REDoS Detection in "Domino" Regular Expressions by Ambiguity Analysis

[1] *Antonina Nepeivoda,* ORCID: *0000-0003-3949-2164 <a_nevod@mail.ru>*
[2] *Yulia Belikova,* ORCID: *0009-0007-7829-1249, <ju.belikova@gmail.com>*
[2] *Kirill Shevchenko,* ORCID: *0009-0007-2868-153X <k.sh3vch3nko@yandex.ru>*
[2] *Mikhail Teriukha,* ORCID: *0009-0005-2825-8171 <misha37a999@yandex.ru>*
[2] *Danila Knyazihin,* ORCID: *0009-0009-6343-6809 <dak151449@gmail.com>*
[2] *Aleksandr Delman,* ORCID: *0009-0009-6885-8429 <adelman2112@gmail.com>*
[2] *Anna Terentyeva,* ORCID: *0009-0006-8547-3959 <mathhyyn@gmail.com>*

[1] *Aylamazyan Program Systems Institute of the Russian Academy of Sciences,*
4a, Petra I st., Veskovo, Pereslavsky District, Yaroslavl Oblast, 152024, Russia.
[2] *Bauman Moscow State Technical University,*
5, 2-nd Baumanskaya, Moscow, 105005, Russia.

Abstract: The Regular Expression Denial of Service (REDoS) problem refers to a time explosion caused by the high computational complexity of matching a string against a regex pattern. This issue is prevalent in popular regex engines, such as PYTHON, JAVASCRIPT, and C++. In this paper, we examine several existing open-source tools for detecting REDoS and identify a class of regexes that can create REDoS situations in popular regex engines but are not detected by these tools. To address this gap, we propose a new approach based on ambiguity analysis, which combines a strong star-normal form test with an analysis of the transformation monoids of Glushkov automata orbits. Our experiments demonstrate that our implementation outperforms the existing tools on regexes with polynomial matching complexity and complex subexpression overlap structures.

Keywords: regular expressions; ambiguity; REDoS; Glushkov automaton; transformation monoid; strong star-normal form.

Выявление REDoS ситуаций в регулярных выражениях структуры «домино»

[1] *А.Н. Непейвода, ORCID: 0000-0003-3949-2164 <a_nevod@mail.ru>,*
[2] *Ю.А. Беликова, ORCID: 0009-0007-7829-1249 <ju.belikova@gmail.com>*
[2] *К.К. Шевченко, ORCID: 0009-0007-2868-153X <k.sh3vch3nko@yandex.ru>*
[2] *М.Р. Терюха, ORCID: 0009-0005-2825-8171 <misha37a999@yandex.ru>*
[2] *Д.П. Князихин, ORCID: 0009-0009-6343-6809 <dak151449@gmail.com>*
[2] *А.Д. Дельман, ORCID: 0009-0009-6885-8429 <adelman2112@gmail.com>*
[2] *А.С. Терентьева, ORCID: 0009-0006-8547-3959 <mathhyyn@gmail.com>*

[1] *Институт программных систем РАН им. А.К. Айламазяна,*
Россия, 152024, Ярославская обл., с. Веськово, ул. Петра I, д. 4а.
[2] *Московский государственный технический университет имени Н.Э. Баумана,*
105005, Россия, Москва, ул. Бауманская 2-я, д. 5/1.

Аннотация. Ситуация отказа в обслуживания регулярных выражений (REDoS) возникает в случае высокой вычислительной сложности сопоставления строки с выражением и встречается во многих библиотеках регулярных выражений таких языков, как PYTHON, JAVASCRIPT, C++. В данной статье рассматривается класс регулярных выражений, которые создают угрозу возникновения REDoS, однако не распознаются как уязвимые рядом существующих программных систем. Предлагается производить оценку степени неоднозначности таких выражений посредством комбинирования проверки на строгую звёздную нормальную форму и анализа трансформационного моноида автомата Глушкова, построенного по входному регулярному выражению. Эксперименты показывают, что данный подход оказывается эффективен при оценке полиномиальных неоднозначностей в регулярных выражениях со сложной структурой перекрытий.

Ключевые слова: регулярные выражения; неоднозначность; REDoS; автомат Глушкова; трансформационный моноид; сильная звёздная нормальная форма.

1. Introduction

Popular regular expression (regex) engines typically use non-deterministic finite automata (NFA) as their internal representation for regexes. This choice is motivated by the flexibility of the NFA concept, which can be extended to support a wider range of regex operations with little effort. For instance, back-references and lookaheads can be easily added to the NFA model. Although, in theory, every string can be matched against a regex in linear time using deterministic finite automata (DFA) conversion, popular regex engines may admit exponential matching time due to a phenomenon called "catastrophic backtracking".

This phenomenon occurs only for a specific class of regular expressions. For example, consider the regex `(a|b)*a`, which is non-deterministic due to the unavoidable non-determinism in the transition to the last occurrence of the letter `a`. However, every string has a unique parsing tree with respect to this regex. In contrast, the regex `(a*b*)*` has an infinite number of accepting parsing trees for any given string, as inner Kleene stars can degenerate to the empty word, causing a combinatorial explosion of parse paths. Intuitively, the latter regex can be considered "bad", while the former is considered "good".

Matching against "bad" regexes can yield a situation called a Regular Expression Denial of Service (REDoS), when the matching time grows super-linearly and can cause performance issues in, for

110

instance, a web service that uses such a regex to parse user input. To avoid these situations, it is essential to detect unsafe regexes and replace them with safe equivalents.

