.*

Proof Membership in **PSPACE**: We non-deterministically generate a subset R of $Q[\mathcal{A}]$ having n states and test (in **PSPACE**, as shown in Lemma 1) that $d_{\mathcal{P}}(\mathcal{A}, \mathcal{A}_{|R}) \leq \epsilon$. This shows the problem is in **NPSPACE** = **PSPACE**.

PSPACE-hardness: We use a reduction from the **PSPACE**-complete problem of checking universality of an NFA $\mathcal{A} = (Q, \delta, I, F)$ over Σ . Consider a symbol $x \notin \Sigma$. Let us construct an NFA \mathcal{A}' over $\Sigma \cup \{x\}$ s.t. $L(\mathcal{A}') = x^*.L(\mathcal{A})$. \mathcal{A}' is constructed by adding a fresh state q_{new} to \mathcal{A} that can loop over x and make a transition to any initial state of \mathcal{A} over x : $\mathcal{A}' = (Q \cup \{q_{new}\}, \delta \cup \{q_{new} \xrightarrow{x} q \mid q \in I \cup \{q_{new}\}\}, I \cup \{q_{new}\}, F)$. We set $n = |\mathcal{A}'| + 1$. Further, we also construct an $(n+1)$ -state NFA \mathcal{B} accepting the language $x^n.\Sigma^*$ defined as $\mathcal{B} = (Q_{\mathcal{B}}, \delta_{\mathcal{B}}, \{q_1\}, \{q_{n+1}\})$ where $Q_{\mathcal{B}} = \{q_1, \dots, q_{n+1}\}$ and $\delta_{\mathcal{B}} = \{q_i \xrightarrow{x} q_{i+1} \mid 1 \leq i \leq n\} \cup \{q_{n+1} \xrightarrow{a} q_{n+1} \mid a \in \Sigma\}$. Moreover, let \mathcal{P} be a PA representing a distribution $\text{Pr}_{\mathcal{P}}$ that is defined for each $w \in (\Sigma \cup \{x\})^*$ as

$$\text{Pr}_{\mathcal{P}}(w) = \begin{cases} \mu^{|w'|+1} & \text{for } w = x^n.w', w' \in \Sigma^*, \\ & \text{and } \mu = \frac{1}{|\Sigma|+1}, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Note that $\text{Pr}_{\mathcal{P}}(x^n.w) = \text{Pr}_{\mathcal{P}_{Exp}}(w)$ for $w \in \Sigma^*$, and $\text{Pr}_{\mathcal{P}}(u) = 0$ for $u \notin x^n.\Sigma^*$ (\mathcal{P} can be easily constructed from \mathcal{P}_{Exp}). Also note that \mathcal{B} accepts exactly those words w such that $\text{Pr}_{\mathcal{P}}(w) \neq 0$ and that $\text{Pr}_{\mathcal{P}}(L(\mathcal{B})) = 1$. Using the automata defined above, we construct an NFA $\mathcal{C} = \mathcal{A}' \cup \mathcal{B}$ where the union of two NFAs is

defined as $\mathcal{A}_1 \cup \mathcal{A}_2 = (Q[\mathcal{A}_1] \uplus Q[\mathcal{A}_2], \delta[\mathcal{A}_1] \uplus \delta[\mathcal{A}_2], I[\mathcal{A}_1] \uplus I[\mathcal{A}_2], F[\mathcal{A}_1] \uplus F[\mathcal{A}_2])$. NFA \mathcal{C} has $2n$ states, the language of \mathcal{C} is $L(\mathcal{C}) = x^*.L(\mathcal{A}) \cup x^n.\Sigma^*$ and its probability is $\text{Pr}_{\mathcal{P}}(L(\mathcal{C})) = 1$.

The important property of \mathcal{C} is that if there exists a set $R \subseteq Q[\mathcal{C}]$ of the size $|R| = n$ s.t. $d_{\mathcal{P}}(\mathcal{C}, \mathcal{C}_{|R}) \leq 0$, then $L(\mathcal{A}) = \Sigma^*$. The property holds because since $|Q[\mathcal{A}']| = n - 1$, when we remove n states from \mathcal{C} , at least one state from $Q[\mathcal{B}]$ is removed, making the whole subautomaton of \mathcal{C} corresponding to \mathcal{B} useless, and, therefore, $L(\mathcal{C}_{|R}) \subseteq x^*.L(\mathcal{A})$. Because $d_{\mathcal{P}}(\mathcal{C}, \mathcal{C}_{|R}) \leq 0$, we know that $\text{Pr}_{\mathcal{P}}(L(\mathcal{C}_{|R})) = 1$, so $x^n.\Sigma^* \subseteq x^*.L(\mathcal{A}) = L(\mathcal{C}_{|R})$ and, therefore, $L(\mathcal{A}) = \Sigma^*$. For the other direction, if $L(\mathcal{A}) = \Sigma^*$, then there exists a set $R \subseteq Q[\mathcal{A}'] \cup Q[\mathcal{B}]$ of the size $|R| = n$ s.t. $d_{\mathcal{P}}(\mathcal{C}, \mathcal{C}_{|R}) \leq 0$ (in particular, R can be such that $R \subseteq Q[\mathcal{B}]$). \square

Although Lemma 6 shows that the pruning reduction is as hard as a general reduction (cf. Lemma 3), the pruning reduction is more amenable to using heuristics like the greedy algorithms from §4.1 and §4.2. We instantiate *reduce*, *error*, and *label* in these high-level algorithms in the following way (the subscript p stands for *pruning*):

$$\begin{aligned} \text{reduce}_p(\mathcal{A}, V) &= \text{trim}(\mathcal{A}_{|Q \setminus V}), \\ \text{error}_p(\mathcal{A}, V, \ell) &= \min_{V' \in [V]_p} \sum \{\ell(q) \mid q \in V'\}, \end{aligned}$$

where $[V]_p$ is defined in the rest of this paragraph: Because of the use of *trim* in *reduce* _{p} , for a pair of sets V, V' s.t. $V \subset V'$, it holds that *reduce* _{p} (\mathcal{A}, V) may, in general, yield the same automaton as *reduce* _{p} (\mathcal{A}, V'). Therefore, in order to obtain a tight approximation, we wish to compute the least error that is obtained when removing the states in V . We define a partial order \sqsubseteq_p on 2^Q as $V_1 \sqsubseteq_p V_2$ iff *reduce* _{p} (\mathcal{A}, V_1) = *reduce* _{p} (\mathcal{A}, V_2) and $V_1 \subseteq V_2$, and use $[V]_p$ to denote the set of minimal elements of the set of elements that are smaller than V (wrt \sqsubseteq_p). The value of the approximation *error* _{p} (\mathcal{A}, V, ℓ) is therefore the minimum of the sum of errors over all sets from $[V]_p$.

Note that the size of $[V]_p$ can again be exponential, and thus we employ a greedy approach for guessing an optimal V' . Clearly, this cannot affect the soundness of the algorithm, but only decreases the precision of the bound on the distance. Our experiments indicate that for automata appearing in NIDSes, this simplification has typically only a negligible impact on the precision of the bounds.