The number of research papers mentioning the REDoS problem has increased rapidly in the last decade [1]–[7]. Several tools have been developed to detect REDoS, using both static analysis and random search. Some of these tools aim to detect the entire class of extended regexes, while others focus on academic ones. However, for a class of simple regexes, which are not safe in theory, the tools considered either take too long time to process, or give an incorrect answer, falsely witnessing their safety. These regexes usually have overlapping, but not completely coinciding, structure of the expressions under the Kleene stars (being a simple analogue of dominoes in the Post Correspondence Problem). An example of such a regex is `(baa|ab)*(b|ε)(a(ba|a)ba*b)*(aab)*`: the ambiguity occurs both in prefixes `(baa)`^{n} and `(ab)`^{n}, which can be constructed in several ways from primitive "dominoes".

Thus, the two natural research questions arise:

- do the "domino" regexes really contain REDoS situations w.r.t. the modern regex engines?

- if the answer is yes, what methods can deal with such regexes in order to analyse them without blow-up of the analysis time because of the overlaps?

The main contributions of the paper are:

- a method for REDoS situations detection, utilizing properties of non-deterministic finite automata and their transition monoids. This approach is novel, since previous static-analysis-based methods use NFA intersection. For "domino" regexes our method is shown to perform better than the open-source analogues REGEX STATIC ANALYZER [3], RESCUE [5], and REVEALER [2].

- experimental testing of the relevance of the NFA model used and the vulnerabilities found, by investigating real regex engines behaviour on the attack strings.

The method is implemented only for the academic regexes for now. Surprisingly, for this case, the tested open-source tools perform significantly worse on domino tests, especially for polynomial REDoS situations.

The paper is organized as follows. Section 2 contains preliminaries on finite automata, and theoretical concepts that are used further. The proposed REDoS detection method is given in Section 3, preceded by lemmas used for its optimisation. Section 4 discusses relevance of the chosen model with respect to the real regex matching engines, and provides a result of comparative testing of our method and three other open-source REDoS detection tools. We discuss the results of the experiments and the related works in more detail in Section 5. Section 6 concludes the paper.

2. Preliminaries

We denote automata with calligraphic $\mathcal{A}$; states are denoted with the letters q and Q, or with the set of these letters (if an automaton is a result of a closure operation). The empty word is denoted by ε; concrete elements from the input alphabet are denoted with a, b, c, ..., and letter parameters are denoted with γ; ω and η denote word parameters. We use only the basic academic regular expression constructing operations: concatenation (which is omitted in notation), alternation (denoted with |), and Kleene star (denoted with *). If r is a regex, $\mathcal{L}(r)$ denotes its language.

Let us recall basic definitions and describe the finite automata models used in this paper.

2.1 Finite Automata

Definition 1. *A non-deterministic finite automaton (NFA) is a tuple <S, Σ, q$_0$, F, δ>, where:*

- *S is a state set;*

- Σ *is a terminal alphabet;*

- δ *is a set of transitions of the form* $<q_i, (\gamma_i \mid \varepsilon), M_i>$*, where* $q_i \in S$*,* $\gamma_i \in \Sigma$*,* $M_i \in 2^S$*;*

- $q_0 \in S$ *is the initial state;*

- $F \subseteq S$*, is a set of final states.*

Every transition in an NFA maps a pair $<q_i, (\gamma_i \mid \varepsilon)>$ into a set of states, contrary to transitions in a deterministic finite automaton (DFA), which map every pair (q_i, γ_i) (where γ_i is essentially not equal to ε) to a single state. Thus, if a word is parsed by a DFA, the parse trace is always unique (i.e., DFAs are unambiguous); in an NFA, there can be a set of parse traces for a single word. This set can even be infinite in case of NFA with ε-transitions. The notation $q_i \rightarrow^\gamma \dots$ is overloaded to denote either NFA transition $<q_i, \gamma, M_i>$ (written as $q_i \rightarrow^\gamma M_i$) or a transition to a single state belonging to M_i (written as $q_i \rightarrow^\gamma q_j$). Existence of a path from q_i to q_j marked by ω in Σ^* is also denoted by $q_i \rightarrow^\omega q_j$.

An NFA can be transformed into an equivalent DFA using a textbook subset-constructing algorithm `Determinize`, which generates states of the DFA corresponding to the sets of the states of the initial NFA resulted in the transitions along the same input symbols.

The NFA models used in regex engines are primarily based on the classical Thompson construction, which provides an algorithm for transforming a regex into an NFA that recognizes the same language. While the implementation details of the transformation may vary, the experiments presented in Section 4 provide evidence that the Thompson model remains relevant for identifying inefficient regexes with respect to NFA-based parsing engines.

In the following descriptions, we only give details of the constructed NFAs in terms of their states and transitions, without mentioning the alphabet construction.

Definition 2. *Thompson NFA (denoted with* `Thompson(r)`*) is constructed from a regex* `r` *as follows. At any construction step except processing concatenations, the new initial state* q_r *and the new final state* Q_r *are introduced, and the transition set is updated depending on the regex operation.*

- *Every single letter* γ *generates a primitive automaton with the only transition* $q_\gamma \rightarrow^\gamma \{Q_\gamma\}$*.*

- *If* $\mathcal{A}_1$ = `Thompson(r₁)`*,* $\mathcal{A}_2$ =`Thompson(r₂)`*, and* q_i *and* Q_i *are their initial and final states, respectively, then* `Thompson(r₁ | r₂)` *is constructed by merging the* $\mathcal{A}_1$ *and* $\mathcal{A}_2$ *states sets and transitions sets, and introducing the transitions* $q_{alt} \rightarrow^\varepsilon \{q_1, q_2\}$*;* $Q_1 \rightarrow^\varepsilon \{Q_{alt}\}$*;* $Q_2 \rightarrow^\varepsilon \{Q_{alt}\}$*.*