For computing the state labelling, we provide the following three functions, which differ in the precision they provide and the difficulty of their computation (naturally, more precise labellings are harder to compute): *label* _{p} ¹, *label* _{p} ², and *label* _{p} ³. Given an NFA \mathcal{A} and a PA \mathcal{P} , they

generate the labellings ℓ_p^1 , ℓ_p^2 , and ℓ_p^3 , respectively, defined as

$$\begin{aligned}\ell_p^1(q) &= \sum \left\{ \Pr_{\mathcal{P}}(L_{\mathcal{A}}^b(q')) \mid q' \in \text{reach}(\{q\}) \cap F \right\}, \\ \ell_p^2(q) &= \Pr_{\mathcal{P}} \left(L_{\mathcal{A}}^b(F \cap \text{reach}(q)) \right), \\ \ell_p^3(q) &= \Pr_{\mathcal{P}} \left(L_{\mathcal{A}}^b(q) \cdot L_{\mathcal{A}}(q) \right).\end{aligned}$$

A state label $\ell(q)$ approximates the error of the words removed from $L(\mathcal{A})$ when q is removed. More concretely, $\ell_p^1(q)$ is a rough estimate saying that the error can be bounded by the sum of probabilities of the languages of all final states reachable from q (in the worst case, all those final states might become unreachable). Note that $\ell_p^1(q)$ (1) counts the error of a word accepted in two different final states of $\text{reach}(q)$ twice and (2) it also considers words that are accepted in some final state in $\text{reach}(q)$ without going through q . The labelling ℓ_p^2 deals with (1) by computing the total probability of the language of the set of all final states reachable from q , and the labelling ℓ_p^3 in addition also deals with (2) by only considering words that traverse through q (they can, however, be accepted in some final state not in $\text{reach}(q)$ by a run completely disjoint from q and $\text{reach}(q) \cap F$, so even ℓ_p^3 can still be imprecise). Note that if \mathcal{A} is unambiguous, then $\ell_p^1 = \ell_p^2$.

Each state labelling is given as the probability (or the sum of probabilities in the case of ℓ_p^1) of the language related to q . Therefore, when computing the particular label of q , we first modify \mathcal{A} to obtain \mathcal{A}' accepting the language related to the labelling. Then, we compute the value of $\Pr_{\mathcal{P}}(L(\mathcal{A}'))$ using the algorithm from §3.1. Recall that this step is in general costly, due to the determinisation/disambiguation of \mathcal{A}' . The key property of the labelling computation resides in the fact that if \mathcal{A} is composed of several disjoint sub-automata, the automaton \mathcal{A}' is typically much smaller than \mathcal{A} and thus the computation of the label is considerably less demanding. Since the automata appearing in regex matching for NIDS are composed of the union of “tentacles”, the particular \mathcal{A}' s are very small, which enables an efficient component-wise computation of the labels.

The following lemma states the correctness of using the pruning reduction as an instantiation of Algorithms 1 and 2 and also the relation among ℓ_p^1 , ℓ_p^2 , and ℓ_p^3 .

Lemma 7 *For every $x \in \{1, 2, 3\}$, the functions reduce_p , error_p , and label_p^x satisfy **C1**. Moreover, consider an NFA \mathcal{A} , a PA \mathcal{P} , and let $\ell_p^x = \text{label}_p^x(\mathcal{A}, \mathcal{P})$ for $x \in \{1, 2, 3\}$. Then, for each $q \in Q[\mathcal{A}]$, we have $\ell_p^1(q) \geq \ell_p^2(q) \geq \ell_p^3(q)$.*

Proof We start by proving the inequalities $\ell_p^1(q) \geq \ell_p^2(q) \geq \ell_p^3(q)$ for each $q \in Q[\mathcal{A}]$, which will then help us prove the first part of the lemma. The first inequality follows from the fact that if the languages of reachable final states are not disjoint, in the case of ℓ_p^1 , we may sum probabilities of the

same words multiple times. The second inequality follows from the inclusion $L_{\mathcal{A}}^b(q) \cdot L_{\mathcal{A}}(q) \subseteq L_{\mathcal{A}}^b(F \cap \text{reach}(q))$.

Second, we prove that the functions reduce_p , error_p , and label_p^x satisfy the properties of **C1**:

- **C1(a)**: In order to show the inequality

$$\text{error}_p(\mathcal{A}, V, \text{label}_p^x(\mathcal{A}, \mathcal{P})) \geq d_{\mathcal{P}}(\mathcal{A}, \text{reduce}_p(\mathcal{A}, V)),$$

we prove it for $\ell_p^3 = \text{label}_p^3(\mathcal{A}, \mathcal{P})$; the rest follows from $\ell_p^1(q) \geq \ell_p^2(q) \geq \ell_p^3(q)$, which is proved above.

Consider some set of states $V \subseteq Q[\mathcal{A}]$ and a set $V' \subseteq [V]_p$ s.t. for any $V'' \in [V]_p$, it holds that $\sum \{\ell_p^3(q) \mid q \in V'\} \leq \sum \{\ell_p^3(q) \mid q \in V''\}$. We have

$$\begin{aligned}L(\mathcal{A}) \triangle L(\text{reduce}_p(\mathcal{A}, V)) &= L(\mathcal{A}) \triangle L(\text{reduce}_p(\mathcal{A}, V')) \quad \{\text{def. of } \sqsubseteq_p\} \\ &= L(\mathcal{A}) \setminus L(\text{reduce}_p(\mathcal{A}, V')) \\ &\quad \{L(\mathcal{A}) \supseteq L(\text{reduce}_p(\mathcal{A}, V'))\} \\ &\subseteq \bigcup_{q \in V'} L_{\mathcal{A}}^b(q) \cdot L_{\mathcal{A}}(q). \quad \{\text{def. of } \text{reduce}_p\}\end{aligned} \tag{4}$$

Finally, using (4), we obtain

$$\begin{aligned}d_{\mathcal{P}}(\mathcal{A}, \text{reduce}_p(\mathcal{A}, V)) &= \Pr_{\mathcal{P}}(L(\mathcal{A}) \triangle L(\text{reduce}_p(\mathcal{A}, V'))) \\ &\quad \{\text{def. of } d_{\mathcal{P}}\} \\ &\leq \sum_{q \in V'} \Pr_{\mathcal{P}}(L_{\mathcal{A}}^b(q) \cdot L_{\mathcal{A}}(q)) \quad \{(4)\} \\ &= \sum \{\ell_p^3(q) \mid q \in V'\} \quad \{\text{def. of } \ell_p^3\} \\ &= \min_{V'' \in [V]_p} \sum \{\ell_p^3(q) \mid q \in V''\} \quad \{\text{def. of } V'\} \\ &= \text{error}_p(\mathcal{A}, V, \ell_p^3). \quad \{\text{def. of } \text{error}_p\}\end{aligned}$$

- **C1(b)**: $|\text{reduce}_p(\mathcal{A}, Q[\mathcal{A}])| \leq 1$ because

$$|\text{reduce}_p(\mathcal{A}, Q[\mathcal{A}])| = |\text{trim}(\mathcal{A}_{|\emptyset})| = 0.$$

- **C1(c)**: $\text{reduce}_p(\mathcal{A}, \emptyset) = \mathcal{A}$ since

$$\text{reduce}_p(\mathcal{A}, \emptyset) = \text{trim}(\mathcal{A}_{|Q[\mathcal{A}]}) = \mathcal{A}$$

(we assume that \mathcal{A} is trimmed at the input). \square

4.4 Self-loop Reduction

The main idea of the self-loop reduction is to over-approximate the language of \mathcal{A} by adding self-loops over every symbol at selected states. This makes some states of \mathcal{A} redundant, allowing them to be removed without introducing any more error. Given an NFA $\mathcal{A} = (Q, \delta, I, F)$, the self-loop reduction searches for a set of states $R \subseteq Q$,

which will have self-loops added, and removes other transitions leading out of these states, making some states unreachable. The unreachable states are then removed.