- `Thompson(r₁r₂)` *is again constructed by merging* `Thompson(rᵢ)` *states and transitions sets, and making* q_1 *the initial state,* Q_2 *the final state, with the additional transition* $Q_1 \rightarrow^\varepsilon \{q_2\}$*.*

- `Thompson(r₁*)` *is constructed introducing transitions* $q_* \rightarrow^\varepsilon \{q_1, Q_*\}$*,* $Q_1 \rightarrow^\varepsilon \{q_1, Q_*\}$*.*

The Thompson construction algorithm ensures that any NFA produced by the algorithm has a unique final state and that each state has at most two outgoing and two incoming transition arcs. The uniqueness of the final state implies that the reverse NFA for `Thompson(r)` is exactly `Thompson(rᴿ)`, where r^R is the reverse of the regex `r`. Additionally, all subregex automata can be treated as isolated directed acyclic graphs, which makes the construction easily extensible and decomposable. An example of a Thompson automaton for a regex is shown in *Fig. 1*. The states labels follow the corresponding regex operations given in Definition 2.

Fig. 1. Thompson automaton for `(a|b)*a`

One drawback of the Thompson construction is that it introduces non-deterministic transitions corresponding to alternating operations (i.e., alternatives or Kleene stars), even in the cases when the regex itself imposes no non-determinism (e.g. for the regex `a(a|b)*`, which is a reverse of the regex shown in *Fig. 1*). To avoid the redundant non-determinism, the regex engine RE2 [8] processes such strongly deterministic regexes (also known as 1-unambiguous regexes [9]) constructing another NFA based on the regex structure, but without ε-transitions. This automaton is known as the Glushkov automaton since 1960s, and in the last two decades it attracted considerable interest, shown to be efficient and extensible to construct deterministic parsing engines for a larger class of regexes (such as memory finite automata for the regexes with back-references [10]). The Glushkov automaton is shown in *Fig. 2*.


Fig. 2a. `Thompson(a(a|b)*)` *with colored ε-closures*


Figure 2b. `Glushkov(a(a|b)*)`

The classical Glushkov construction is based on so-called follow-relation on linearised regexes. By construction, every state in the Glushkov automaton except the initial state corresponds to an occurrence of some γ in Σ in the input regex r; conversely, any letter occurrence in the regex r corresponds to exactly one state in `Glushkov(r)`, whose incoming arcs are all marked with γ. Now we can reformulate this property in the terms of Thompson and Glushkov automata.

Proposition 1. *There is a bijection from state set in* `Glushkov(r)` *minus the initial state to state set Q_γ in* `Thompson(r)` *(where Q_γ are final states of the primitive automata reading γ).*

In the paper [11], it was shown that `Glushkov(r)` could also be obtained from `Thompson(r)` merging its ε-closures.

Definition 3. *Given an NFA $\mathcal{A}$ and its state q, ε-closure of q is the maximal set of states reachable from q following only ε-transitions.*

Closure-merging[1] ε-free automaton (denoted with $\mathtt{RemEps}(\mathcal{A})$*) is constructed from* $\mathcal{A}$ *as follows:*

- *its states are ε-closures of the states of $\mathcal{A}$;*
- *if state q_1 belongs to closure C_i, state q_2 belongs to C_j, and there is a transition $q_1 \rightarrow^{\gamma} \{...,$ $q_2, ...\}$ ($\gamma \neq \varepsilon$) in $\mathcal{A}$, then there is a transition $C_i \rightarrow \{..., C_j, ...\}$ in* $\mathtt{RemEps}(\mathcal{A})$.

An example of closure-merging operation is given in *Fig. 1* and *Fig. 2*, the nodes belonging to a closure are highlighted with the same color.

2.2 Transformation Monoid of NFA

Let us consider an automaton with no useless states and ε-transitions. Its transitions over the alphabet Σ and the states set 2^Q form the function $F : \Sigma \times S \rightarrow 2^S$ taking a pair $<\gamma, q_i>$. This function, when curried and specialized in the first argument, becomes $F_\gamma : S \rightarrow 2^S$ (where $\gamma \in \Sigma$). We can form a monoid over the set of such partially specialized functions (transformations) if we continue them on strings as follows: $F_{\omega 1} \circ F_{\omega 2} = F_{\omega 2 \omega 1}$. Then associativity is provided "for free", given associativity of string concatenation, and ε becomes the monoid unit, because $F_\omega \circ F_\varepsilon = F_{\omega \varepsilon} = F_\omega = F_{\varepsilon \omega} = F_\varepsilon \circ F_\omega$ holds. The state transformations are denoted by the corresponding strings ω.

The formal definition is as follows [12].

Definition 4. *Given an ε-free automaton $\mathcal{A}$ over the alphabet Σ, its transformation monoid $\mathcal{M} =$* $\mathtt{TransMonoid}(\mathcal{A})$ *is the monoid of transformations imposed by elements of Σ^* on the states of $\mathcal{A}$.*

The monoid construction does not depend on the choice of the final or initial states of $\mathcal{A}$ (except the condition that all the states are useful, i.e. reachable and producing), thus, instead of classical NFAs, the monoid is based on a labelled transition system. Since the set of functions $S \rightarrow 2^S$ is finite, the transformation monoid of an NFA always contains a finite number of equivalence classes. The pair $<M, R>$, where M is a finite set of lexicographically minimal elements of the equivalence classes and R is a set of simplification rules is considered a standard representation of the transformation monoid. Such a representation for $\mathtt{TransMonoid(Glushkov(}a\,(a|b)^*\mathtt{))}$ is given in *Fig. 3a*, *Fig. 3b*, and *Table 1*. The monoid representation uncovers some useful NFA properties. For example, we can immediately conclude that the words $\mathtt{aa}$ and $\mathtt{ab}$ are synchronizing, since for all q_i, $q_i \rightarrow^{\mathtt{aa}} q_2$, $q_i \rightarrow^{\mathtt{ab}} q_3$, and no other transition is possible.

Fig. 3a. Labelled transition system of NFA

```
bb  → b        aaa  → aa
aab → ab       aba  → aa
baa → ba       bab  → bb
```

Fig. 3b. Rewriting rules of NFA

---

[1] This ε-removal construction differs from the standard textbook ε-removal algorithm, since it changes states, and not only transitions. This strategy allows the algorithm to succeed in conversion from Thompson to Glushkov.

Table 1. Equivalence classes of NFA

	q_0	q_1	q_2	q_3
a	$\{q_1\}$	$\{q_2\}$	$\{q_2\}$	$\{q_2\}$
b	$\{\ \}$	$\{q_3\}$	$\{q_3\}$	$\{q_3\}$
aa	$\{q_2\}$	$\{q_2\}$	$\{q_2\}$	$\{q_2\}$
ab	$\{q_3\}$	$\{q_3\}$	$\{q_3\}$	$\{q_3\}$
ba	$\{\ \}$	$\{q_2\}$	$\{q_2\}$	$\{q_2\}$

2.3. Ambiguity of NFAs and REDoS

Intuitively, the worst-case scenario for backtracking-based matching of a string against a regex r occurs when the matched string has a prefix η_1 with a large set of parse paths, and a suffix η_2 such that $\eta_1\eta_2 \notin \mathcal{L}(r)$. In this case, in order to determine that $\eta_1\eta_2$ is not recognizable by r, a regex engine must backtrack through all the parse variants of η_1. Obviously, we can choose such a suffix η_3 that $\eta_1\eta_3 \in \mathcal{L}(r)$, and $\eta_1\eta_3$ will still have a large number of parse trees (although the regex engine will report a success after finding a first one).

Therefore, worst-case matching time depends on the upper bound on the parse paths in a regex.

In the domain of finite automata, the following definition is used [13], [14].

Definition 5. *A degree of ambiguity of an NFA $\mathcal{A}$ is a worst-case bound on the number of paths recognizing an input string (in a length of the string).*

The ambiguity of NFAs is known to be either a constant, an exponential, or a polynomial [13]. If the ambiguity degree of $\mathcal{A}$ is non-constant, it is said $\mathcal{A}$ has an infinite degree of ambiguity (IDA). A standard acronym for exponential ambiguity degree is EDA.

A minimal EDA-generating regex example is `(a|a)*`. A minimal example of a regex producing IDA but not EDA automaton is `a*a*`. For regexes such that `(a*b*)*`, `Glushkov(r)` is unambiguous, despite `Thompson(r)` is EDA. We can notice that in `Thompson((a*b*)*)`, a special situation occurs: there is a loop inside an ε-closure of a state (i.e., there is at least one Kleene star in a regex iterating over an expression r_E s.t. $\varepsilon \in \mathcal{L}(r_E)$). Further we show that such a case is one of the few possible exceptions when `Thompson(r)` and `Glushkov(r)` have distinct ambiguity degrees.

The following criterion estimates the degree of ambiguity in any NFA.

Theorem 1. • NFA $\mathcal{A}$ *satisfies IDA condition iff there exist states q_1, q_2 in $\mathcal{A}$, and a word ω such that $\mathcal{A}$ contains paths from q_1 and q_2 to themselves, and a path from q_1 to q_2 all accepting the word ω.*

• NFA $\mathcal{A}$ *satisfies EDA condition iff there exists a state q_1 in $\mathcal{A}$, and a word ω such that $\mathcal{A}$ contains two distinct loops from q_1 to itself both accepting the word ω.*

We can also say than if EDA occurs in an NFA, then $\exists\ q_i$, q_j, q_k, where q_j and q_k are distinct, such that there exist words ω_1 and ω_2 such that both q_k and q_j are reachable from q_i following a path reading the word ω_1, and q_i is reachable from both q_k and q_j following a path reading the word ω_2 (see *Fig. 4*).

After the work [9], we use the term "<u>orbit of state q</u>" for the maximal strongly connected component containing q. We assume that orbits are non-trivial, i.e. contain at least one transition. If a state q of $\mathcal{A}$ satisfies EDA criterion for some ω, then all states belonging to its orbit also satisfy EDA. Thus, to check the EDA condition, it is sufficient to check if any state of some strongly connected

component of an NFA satisfies EDA; for the IDA condition, it is sufficient to check if there are two strongly connected components satisfying it.


Fig. 4a. The EDA situation

Fig. 4b. The IDA situation

An approach to the IDA and EDA detection used in the REDoS analysers [3], [4] tests the above criterion constructing single or double intersections of automaton $\mathcal{A}$ with itself. Although the intersection construction can be done in polynomial time on an NFA size, it may lead to large NFAs if there are many crossing components (i.e., matching the same string sets) in the initial NFA.

The IDA criterion can be also reformulated in the terms of transformation monoids.