Formally, let $sl(\mathcal{A}, R)$ be the NFA $(Q \cup \{s\}, \delta', I, F \cup \{s\})$ where $s \notin Q$ and the transition function δ' is defined such that $\delta'(s, a) = \{s\}$ and, for all states $p \in Q$ and symbols $a \in \Sigma$, $\delta'(p, a) = (\delta(p, a) \setminus R) \cup \{s\}$ if $\delta(p, a) \cap R \neq \emptyset$ and $\delta'(p, a) = \delta(p, a)$ otherwise. Similarly to the pruning reduction, the natural decision problem corresponding to the self-loop reduction is also **PSPACE**-complete.

Lemma 8 *Consider an NFA \mathcal{A} , a PA \mathcal{P} , a bound on the number of states $n \in \mathbb{N}$, and an error bound $\epsilon \in (0, 1)$. It is **PSPACE**-complete to determine whether there exists a subset of states $R \subseteq Q[\mathcal{A}]$ of size $|R| = n$ such that $d_{\mathcal{P}}(\mathcal{A}, sl(\mathcal{A}, R)) \leq \epsilon$.*

Proof Membership in **PSPACE** can be proved in the same way as in the proof of Lemma 6.

PSPACE-hardness: We reduce from the **PSPACE**-complete problem of checking universality of an NFA $\mathcal{A} = (Q, \delta, I, F)$. First, we check whether $I[\mathcal{A}] \neq \emptyset$. We have that $L(\mathcal{A}) = \Sigma^*$ iff there exists a set of states $R \subseteq Q$ of the size $|R| = |Q|$ such that $d_{\mathcal{P}_{Exp}}(\mathcal{A}, sl(\mathcal{A}, R)) \leq 0$ (note that this means that a self-loop is added to every state of \mathcal{A}). \square

The required functions in the error- and size-driven reduction algorithms are instantiated in the following way (the subscript sl stands for *self-loop*):

$$\begin{aligned} reduce_{sl}(\mathcal{A}, V) &= trim(sl(\mathcal{A}, V)), \\ error_{sl}(\mathcal{A}, V, \ell) &= \sum \{\ell(q) \mid q \in \min([V]_{sl})\}, \end{aligned}$$

where $[V]_{sl}$ is defined in a similar manner as $[V]_p$ in the previous section (using a partial order \sqsubseteq_{sl} defined similarly to \sqsubseteq_p ; the difference is that in this case, the order \sqsubseteq_{sl} has a single minimal element, though).

The functions $label_{sl}^1$, $label_{sl}^2$, and $label_{sl}^3$ compute the state labellings ℓ_{sl}^1 , ℓ_{sl}^2 , and ℓ_{sl}^3 for an NFA \mathcal{A} and a PA \mathcal{P} , which are defined as follows:

$$\begin{aligned} \ell_{sl}^1(q) &= weight_{\mathcal{P}}(L_{\mathcal{A}}^b(q)), \\ \ell_{sl}^2(q) &= \Pr_{\mathcal{P}}(L_{\mathcal{A}}^b(q) \cdot \Sigma^*), \\ \ell_{sl}^3(q) &= \ell_{sl}^2(q) - \Pr_{\mathcal{P}}(L_{\mathcal{A}}^b(q) \cdot L_{\mathcal{A}}(q)). \end{aligned}$$

In the definitions above, the function $weight_{\mathcal{P}}(w)$ for a PA $\mathcal{P} = (\alpha, \gamma, \{\Delta_a\}_{a \in \Sigma})$ and a word $w \in \Sigma^*$ is defined as $weight_{\mathcal{P}}(w) = \alpha^T \cdot \Delta_w \cdot \mathbf{1}$ (i.e., similarly as $\Pr_{\mathcal{P}}(w)$ but with the final weights γ discarded), and $weight_{\mathcal{P}}(L)$ for $L \subseteq \Sigma^*$ is defined as $weight_{\mathcal{P}}(L) = \sum_{w \in L} weight_{\mathcal{P}}(w)$.

Intuitively, the state labelling $\ell_{sl}^1(q)$ computes the probability that q is reached from an initial state, so if q is pumped up with all possible word endings, this is the

maximum possible error introduced by the added word endings. This has the following sources of imprecision: (1) the probability of some words may be included twice, e.g., when $L_{\mathcal{A}}^b(q) = \{a, ab\}$, the probabilities of all words from $\{ab\} \cdot \Sigma^*$ are included twice in $\ell_{sl}^1(q)$ because $\{ab\} \cdot \Sigma^* \subseteq \{a\} \cdot \Sigma^*$, and (2) $\ell_{sl}^1(q)$ can also contain probabilities of words already accepted on a run traversing q . The state labelling ℓ_{sl}^2 deals with (1) by considering the probability of the language $L_{\mathcal{A}}^b(q) \cdot \Sigma^*$, and ℓ_{sl}^3 deals also with (2) by subtracting from the result of ℓ_{sl}^2 the probabilities of the words that pass through q and are accepted.

The computation of the state labellings for the self-loop reduction is done in a similar way as the computation of the state labellings for the pruning reduction (cf. §4.3). For a computation of $weight_{\mathcal{P}}(L)$ one can use the same algorithm as for $\Pr_{\mathcal{P}}(L)$, only the final vector for PA \mathcal{P} is set to $\mathbf{1}$. The correctness of Algorithms 1 and 2 when instantiated using the self-loop reduction is stated in the following lemma.

Lemma 9 *For every $x \in \{1, 2, 3\}$, the functions $reduce_{sl}$, $error_{sl}$, and $label_{sl}^x$ satisfy **C1**. Moreover, consider an NFA \mathcal{A} , a PA \mathcal{P} , and let $\ell_{sl}^x = label_{sl}^x(\mathcal{A}, \mathcal{P})$ for $x \in \{1, 2, 3\}$. Then, for each $q \in Q[\mathcal{A}]$, we have $\ell_{sl}^1(q) \geq \ell_{sl}^2(q) \geq \ell_{sl}^3(q)$.*

Proof First, we prove the inequalities $\ell_{sl}^1(q) \geq \ell_{sl}^2(q) \geq \ell_{sl}^3(q)$ for each $q \in Q[\mathcal{A}]$, which we then use to prove the first part of the lemma. We start with the equality $weight_{\mathcal{P}}(w) = \Pr_{\mathcal{P}}(w \cdot \Sigma^*)$, which follows from the fact that for each state p of \mathcal{P} the sum of probabilities of all words, when considering p as the only initial state of \mathcal{P} , is 1. Then, we obtain the equality

$$\sum_{w \in L_{\mathcal{A}}^b(q)} weight_{\mathcal{P}}(w) = \sum_{w \in L_{\mathcal{A}}^b(q)} \Pr_{\mathcal{P}}(w \cdot \Sigma^*),$$

which, in turn, implies

$$\begin{aligned} \ell_{sl}^1(q) &= weight_{\mathcal{P}}(L_{\mathcal{A}}^b(q)) = \sum_{w \in L_{\mathcal{A}}^b(q)} \Pr_{\mathcal{P}}(w \cdot \Sigma^*) \\ &\geq \Pr_{\mathcal{P}}(L_{\mathcal{A}}^b(q) \cdot \Sigma^*) = \ell_{sl}^2(q). \end{aligned} \quad (5)$$

For example, for $L_{\mathcal{A}}^b(q) = \{w, wa\}$ where $w \in \Sigma^*$ and $a \in \Sigma$, we have

$$\begin{aligned} weight_{\mathcal{P}}(L_{\mathcal{A}}^b(q)) &= weight_{\mathcal{P}}(\{w, wa\}) \\ &= weight_{\mathcal{P}}(w) + weight_{\mathcal{P}}(wa) \\ &= \Pr_{\mathcal{P}}(w \cdot \Sigma^*) + \Pr_{\mathcal{P}}(wa \cdot \Sigma^*), \end{aligned} \quad (6)$$

while

$$\Pr_{\mathcal{P}}(L_{\mathcal{A}}^b(q) \cdot \Sigma^*) = \Pr_{\mathcal{P}}(\{w, wa\} \cdot \Sigma^*) = \Pr_{\mathcal{P}}(w \cdot \Sigma^*).$$

The inequality $\ell_{sl}^2 \geq \ell_{sl}^3$ holds trivially.