Proposition 2. *An ε-free automaton $\mathcal{A}$ satisfies IDA iff its transformation monoid contains an equivalence class ω such that for some states q_i, q_j, $q_i \in F_\omega(q_i)$, $q_j \in F_\omega(q_j)$, and $q_j \in F_\omega(q_i)$.*

Using this criterion for an initial NFA "as is" is highly impractical: if the NFA contains non-crossing components, the transformation monoid becomes exponentially huge. However, with some refinements, we observed that the monoid criterion can be applied (and even be fast) in the cases when the intersection criterion is slow. Moreover, Proposition 2 provides explicit construction of a string with the ambiguity, allowing the analysing algorithm to reconstruct the REDoS situation easily. First, take any NFA path from the initial state of $\mathcal{A}$ to q_i, recognizing some prefix η_1. Then pump ω to construct an infix with superlinear number of parse trees, and then take some string η_2 such that any path from q_j recognizing η_2 does not end in a final state of $\mathcal{A}$. The string $\eta_1\omega^n\eta_2$ will force an NFA parsing device to do superlinear backtracking.

If the monoid criterion is applied to the orbit automaton of state q, the REDoS pump can be constructed as well. Just choose some η_1, η_2 such that $q_0 \rightarrow^{\eta_1} q$, and $\forall q_F \in F$ in the condition $q \rightarrow^{\eta_2} q_F$ is not satisfied.

3. Our approach

As a starting point, we prefer to use the Thompson automaton as a preliminary NFA model for a regex since regex matching engines rely on it in their internal algorithms, and experiments in Section 4 demonstrate that the Thompson construction is suitable for analysing real REDoS. However, in order to apply the monoid criterion, we must first eliminate ε-transitions in the regex and ensure that the removal of ε-transitions does not affect the degree of ambiguity.

Definition 6. *A regular expression r is said to be in a star-normal form (SNF) if for each its subexpression $(r')^* \varepsilon \notin \mathcal{L}(r')$.*

Let us say that r *is in the strong star-normal form (SSNF) if it is in SNF and none of its subexpressions* $(r')^*$ *contains an alternation* $r_1 | r_2$, *where* $\varepsilon \in \mathcal{L}(r_1)$ & $\varepsilon \in \mathcal{L}(r_2)$.

The following theorem is the main theoretical result of the paper.

Theorem 2. *If* r *is SSNF, and* Thompson(r) *is infinitely ambiguous, then* Glushkov(r) *is also infinitely ambiguous. Moreover, the overall ambiguity degrees (exponential or polynomial) of* Thompson(r) *and* Glushkov(r) *coincide.*

In order to prove Theorem 2, we use the statement proven in [11] mentioned above: RemEps(Thompson(r)) = Glushkov(r), where RemEps is the transformation described in Section 2, and the following auxiliary proposition.

Proposition 3. *If* a_1 *and* a_2 *are distinct letter occurrences in* r, *and* q_{A1}, *and* q_{A2} *are final states of the elementary Thompson automata for* a_1 *and* a_2, *then* q_{A1} *and* q_{A2} *never belong to a single* ε-closure *of a state in* Thompson(r).

Proof of Proposition 3:

Every final state of the elementary automaton for a single letter has a unique ingoing edge, marked with the given letter. No other construction adds ingoing edges to the final states of the subautomata used in the construction. Thus, the states q_{A1} cannot be reached from q_{A2} along ε-transitions, and vice versa.

Proposition 3 allows us to construct the bijection between final states of the elementary subautomata of Thompson(r) and all the states except the initial one in Glushkov(r), mentioned in Section 2.

Proof of Theorem 2:

Let r be in the strong star-normal form. All the strongly connected components of Thompson(r) and Glushkov(r) correspond to expressions under Kleene stars.

If some strongly connected component of Thompson(r) contains an EDA, then there exists a state q, two distinct states q_1 и q_2 and $\alpha_1, \alpha_2 \in \Sigma \cup \{\varepsilon\}$, words ω_1, ω_2 such that $<q, \alpha_1, q_1> \in \delta$, $<q, \alpha_2, q_2> \in \delta$, $q_1 \rightarrow^{\omega_1} q$, $q_2 \rightarrow^{\omega_2} q$, satisfying $\alpha_1 \omega_1 = \alpha_2 \omega_2$. Let us denote the path from q to itself following through q_1 by P_1, and the similar path following through q_2 by P_2.

If $\alpha_1 \omega_1 = \varepsilon$, then there is an ε-loop from q to itself, which contradicts the SSNF condition. Thus, we can take the first letter of $\alpha_1 \omega_1$ belonging to Σ, say a. Let us consider the final states q_1', q_2' of elementary Thompson automata for a in the paths P_1 и P_2.

If q_1' is not equal to q_2', then their ε-closures are also distinct, which implies the EDA situation in Glushkov(r).

Let q_1' and q_2' coincide. We recall that we chose the states q, q_1, q_2, such that the first edges in the paths P_1 and P_2 outgoing from q (and ingoing in q_1 and q_2), are distinct.

Let $q_1'=q_1$ (or $q_1'=q_2$). The state q_1' has a single ingoing edge, namely the one outgoing from q and marked with a. But $q_1'=q_2'$, and q, being a predecessor of q_2' in the path P_2, must occur in its initial fragment twice, thus, there is a path from q to q recognizing ε. This contradicts the SSNF condition.

Let q_1' to be distinct both from q_1 and q_2, but to coincide with q_2'. Let us again consider the ingoing edge in q_1' marked with a. Let this edge to outgo from some state q_0. Then there are the two distinct paths from q into q_0 reading the empty word, again contradicting SSNF. Thus, EDA in the Thompson automaton leads to EDA in the Glushkov automaton.

Now let Thompson(r) contain IDA, but not an EDA. Then r contains the distinct subexpressions r_1 and r_2 under the Kleene stars, both recognizing a same word ω, such that the states corresponding to r_1 are not reachable from the states generated by r_2. Thus, r_1 and r_2 contain the same letter a with occurrences in the positions i and j, $i<j$, and the state for a_i in Glushkov(r) is not reachable from the state for a_j. Then Glushkov(r) contains an IDA, which is not an EDA. This completes the proof.

117

Thus, it is sufficient to test r for the strong star-normal form property and then, if necessary, continue the ambiguity analysis operating with the Glushkov automaton, having significantly less states. If there are loops in ε-closures, the further analysis is not needed: these loops already produce EDA situations.

Given a state q in $\mathcal{A}$ and its orbit M, an <u>orbit automaton</u> of q is automaton M_q including all states and transitions from M, having q as is the initial state, and whose final states are either final states of $\mathcal{A}$ or states with outgoing transitions outside the orbit M in $\mathcal{A}$.

If we choose one state q_i from each strongly connected component C_i of $\mathcal{A}$, then testing an IDA criterion for `TransMonoid`(M_{qi}) is enough to reveal all EDA situations.

However, in the case of a polynomial IDA, we must test pairs of the strongly connected components (together with the transitions from one component to another) and building a monoid for any such pair-generated NFA is too time- consuming. Thus, we use the following simple necessary condition for the polynomial IDA.

Proposition 4. *Let C_1, C_2 be distinct strongly connected components of $\mathcal{A}$. If $\mathcal{A}$ contains a polynomial IDA within the components, then there exist two states, $q_1 \in C_1$, $q_2 \in C_2$, such that* `Determinize(`$\mathcal{A}$`)` *contains a subset state including both q_1 and q_2. Moreover, such a subset state occurs also in* `Determinize(Reverse(`$\mathcal{A}$`))`.

Although the determinization algorithm is exponentially hard in the worst case, it is known to be fast in most practical cases [16]. Thus, the subset test accelerates candidates search for the polynomial IDA. However, it is not sufficient, which can be shown by analysing regex `(a|b)*(b|c)(a|c)*` whose Thompson automaton contains no IDA.

The pseudocode of the complete algorithm[2] is given in *Fig. 5*. There $\mathcal{A}_{q1+q2}$ includes the orbit automata M_{q1} and M_{q2} of q_1 and q_2, and all states reachable from M_{q1} and reaching M_{q2} together with their transitions. Its initial state coincides with initial state of M_{q1}, and its final states are final states of M_{q2} (ignoring final states of $\mathcal{A}$ belonging either to M_{q1} or to the intermediate states). The condition "c_1 `reaches` c_2" ensures that the component c_2 is reachable from c_1, and they do not coincide. Operator $c[1]$ takes a first state from the component c (since the `Ambiguity.TransMonoid` and determinization tests results do not depend on the choice of the initial state in the orbit automata[3]). Function `SCC`($\mathcal{A}$) returns all strongly connected components of $\mathcal{A}$.

4. Experiments

4.1 Data Set

In order to evaluate the effectiveness of our approach on the "domino" regexes, a dataset of 100 academic regexes was generated. The regexes satisfy the following properties:

- their length and alphabet are small (not more than 50 terms and not more than 5 distinct letters);
- they have iterated elements;
- all are in SSNF.

The first condition allows significant subexpression languages overlap, without blowing up the regex length. However, the test set contains not only complex dominoes, but also regexes with simple ambiguity situations like `b*c(ac|(aa|a)*d)*`. The second condition is necessary for

---

[2] The trial implementation of the method is given on https://github.com/bmstu-iu9/Chipo-Kleene/tree/ambiguity.

[3] Absence of any useless states is guaranteed, because all the states are reachable from each other.

REDoS situations. The third condition mostly excludes the trivial SSNF test, returning EDA value using our method too quickly.

```
if ¬ SSNF(r) then
    return EDA
    A ← Glushkov(r)
end if
C ← SCC(A)
for c ∈ C do
    q0 ← c[1]
    if Ambiguity(TransMonoid(M_q0)) then
        return EDA
    end if
end for
for c1, c2 ∈ C do
    if c1 reaches c2 then
        q1 ← c1[1]
        q2 ← c2[1]
        if SubsetPairs(Determinize(A_{q1+q2}))
        ∩ SubsetPairs(Determinize(Reverse(A_{q1+q2}))) ≠ ∅
        then
            if Ambiguity(TransMonoid(A_{q1+q2})) then
                return IDA
            end if
        end if
    end if
end for
return Safe
```

Fig. 5. The overall algorithm schema

We explored the dependence of the regexes matching time from the input length on the popular engines in PYTHON, JAVASCRIPT, C++, JAVA 8, JAVA 11, GO, and RUST.

In order to detect super-linear dependencies, it is necessary to generate potentially attacking input, for which the string pumping method is used. The attacking input must match a pattern of the three components: a prefix that satisfies the regular expression, a pumping core whose repetition can lead to a rapid increase in the number of parsing paths (i.e., malicious pump), and a suffix whose mismatch leads to catastrophic backtracking.

The results obtained by applying JAVASCRIPT, PYTHON, C++ and JAVA 8 standard regex engines are the same, according to them, the data set contains 34 exponential, 36 polynomial and 30 safe regexes. In addition, the experiments indicated that JAVA 11 standard regex engine handles some polynomial and exponential cases, but when the length of the input data increases significantly, it throws a stack overflow exception, which may be due to the introduction of the local storage of indexes to the regex module in the 11 version of JAVA.