Second, we prove that the functions $reduce_{sl}$, $error_{sl}$, and $label_{sl}^x$ satisfy the properties of **C1**:

- **C1(a)**: To show that $error_{sl}(\mathcal{A}, V, label_{sl}^x(\mathcal{A}, \mathcal{P})) \geq d_{\mathcal{P}}(\mathcal{A}, reduce_{sl}(\mathcal{A}, V))$, we prove that the inequality holds for $\ell_{sl}^3 = label_{sl}^3(\mathcal{A}, \mathcal{P})$; the rest follows from $\ell_{sl}^1(q) \geq \ell_{sl}^2(q) \geq \ell_{sl}^3(q)$ proved above. Consider some set of states $V \subseteq Q[\mathcal{A}]$ and the set $V' = \min(\lfloor V \rfloor_{sl})$. We can estimate the symmetric difference of the languages of the original and the reduced automaton as

$$\begin{aligned}
L(\mathcal{A}) \triangle L(reduce_{sl}(\mathcal{A}, V)) & \\
&= L(\mathcal{A}) \triangle L(reduce_{sl}(\mathcal{A}, V')) \quad \{\text{def. of } \lfloor \cdot \rfloor_{sl}\} \\
&= L(reduce_{sl}(\mathcal{A}, V')) \setminus L(\mathcal{A}) \\
&\quad \cup \{L(\mathcal{A}) \setminus L(reduce_{sl}(\mathcal{A}, V'))\} \\
&\subseteq \bigcup_{q \in V'} L_{\mathcal{A}}^b(q) \cdot \Sigma^* \setminus \bigcup_{q \in V'} L_{\mathcal{A}}^b(q) \cdot L_{\mathcal{A}}(q) \\
&\quad \{\text{def. of } reduce_{sl}\} \tag{7}
\end{aligned}$$

The last inclusion holds because $sl(\mathcal{A}, V)$ adds self-loops to the states in V , so the newly accepted words are for sure those that traverse through V , and they are for sure not those that could be accepted by going through V before the reduction (but they could be accepted without touching V , hence the inclusion). We can estimate the probabilistic distance of \mathcal{A} and $reduce_{sl}(\mathcal{A}, V)$ as

$$\begin{aligned}
d_{\mathcal{P}}(\mathcal{A}, reduce_{sl}(\mathcal{A}, V)) & \\
&\leq \Pr_{\mathcal{P}} \left(\bigcup_{q \in V'} L_{\mathcal{A}}^b(q) \cdot \Sigma^* \setminus \bigcup_{q \in V'} L_{\mathcal{A}}^b(q) \cdot L_{\mathcal{A}}(q) \right) \quad \{\text{(7)}\} \\
&\leq \Pr_{\mathcal{P}} \left(\bigcup_{q \in V'} \left(L_{\mathcal{A}}^b(q) \cdot \Sigma^* \setminus L_{\mathcal{A}}^b(q) \cdot L_{\mathcal{A}}(q) \right) \right) \\
&\quad \{\text{properties of union and set difference}\} \\
&\leq \sum_{q \in V'} \Pr_{\mathcal{P}} \left(L_{\mathcal{A}}^b(q) \cdot \Sigma^* \setminus L_{\mathcal{A}}^b(q) \cdot L_{\mathcal{A}}(q) \right) \\
&\quad \{\text{union bound}\} \\
&= \sum_{q \in V'} \left(\Pr_{\mathcal{P}} \left(L_{\mathcal{A}}^b(q) \cdot \Sigma^* \right) - \Pr_{\mathcal{P}} \left(L_{\mathcal{A}}^b(q) \cdot L_{\mathcal{A}}(q) \right) \right) \\
&\quad \{\text{prop. of Pr and the fact that } L_{\mathcal{A}}^b(q) \cdot L_{\mathcal{A}}(q) \subseteq L_{\mathcal{A}}^b(q) \cdot \Sigma^*\} \\
&= \sum \{ \ell_{sl}^3(q) \mid q \in \min(\lfloor V \rfloor_{sl}) \} \\
&\quad \{\text{def. of } \ell_{sl}^3 \text{ and } V'\} \\
&= error_{sl}(\mathcal{A}, V, \ell_{sl}^3). \tag{8}
\end{aligned}$$

- **C1(b)**: $|reduce_{sl}(\mathcal{A}, Q[\mathcal{A}])| \leq 1$ because, from the definition, $|reduce_{sl}(\mathcal{A}, Q[\mathcal{A}])| = |trim(sl(\mathcal{A}, Q[\mathcal{A}]))| \leq 1$.
- **C1(c)**: $reduce_{sl}(\mathcal{A}, \emptyset) = \mathcal{A}$ since

$$reduce_{sl}(\mathcal{A}, \emptyset) = trim(sl(\mathcal{A}, \emptyset)) = \mathcal{A}$$

(we assume that \mathcal{A} is trimmed at the input). \square

5 Reduction of NFAs in Network Intrusion Detection Systems

We have implemented our approach in a Python prototype named APPREAL (APProximate REDuction of Automata and Languages)⁶ and evaluated it on the use case of network intrusion detection using SNORT [48], a popular open source NIDS. The version of APPREAL used for the evaluation in the current paper is available as an artifact [52] for the TACAS'18 artifact virtual machine [20].

5.1 Network Traffic Model

The reduction we describe in this paper is driven by a probabilistic model representing a distribution over the words from Σ^* , and the formal guarantees are also wrt this model. We use *learning* to obtain a model of network traffic over the 8-bit ASCII alphabet at a given network point. Our model is created from several gigabytes of network traffic from a measuring point of the CESNET Internet provider connected to a 100 Gbps backbone link (unfortunately, we cannot provide the traffic dump since it may contain sensitive data).

Learning a PA representing the network traffic faithfully is hard. The PA cannot be too specific—although the number of different packets that can occur is finite, it is still extremely large (a conservative estimate assuming the most common scenario Ethernet/IPv4/TCP would still yield a number over $2^{10,000}$). If we assigned non-zero probabilities only to the packets from the dump (which are less than 2^{20}), the obtained model would completely ignore virtually all packets that might appear on the network, and, moreover, the model would also be very large (millions of states), making it difficult to use in our algorithms. A generalization of the obtained traffic is therefore needed.

A natural solution is to exploit results from the area of PA learning, such as [10, 49]. Indeed, we experimented with the use of ALERGIA [10], a learning algorithm that constructs a PA from a prefix tree (where edges are labelled with multiplicities) by merging nodes that are “similar.” The automata that we obtained were, however, *too* general. In particular, the constructed automata destroyed the structure of network protocols—the merging was too permissive and the generalization merged distant states, which introduced loops over a very large substructure in the automaton (such a case usually does not correspond to the design of network protocols). As a result, the obtained PA more or less represented the Poisson distribution, having essentially no value for us.