The regexes are safe for GO and RUST engines, which are based on the deterministic structures. Nevertheless, it was noted that there are frequent single outliers in trends when matching strings in GO.

During testing, we observed that polynomial regexes only lead to critical matching times (more than 1 minute) with significant input string lengths (approximately more than 500 characters), while expressions that have exponential matching complexity can reach critical time when parsing even

relatively small input strings. In the simplest case, such a time explosion can be achieved with regexes that have large star nesting or multiple alternatives under a star *q*uantifier. For instance, the PYTHON, JAVASCRIPT, JAVA 8, and C++ regex engines are vulnerable to attacks in the case of the `((a*)*)*` regex, and even the optimized JAVA 11 engine, which successfully handles double star nesting, reaches critical time processing such an expression.

However, more non-trivial cases were encountered in the proposed data set. For example, the regex `b(ab((a|b(a*a)*)a*b*)*|a*aaaa*)*`, when matched against the input of 32 characters that satisfies the pattern with prefix – `b`, pump – `abab`, suffix – `bbd`, achieves the following timings: PYTHON engine – over 3 minutes, JAVA 8 – over 3 minutes, JAVA 11 – 0.80 minutes, C++ – over 3 minutes, JAVASCRIPT – 1.73 minutes.

In general, the REDoS vulnerability degree coincides with the theoretical expectations, taking into account the asymptotic growth of the ambiguity function for the corresponding Thompson automata. Non-SSNF regexes cause critical time explosion, which is evidence that the regex engines do not apply SSNF transformation to their input. In addition to non-SSNF regexes, critical REDoS situations occur on polynomial ambiguities iterated under a Kleene star.

4.2 Comparing with other tools

We evaluated the effectiveness of the proposed approach by comparing it with three state-of-the-art open-source tools for detecting vulnerabilities in regexes: RSA [3], [17], a static analysis tool, RESCUE [5], [18], a genetic fuzzing tool, REVEALER [2], [19], an automated hybrid analysis tool that uses static and dynamic approaches.

The qualitative results of the experiments are described in *Table 2*. To evaluate the effectiveness of detection of vulnerable and safe regexes, we used F_1-score, where true positive values are all vulnerable regular expressions that were classified as exponential or polynomial, the absence of results due to a timeout is taken into account as a false result, also we used the error rate, where a cumulative error on all classes of regexes – total error rate and a classification error among vulnerable regexes – vulnerable error rate. It should be noted that RESCUE does not support the exponential-polynomial classification, therefore, not all values were calculated for this tool.

Table 2. Evaluation results

Tool	F_1-score	Total error rate	Vulnerable error rate
RSA	0.90	0.13	0.00
ReScue	0.39	-	-
Revealer	0.55	0.45	0.04
Our method	1.00	0.00	0.00

The results of measuring the execution time for the considered tools are shown in Table 3. When measuring time, all extended features of the tools were disabled, and their parameters were optimized. For each class of correctly classified regexes: exponential, polynomial, safe, unsafe (union of vulnerable regexes), the average running time (μ) and the standard deviation (σ) of this value were estimated, the number of timeouts was also calculated.

Additionally, we chose 25 regexes with non-SSNF structure, which are analysed in our method by the preliminary ε-loop test. While our approach proved to be the fastest (which is not a surprise, provided the algorithm structure), the static part of REVEALER also had 100% success rate on this set, although, taking at average 4x more time.

It is important to note that the theoretical results obtained by using static analysis methods, determining ambiguity degree of the Thompson automata, completely coincide with the experimental results obtained when testing the domino regexes on the PYTHON, JAVASCRIPT, JAVA 8, and C++ regex engines. This is a strong witness that regexes declared safe by dynamic or combined methods are their false negatives.

Table 3. Time measurements

Tool	Exp		Poly		Safe		Unsafe		Timeouts
	$\mu(s)$	$\sigma(s)$	$\mu(s)$	$\sigma(s)$	$\mu(s)$	$\sigma(s)$	$\mu(s)$	$\sigma(s)$	
RSA	1.895	2.614	3.480	3.748	0.836	0.341	2.578	3.221	13
ReScue	-	-	-	-	0.940	1.724	8.803	6.263	43
Revealer	0.410	0.035	0.402	0.021	0.320	0.065	0.409	0.033	0
Our method	0.836	1.059	1.178	1.259	0.484	0.400	1.014	1.169	0

From the test results, we can conclude that the detection efficiency of the static analyser is high, but in non-trivial exponential or polynomial cases such as `(baa|ab)*b(a(b|a)ba*b*(aab)*`, timeouts occur. The recognition efficiency of RESCUE and REVEALER tools on this data set is low. However, the proposed approach has the maximum *q*uality of vulnerability detection, the average execution time is also superior to other implementations. This is partly explained by its narrow domain: testing only academic regexes. But RSA also aims at the academic regexes, and still has several timeouts; on the other hand, it seems that extension of REDoS-detection tools to non-academic regexes made them to miss almost all polynomial REDoS with domino structure.

5. Discussion and Related Works

Initially, our finite automata transforming tool was not designed to reveal REDoS situations. However, attempts to use open-source tools like Regex Static Analyser or RESCUE to analyze simple academic regexes with non-trivial ambiguity structure failed. The main purpose of the work was educational, so we designed our algorithm in such a way that it not only detects vulnerabilities, but also demonstrates them on the automata graphs, at the cost of longer execution

time. Since the tool was initially designed for demonstrations, only core academic regexes were considered. The algorithms used in the monoid-based approach have poor worst-case complexity, so its efficiency, compared to RSA and RESCUE, was a real surprise.