⁶ <https://github.com/vhavlena/appreal/tree/tacas18>

In §5.2, we focus on the detection of malicious traffic transmitted over HTTP. We take advantage of this fact and create a PA representing the traffic while taking into account the structure of HTTP. We start by manually creating a DFA that represents the high-level structure of HTTP. Then, we proceed by feeding 34,191 HTTP packets from our sample into the DFA, at the same time taking notes about how many times every state is reached and how many times every transition is taken. The resulting PA \mathcal{P}_{HTTP} (of 52 states) is then constructed from the DFA and the labels in the obvious way.

The described method yields automata that are much better than those obtained using ALERGIA in our experiments. A disadvantage of the method is that it is only semi-automatic—the basic DFA needed to be provided by an expert. We have yet to find an algorithm that would suit our needs for learning more general network traffic.

5.2 Evaluation

We start this section by introducing the experimental setting, namely, the integration of our reduction techniques into the tool chain implementing efficient regex matching, the concrete settings of APPREAL, and the evaluation environment. Afterwards, we discuss the results evaluating the quality of the obtained approximate reductions as well as of the provided error bounds. Finally, we present the performance of our approach and discuss its key aspects. We selected the most interesting results demonstrating the potential as well as the limitations of our approach.

General setting. SNORT detects malicious network traffic based on *rules* that contain *conditions*. The conditions take into consideration, among others, network addresses, ports, or Perl compatible regular expressions (PCREs) that the packet payload should match. In our evaluation, we select a subset of SNORT rules, extract the PCREs from them, and use NETBENCH [43] to transform them into a single NFA \mathcal{A} . Before applying APPREAL, we use the state-of-the-art NFA reduction tool REDUCE [36] to reduce \mathcal{A} . REDUCE performs a language-preserving reduction of \mathcal{A} using advanced variants of simulation [35] (in the experiment reported in Table 3, we skip the use of REDUCE at this step as discussed later in the performance evaluation). The automaton \mathcal{A}^{RED} obtained as the result of REDUCE is the input of APPREAL, which performs one of the approximate reductions from §4 wrt the traffic model \mathcal{P}_{HTTP} , yielding \mathcal{A}^{APP} . After the approximate reduction, we, one more time, use REDUCE and obtain the result \mathcal{A}' .

Settings of APPREAL. In the use case of NIDS pre-filtering, it may be important to never introduce a false negative, i.e., to never drop a malicious packet. Therefore, we focus our evaluation on the *self-loop reduction* (§4.4).

In particular, we use the state labelling function $label_{sl}^2$, since it provides a good trade-off between the precision and the computational demands (recall that the computation of $label_{sl}^2$ can exploit the “tentacle” structure of the NFAs we work with). We give more attention to the *size-driven reduction* (§4.1) since, in our setting, a bound on the available FPGA resources is typically given and the task is to create an NFA with the smallest error that fits inside. The order $\preceq_{\mathcal{A}, \ell_{sl}^2}$ over states used in §4.1 and §4.2 is defined as $s \preceq_{\mathcal{A}, \ell_{sl}^2} s' \Leftrightarrow \ell_{sl}^2(s) \leq \ell_{sl}^2(s')$.

Evaluation environment. All experiments ran on a 64-bit LINUX DEBIAN workstation with the Intel Core(TM) i5-661 CPU running at 3.33 GHz with 16 GiB of RAM.

Description of tables. In the caption of every table, we provide the name of the input file (in the directory `regexps/tacas18/` of the repository of APPREAL) with the selection of SNORT regexes used in the particular experiment, together with the type of the reduction (size- or error-driven). All reductions are over-approximating (self-loop reduction). We further provide the size of the input automaton $|\mathcal{A}|$, the size after the initial processing by REDUCE ($|\mathcal{A}^{\text{RED}}|$), and the time of this reduction ($time(\text{REDUCE})$). Finally, we list the times of computing the state labelling $label_{sl}^2$ on \mathcal{A}^{RED} ($time(label_{sl}^2)$), the exact probabilistic distance ($time(\text{Exact})$), and also the number of *look-up tables* ($LUTs(\mathcal{A}^{\text{RED}})$) consumed on the targeted FPGA (Xilinx Virtex 7 H580T) when \mathcal{A}^{RED} was synthesized (more on this in §5.3). The meaning of the columns in the tables is the following:

k/ϵ is the parameter of the reduction. In particular, k is used for the size-driven reduction and denotes the desired reduction ratio $k = \frac{n}{|\mathcal{A}^{\text{RED}}|}$ for an input NFA \mathcal{A}^{RED} and the desired size of the output n . On the other hand, ϵ is the desired maximum error on the output for the error-driven reduction.

$|\mathcal{A}^{\text{APP}}|$ shows the number of states of the automaton \mathcal{A}^{APP} after the reduction by APPREAL and the time the reduction took (we omit it when it is not interesting).

$|\mathcal{A}'|$ contains the number of states of the NFA \mathcal{A}' obtained after applying REDUCE on \mathcal{A}^{APP} and the time used by REDUCE at this step (omitted when not interesting).

Error bound shows the estimation of the error of \mathcal{A}' as determined by the reduction itself, i.e., it is the probabilistic distance computed by the corresponding function *error* from §4.

Exact error contains the values of $d_{\mathcal{P}_{HTTP}}(\mathcal{A}, \mathcal{A}')$ that we computed *after* the reduction in order to evaluate the precision of the result given in **Error bound**. The computation of this value is very expensive ($time(\text{Exact})$) since it inherently requires determination of the whole automaton \mathcal{A} . We do not provide it in Table 3 (presenting the results for the automaton \mathcal{A}_{bd} with 1,352 states) be-

Table 1: Results for the `http-malicious regex`, $|\mathcal{A}_{\text{mal}}| = 249$, $|\mathcal{A}_{\text{mal}}^{\text{RED}}| = 98$, $\text{time}(\text{REDUCE}) = 3.5$ s, $\text{time}(\text{label}_{\text{sl}}^2) = 38.7$ s, $\text{time}(\text{Exact}) = 3.8\text{--}6.5$ s, and $\text{LUTs}(\mathcal{A}_{\text{mal}}^{\text{RED}}) = 382$.

(a) size-driven reduction							(b) error-driven reduction					
k	$ \mathcal{A}_{\text{mal}}^{\text{APP}} $	$ \mathcal{A}'_{\text{mal}} $	Error bound	Exact error	Traffic error	LUTs	ϵ	$ \mathcal{A}_{\text{mal}}^{\text{APP}} $	$ \mathcal{A}'_{\text{mal}} $	Error bound	Exact error	Traffic error
0.1	9 (0.65 s)	9 (0.4 s)	0.0704	0.0704	0.0685	—	0.08	3	3	0.0724	0.0724	0.0720
0.2	19 (0.66 s)	19 (0.5 s)	0.0677	0.0677	0.0648	—	0.07	4	4	0.0700	0.0700	0.0683
0.3	29 (0.69 s)	26 (0.9 s)	0.0279	0.0278	0.0598	154	0.04	35	32	0.0267	0.0212	0.0036
0.4	39 (0.68 s)	36 (1.1 s)	0.0032	0.0032	0.0008	—	0.02	36	33	0.0105	0.0096	0.0032
0.5	49 (0.68 s)	44 (1.4 s)	2.8e-05	2.8e-05	4.1e-06	—	0.001	41	38	0.0005	0.0005	0.0003
0.6	58 (0.69 s)	49 (1.7 s)	8.7e-08	8.7e-08	0.0	224	1e-04	47	41	7.7e-05	7.7e-05	1.2e-05
0.8	78 (0.69 s)	75 (2.7 s)	2.4e-17	2.4e-17	0.0	297	1e-05	51	47	6.6e-06	6.6e-06	0.0

cause the determinisation ran out of memory (the step is not required in the reduction process).

Traffic error shows the error that we obtained when compared \mathcal{A}' with \mathcal{A} on an HTTP traffic sample, in particular the ratio of packets misclassified by \mathcal{A}' to the total number of packets in the sample (242,468). Comparing **Exact error** with **Traffic error** gives us a feedback about the fidelity of the traffic model $\mathcal{P}_{\text{HTTP}}$.

We note that there are no guarantees on the relationship between **Exact error** and **Traffic error**.

LUTs is the number of LUTs consumed by \mathcal{A}' when synthesized into the target FPGA. Hardware synthesis is a costly step, therefore we provide this value only for selected interesting NFAs.

Table 2: Results for the `http-attacks regex`, size-driven reduction, $|\mathcal{A}_{\text{att}}| = 142$, $|\mathcal{A}_{\text{att}}^{\text{RED}}| = 112$, $\text{time}(\text{REDUCE}) = 7.9$ s, $\text{time}(\text{label}_{\text{sl}}^2) = 28.3$ min, $\text{time}(\text{Exact}) = 14.0\text{--}16.4$ min.

k	$ \mathcal{A}_{\text{att}}^{\text{APP}} $	$ \mathcal{A}'_{\text{att}} $	Error bound	Exact error	Traffic error
0.1	11 (1.1s)	5 (0.4s)	1.0	0.9972	0.9957
0.2	22 (1.1s)	14 (0.6s)	1.0	0.8341	0.2313
0.3	33 (1.1s)	24 (0.7s)	0.081	0.0770	0.0067
0.4	44 (1.1s)	37 (1.6s)	0.0005	0.0005	0.0010
0.5	56 (1.1s)	49 (1.2s)	3.3e-06	3.3e-06	0.0010
0.6	67 (1.1s)	61 (1.9s)	1.2e-09	1.2e-09	8.7e-05
0.7	78 (1.1s)	72 (2.4s)	4.8e-12	4.8e-12	1.2e-05
0.9	100 (1.1s)	93 (4.7s)	3.7e-16	1.1e-15	0.0

5.2.1 Approximation errors

Table 1 presents the results of the self-loop reduction for the NFA \mathcal{A}_{mal} describing regexes from `http-malicious`. We can observe that the differences between the upper bounds on the probabilistic distance and its real value are negligible (typically in the order of 10^{-4} or less). We can also see that the probabilistic distance agrees with the traffic error. This indicates a good quality of the traffic model employed in the reduction process. Further, we can see that our approach can provide useful trade-offs between the reduction error and the reduction factor. Finally, Table 1b shows that a significant reduction is obtained when the error threshold ϵ is increased from 0.04 to 0.07.

Table 2 presents the results of the size-driven self-loop reduction for NFA \mathcal{A}_{att} describing `http-attacks` regexes. We can observe that the error bounds provide again a very good approximation of the real probabilistic distance. On the other hand, the difference between the probabilistic distance and the traffic error is larger than that for \mathcal{A}_{mal} . Since all experiments use the same probabilistic automaton and the same traffic, this discrepancy is accounted to the different set of packets that are incorrectly accepted by $\mathcal{A}_{\text{att}}^{\text{RED}}$. If the probability of these packets is adequately captured in the traffic model, the difference between the

distance and the traffic error is small and vice versa. This also explains an even larger difference in Table 3 (presenting the results for \mathcal{A}_{bd} constructed from `http-backdoor` regexes) for $k \in \langle 0.2, 0.4 \rangle$. Here, the traffic error is very small and caused by a small set of packets (approx. 70), whose probability is not correctly captured in the traffic model. Despite this problem, the results clearly show that our approach still provides significant reductions while keeping the traffic error small: about a 5-fold reduction is obtained for the traffic error 0.03 % and a 10-fold reduction is obtained for the traffic error 6.3 %. We discuss the practical impact of such a reduction in §5.3.

5.2.2 Performance of the approximate reduction

In all our experiments (Tables 1–3), we can observe that the most time-consuming step of the reduction process is the computation of state labellings (it takes at least 90 % of the total time). The crucial observation is that the structure of the NFAs fundamentally affects the performance of this step. Although after REDUCE, the size of \mathcal{A}_{mal} is very similar to the size of \mathcal{A}_{att} , computing $\text{label}_{\text{sl}}^2$ takes more time (28.3 min vs. 38.7 s). The key reason behind this slowdown is the determinisation

(or alternatively disambiguation) process required by the product construction underlying the state labelling computation (cf. §4.4). For \mathcal{A}_{att} , the process results in a significantly larger product when compared to the product for \mathcal{A}_{mal} . The size of the product directly determines the time and space complexity of solving the linear equation system required for computing the state labelling.

As explained in §4, the computation of the state labelling $\text{label}_{s_l}^2$ can exploit the “tentacle” structure of the NFAs appearing in NIDSes and thus can be done component-wise. On the other hand, our experiments reveal that the use of REDUCE typically breaks this structure and thus the component-wise computation cannot be effectively used. For the NFA \mathcal{A}_{mal} , this behaviour does not have any major performance impact as the determinisation leads to a moderate-sized automaton and the state labelling computation takes less than 40 s. On the other hand, this behaviour has a dramatic effect for the NFA \mathcal{A}_{att} . By disabling the initial application of REDUCE and thus preserving the original structure of \mathcal{A}_{att} , we were able to speed up the state label computation from 28.3 min to 1.5 min. Note that other steps of the approximate reduction took a similar time as before disabling REDUCE and also that the trade-offs between the error and the reduction factor were similar. Surprisingly, disabling REDUCE caused that the computation of the exact probabilistic distance became computationally infeasible because the determinisation ran out of memory.

Due to the size of the NFA \mathcal{A}_{bd} , the impact of disabling the initial application of REDUCE is even more fundamental. In particular, computing the state labelling took only 19.9 min, in contrast to running out of memory when the REDUCE is applied in the first step (therefore, the input automaton is not processed by REDUCE in Table 3; we still give the number of LUTs of its reduced version for comparison, though). Note that the size of \mathcal{A}_{bd} also slows down other reduction steps (the greedy algorithm and the final REDUCE reduction). We can, however, clearly see that computing the state labelling is still the most time-consuming step of the process.

5.3 The Real Impact in an FPGA-Accelerated NIDS

To demonstrate the practical usefulness and impact of the proposed approximation techniques, we employ the reduced automata in a real use case from the area of HW-accelerated deep packet inspection. We consider the framework of [34] implementing a high-speed NIDS pre-filter in an FPGA. The crucial challenge is to obtain a pre-filter with a sufficiently small false positive rate (and no false negatives), while being able to handle the traffic of current networks operating on 100 Gbps and beyond. The implementation of NFAs performing regex matching in FPGAs uses two types of HW

Table 3: Results for `http-backdoor`, size-driven reduction, $|\mathcal{A}_{\text{bd}}| = 1,352$, $\text{time}(\text{label}_{s_l}^2) = 19.9$ min, $\text{LUTs}(\mathcal{A}_{\text{bd}}^{\text{RED}}) = 2,266$.

k	$ \mathcal{A}_{\text{bd}}^{\text{APP}} $	$ \mathcal{A}'_{\text{bd}} $	Error bound	Traffic error	LUTs
0.1	135 (1.2 m)	8 (2.6 s)	1.0	0.997	202
0.2	270 (1.2 m)	111 (5.2 s)	0.0012	0.0631	579
0.3	405 (1.2 m)	233 (9.8 s)	3.4e-08	0.0003	894
0.4	540 (1.3 m)	351 (21.7 s)	1.0e-12	0.0003	1,063
0.5	676 (1.3 m)	473 (41.8 s)	1.2e-17	0.0	1,249
0.7	946 (1.4 m)	739 (2.1 m)	8.3e-30	0.0	1,735
0.9	1216 (1.5 m)	983 (5.6 m)	1.3e-52	0.0	2,033

resources: LUTs, which are used to build the combinational circuit representing the NFA transition function, and flip-flops, representing NFA states. In our use case, we omit the analysis of flip-flop consumption because it is always dominated by the LUT consumption.

In our setting, the amount of resources available for the FPGA-based regex matching engine is 15,000 LUTs and the frequency of the engine is 200 MHz using a 32-bit-wide data path. As explained in [34], the engine containing a single unit (i.e. the single NFA implementation) can achieve the throughput of 6.4 Gbps ($200 \text{ MHz} \times 32 \text{ b}$). Therefore, 16 units are required for the desired link speed of 100 Gbps and 63 units are needed to handle 400 Gbps. With the given amount of LUTs, the size of a single NFA is thus bounded by 937 LUTs ($15,000/16$) for 100 Gbps and 238 LUTs for 400 Gbps, respectively. These bounds directly limit the complexity of regexes the engine can handle.

We now analyse the resource consumption of the matching engine for two automata, `http-backdoor` ($\mathcal{A}_{\text{bd}}^{\text{RED}}$) and `http-malicious` ($\mathcal{A}_{\text{mal}}^{\text{RED}}$), and evaluate the impact of the reduction techniques. Recall that the automata represent two important sets of know network attacks from SNORT [48].

- **100 Gbps:** For this speed, $\mathcal{A}_{\text{mal}}^{\text{RED}}$ can be used without any approximate reduction as it is small enough (it has 382 LUTs) to fit in the available space. On the other hand, $\mathcal{A}_{\text{bd}}^{\text{RED}}$ without the approximate reduction is way too large to fit (it has 2,266 LUTs and thus at most 6 units fit inside the available space, yielding the throughput of only 38.4 Gbps, which is unacceptable). The column **LUTs** in Table 3 shows that using our framework, we are able to reduce $\mathcal{A}_{\text{bd}}^{\text{RED}}$ such that it uses 894 LUTs (for $k = 0.3$), and so all of the 16 needed units fit into the FPGA, yielding the throughput over 100 Gbps and the theoretical error bound of a false positive $\leq 3.4 \times 10^{-8}$ wrt the network traffic model $\mathcal{P}_{\text{HTTP}}$.
- **400 Gbps:** Regex matching at this speed is extremely challenging. In the case of $\mathcal{A}_{\text{bd}}^{\text{RED}}$, the reduction $k = 0.1$ is required to fit 63 units in the available space. As such a reduction has error bound almost 1, this solution is

not useful due to a prohibitively high false positive rate. The situation is better for $\mathcal{A}_{\text{mal}}^{\text{RED}}$. In the exact version, at most 39 units can fit inside the FPGA with the maximum throughput of 249.6 Gbps. On the other hand, when using our reduced automata, we are able to place 63 units into the FPGA, each of the size 224 LUTs ($k = 0.6$), and achieve a throughput of over 400 Gbps with the theoretical error bound of a false positive $\leq 8.7 \times 10^{-8}$ wrt the model $\mathcal{P}_{\text{HTTP}}$.

6 Conclusion

We have proposed a novel approach for approximate reduction of NFAs used in network traffic filtering. Our approach is based on a proposal of a probabilistic distance of the original and reduced automaton using a probabilistic model of the input network traffic, which characterizes the significance of particular packets. We characterized the computational complexity of approximate reductions based on the described distance and proposed a sequence of heuristics allowing one to perform the approximate reduction in an efficient way. Our experimental results are quite encouraging and show that we can often achieve a very significant reduction for a negligible loss of precision. We showed that using our approach, FPGA-accelerated network filtering on large traffic speeds can be applied on regexes of malicious traffic where it could not be applied before.

In the future, we plan to investigate other approximate reductions of the NFAs, maybe using some variant of abstraction from abstract regular model checking [7], adapted for the given probabilistic setting. Another important issue for the future is to develop better ways of learning a suitable probabilistic model of the input traffic.

Acknowledgements We thank Jan Kořenek, Vlastimil Kořaň, and Denis Matoušek for their help with translating regexes into automata and synthesis of FPGA designs, and Martin Žádník for providing us with the backbone network traffic. We thank Stefan Kiefer for helping us proving the **PSPACE** part of Lemma 1 and Petr Peringer for testing our artifact. We also thank the anonymous reviewers for their useful comments, which improved the quality of the paper. The work on this paper was supported by the Czech Science Foundation project 16-24707Y, the IT4IXS: IT4Innovations Excellence in Science project (LQ1602), and the FIT BUT internal project FIT-S-17-4014.

References

1. Angluin, D.: Learning regular sets from queries and counterexamples. *Inf. Comput.* **75**(2), 87–106 (1987). DOI 10.1016/0890-5401(87)90052-6. URL [http://dx.doi.org/10.1016/0890-5401\(87\)90052-6](http://dx.doi.org/10.1016/0890-5401(87)90052-6)
2. Baier, C., Kiefer, S., Klein, J., Klüppelholz, S., Müller, D., Worrell, J.: Markov chains and unambiguous Büchi automata. In: *CAV'16*, pp. 23–42. Springer (2016)
3. Becchi, M., Crowley, P.: A hybrid finite automaton for practical deep packet inspection. In: *CoNEXT'07*, p. 1. ACM (2007)
4. Becchi, M., Crowley, P.: An improved algorithm to accelerate regular expression evaluation. In: *ANCS'07*, pp. 145–154. ACM (2007)
5. Becchi, M., Wiseman, C., Crowley, P.: Evaluating regular expression matching engines on network and general purpose processors. In: *Proceedings of the 5th ACM/IEEE Symposium on Architectures for Networking and Communications Systems, ANCS '09*, pp. 30–39. ACM (2009)
6. Benedikt, M., Lenhardt, R., Worrell, J.: Model checking Markov chains against unambiguous Büchi automata. *CoRR abs/1405.4560v2* (2016)
7. Bouajjani, A., Habermehl, P., Rogalewicz, A., Vojnar, T.: Abstract regular (tree) model checking. *STTT* **14**(2), 167–191 (2012)
8. Brodie, B.C., Taylor, D.E., Cytron, R.K.: A scalable architecture for high-throughput regular-expression pattern matching. In: *ISCA'06*, pp. 191–202. IEEE Computer Society (2006)
9. Bustan, D., Grumberg, O.: Simulation-based minimization. *ACM Trans. Comput. Log.* **4**(2), 181–206 (2003)
10. Carrasco, R.C., Oncina, J.: Learning stochastic regular grammars by means of a state merging method. In: *Proceedings of the Second International Colloquium on Grammatical Inference and Applications, ICGI '94*, pp. 139–152. Springer-Verlag (1994)
11. Champarnaud, J., Coulon, F.: NFA reduction algorithms by means of regular inequalities. *Theor. Comput. Sci.* **327**(3), 241–253 (2004)
12. Clark, C.R., Schimmel, D.E.: Efficient reconfigurable logic circuits for matching complex network intrusion detection patterns. In: *FPL'03, Lecture Notes in Computer Science*, vol. 2778, pp. 956–959. Springer (2003)
13. Clemente, L.: Büchi automata can have smaller quotients. In: *ICALP'11, Lecture Notes in Computer Science*, vol. 6756, pp. 258–270. Springer (2011)
14. Csanky, L.: Fast parallel matrix inversion algorithms. In: *16th Annual Symposium on Foundations of Computer Science*, pp. 11–12 (1975). DOI 10.1109/SFCS.1975.14
15. Deza, M.M., Deza, E.: *Encyclopedia of Distances*. Springer (2009)
16. Etesami, K.: A hierarchy of polynomial-time computable simulations for automata. In: *CONCUR 2002 - Concurrency Theory, 13th International Conference, Brno, Czech Republic, August 20-23, 2002, Proceedings, Lecture Notes in Computer Science*, vol. 2421, pp. 131–144. Springer (2002)
17. Fortune, S., Wyllie, J.: Parallelism in random access machines. In: *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing, STOC '78*, pp. 114–118. ACM, New York, NY, USA (1978). DOI 10.1145/800133.804339. URL <http://doi.acm.org/10.1145/800133.804339>
18. Gange, G., Ganty, P., Stuckey, P.J.: Fixing the state budget: Approximation of regular languages with small DFAs. In: *ATVA'17, Lecture Notes in Computer Science*, vol. 10482, pp. 67–83. Springer (2017)
19. Gawrychowski, P., Jez, A.: Hyper-minimisation made efficient. In: *MFC'S'09, Lecture Notes in Computer Science*, vol. 5734, pp. 356–368. Springer (2009)
20. Hartmanns, A., Wendler, P.: TACAS 2018 artifact evaluation VM. In: *Figshare* (2018). <https://doi.org/10.6084/m9.figshare.5896615>
21. Hogben, L.: *Handbook of Linear Algebra*, 2nd edn. CRC Press (2013)
22. Hopcroft, J.E.: An $N \log N$ algorithm for minimizing states in a finite automaton. *Tech. rep.* (1971)
23. Hutchings, B.L., Franklin, R., Carver, D.: Assisting network intrusion detection with reconfigurable hardware. In: *FCCM'02*, pp. 111–120. IEEE Computer Society (2002)
24. Jiang, T., Ravikumar, B.: Minimal NFA problems are hard. *SIAM Journal on Computing* **22**(6), 1117–1141 (1993)

25. Kaštil, J., Kořenek, J., Lengál, O.: Methodology for fast pattern matching by deterministic finite automaton with perfect hashing. In: 2009 12th Euromicro Conference on Digital System Design, Architectures, Methods and Tools, pp. 823–829 (2009)
26. Kořenek, J., Kobierský, P.: Intrusion detection system intended for multigigabit networks. In: 2007 IEEE Design and Diagnostics of Electronic Circuits and Systems, pp. 1–4 (2007)
27. Kumar, S., Chandrasekaran, B., Turner, J.S., Varghese, G.: Curing regular expressions matching algorithms from insomnia, amnesia, and acalculia. In: ANCS'07, pp. 155–164. ACM (2007)
28. Kumar, S., Dharmapurikar, S., Yu, F., Crowley, P., Turner, J.S.: Algorithms to accelerate multiple regular expressions matching for deep packet inspection. In: SIGCOMM'06, pp. 339–350. ACM (2006)
29. Kumar, S., Turner, J.S., Williams, J.: Advanced algorithms for fast and scalable deep packet inspection. In: ANCS'06, pp. 81–92. ACM (2006)
30. Liu, C., Wu, J.: Fast deep packet inspection with a dual finite automata. IEEE Trans. Computers **62**(2), 310–321 (2013)
31. Luchaup, D., De Carli, L., Jha, S., Bach, E.: Deep packet inspection with DFA-trees and parametrized language overapproximation. In: INFOCOM'14, pp. 531–539. IEEE (2014)
32. Malcher, A.: Minimizing finite automata is computationally hard. Theor. Comput. Sci. **327**(3), 375–390 (2004)
33. Maletti, A., Quernheim, D.: Optimal hyper-minimization. CoRR **abs/1104.3007** (2011)
34. Matoušek, D., Kořenek, J., Puš, V.: High-speed regular expression matching with pipelined automata. In: 2016 International Conference on Field-Programmable Technology (FPT), pp. 93–100 (2016)
35. Mayr, R., Clemente, L.: Advanced automata minimization. In: POPL'13, pp. 63–74. ACM Trans. Comput. Log. (2013)
36. Mayr, R., et al.: Reduce: A tool for minimizing non-deterministic finite-word and Büchi automata. <http://languageinclusion.org/doku.php?id=tools>. [Online; accessed 2017-09-30]
37. Mitra, A., Najjar, W.A., Bhuyan, L.N.: Compiling PCRE to FPGA for accelerating SNORT IDS. In: ANCS'07, pp. 127–136. ACM (2007)
38. Mohri, M.: Edit-distance of weighted automata. In: CIAA'02, *Lecture Notes in Computer Science*, vol. 2608, pp. 1–23. Springer (2002)
39. Mohri, M.: A disambiguation algorithm for finite automata and functional transducers. In: CIAA'12, pp. 265–277. Springer (2012)
40. Paige, R., Tarjan, R.E.: Three partition refinement algorithms. SIAM J. Comput. **16**(6), 973–989 (1987)
41. Papadimitriou, C.M.: Computational complexity. Addison-Wesley (1994)
42. Parker, A.J., Yancey, K.B., Yancey, M.P.: Regular language distance and entropy. CoRR **abs/1602.07715** (2016)
43. Puš, V., Tobola, J., Košář, V., Kaštil, J., Kořenek, J.: Netbench: Framework for evaluation of packet processing algorithms. Symposium On Architecture For Networking And Communications Systems pp. 95–96 (2011)
44. Shützenberger, M.: On the definition of a family of automata. Inf. Control (4), 245–270
45. Sidhu, R.P.S., Prasanna, V.K.: Fast regular expression matching using FPGAs. In: FCCM'01, pp. 227–238. IEEE Computer Society (2001)
46. Solodovnikov, V.I.: Upper bounds on the complexity of solving systems of linear equations. Journal of Soviet Mathematics **29**(4), 1482–1501 (1985)
47. Tan, L., Sherwood, T.: A high throughput string matching architecture for intrusion detection and prevention. In: ISCA'05, pp. 112–122. IEEE Computer Society (2005)
48. The Snort Team: Snort. (<http://www.snort.org>)
49. Thollard, F., Clark, A.: Learning stochastic deterministic regular languages. In: G. Paliouras, Y. Sakakibara (eds.) Grammatical Inference: Algorithms and Applications: 7th International Colloquium, ICGI 2004, Athens, Greece, October 11-13, 2004. Proceedings, pp. 248–259. Springer Berlin Heidelberg, Berlin, Heidelberg (2004). DOI 10.1007/978-3-540-30195-0_22. URL https://doi.org/10.1007/978-3-540-30195-0_22
50. Vardi, M.Y.: Automatic verification of probabilistic concurrent finite state programs. SFCS '85, pp. 327–338. IEEE
51. Češka, M., Havlena, V., Holík, L., Lengál, O., Vojnar, T.: Approximate reduction of finite automata for high-speed network intrusion detection. In: Proc. of TACAS'18, LNCS, vol. 10806. Springer (2018)
52. Češka, M., Havlena, V., Holík, L., Lengál, O., Vojnar, T.: Approximate reduction of finite automata for high-speed network intrusion detection. In: Figshare (2018). <https://doi.org/10.6084/m9.figshare.5907055>
53. Yu, F., Chen, Z., Diao, Y., Lakshman, T.V., Katz, R.H.: Fast and memory-efficient regular expression matching for deep packet inspection. In: ANCS'06, pp. 93–102. ACM (2006)