What features of the analysers caused such a situation? RSA uses NFA intersection construction, based on the well-known paper of Mohri et al [14]. To detect polynomial ambiguities, the algorithm requires self-intersecting an NFA twice. The automata intersection problem is known to be PSPACE-complete [20], [21], thus, every additional intersection results in a significant slowdown. Maybe that is the main cause why the polynomial detection results in timeouts in RSA. The monoid and determinization algorithms are known to be worst-case exponential. However, the determinization is proven to be fast 4 in average [16], while the monoid representation depends heavily on the automata structure and, implemented to orbit automata, generates significantly fewer e*q*uivalence classes, compared to the case when automata are not cyclic. Another well-known problem in static analysers is dealing with ε-transitions, which can ruin the intersection construction, as well as the monoid. Surprisingly, the tools do not use the simple and natural conversion to the Glushkov construction preceded by the SSNF test.

Error rate of static tools is usually much lower than in tools using genetic algorithms and fuzzing, since REDoS-provoking strings can be disguised, requiring several explicit iterations to construct, or be combined from several alternative subexpressions under an iteration. Even using two approaches in REVEALER cannot help to find vulnerabilities, if the malicious pump is hidden in overlaps and crossing occurrences. For example, in paper [6], four REDoS classes are provided, based on a regex structure, and the regex `a*(ab)*a(ba)*` satisfies neither of them, because the vulnerability appears due to the crossing occurrence of the string ab on the border of the two orbits, whereas the expressions under Kleene stars have languages with empty intersection, which makes the regex "seemingly safe". A similar pattern-based approach is used in [7], resulting in the same sort of false negatives. So, regex-based heuristics showed themselves to be too weak as compared to the model NFA analysis in the domino ambiguity cases.

If a malicious pump for a regex is found, the natural question arises: how to correct the regex? We did not consider the whole implementation of the regex correction, but implemented a trial algorithm constructing a 1-unambiguous regex, if it exists [9]. However, for most regexes with overlaps, even if the equivalent 1-unambiguous regex can be built, the algorithm given in [9] produces exponentially longer result, as compared to the input, processing all overlap combinations separately. A more optimistic regex correcting heuristic is the Star Normal Form transformation: it is performed in linear time and produces regexes approximately of the same length. Moreover, the SSNF transformation is rather local, does not require transition to NFA, and can be applied even to extended regexes, which is useful, taking in account that non-SSNF regexes cause critical REDoS w.r.t. PYTHON and JAVASCRIPT regex engines. In general, the question what theoretical results can be used to fix REDoS regexes, is still a subject of research.

6. Conclusion

The research resulted in the following answers to our research questions.

- **RQ1:** *how relevant is NFA static analysis w.r.t. to popular regex engines?*

  Our experiments demonstrated that the Thompson NFA model is entirely suitable for evaluating REDoS situations concerning the most widely used regex engines, including PYTHON, JAVASCRIPT, JAVA, and C++. Interestingly, although the GO regex machine uses conversion to DFA, it still produces surges on some ambiguous regexes with complex structures. The RUST DFA engine proved to be the most stable.

- **RQ2:** *what features of the REDoS analysers considered cause errors and time explosion on the regexes with complex overlap structure? How they can be processed reliably with less risk of time explosion?*

  We found out that considering orbit automata (instead of performing ambiguity analysis on the entire NFA) and using the Glushkov construction, preceded by the Strong Star Normal Form test, do not result in any loss of relevance, but significantly speed up the static analysis.

  Another interesting approach is to use monoid analysis as the primary ambiguity-detecting algorithm instead of NFA intersection analysis. If there are multiple substring overlaps in the orbits, this method performs significantly faster. However, if the overlaps are small, the number of equivalence classes in the monoid increases dramatically, making the intersection method more preferable.

  We also provided experimental evidence that the genetic search REDoS detection methods still miss complex REDoS cases, easily detected by static NFA analysis approaches.

Despite our approach proved itself to be efficient and reliable on the test set of domino regexes, it still requires many refinements. First, the monoid construction may explode if we take large alphabets, so the input regexes may need some alphabet factorization. E.g., if no overlaps are contained within a long string, then this string sometimes can be considered as a single letter. Second, it would be interesting to test the method on extended regexes approximation, and to combine the monoid-based and intersection-based ambiguity detection algorithms.

References

[1]. Davis J. C., Coghlan C. A., Servant F., and Lee D. The impact of regular expression denial of service (ReDoS) in practice: an empirical study at the ecosystem scale. In Proc. of the 2018 ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, 2018. pp. 246-256. DOI: 10.1145/3236024.3236027.

[2]. Liu Y., Zhang M., and Meng W. Revealer: Detecting and exploiting regular expression denial-of-service vulnerabilities. In Proc. of the 2021 IEEE Symposium on Security and Privacy (SP), 2021. pp. 1468-1484. DOI: 10.1109/SP40001.2021.00062.

[3]. Van der Merwe B., Weideman N., and Berglund M. Turning evil regexes harmless. In Proc. of the South African Institute of Computer Scientists and Information Technologists, 2017. pp. 1-10. DOI: 10.1145/3129416.3129440.

[4]. Weideman N., van der Merwe B., Berglund M., Watson B. W. Analyzing matching time behavior of backtracking regular expression matchers by using ambiguity of NFA. In Proc. of the Implementation and Application of Automata - 21st International Conference. 2016. pp. 322-334. DOI: 10.1007/978-3-319-40946-7_27.

[5]. Shen Y., Jiang Y., Xu C., Yu P., Ma X., Lu J. ReScue: Crafting regular expression DoS attacks. In Proc. of the 33rd ACM/IEEE International Conference on Automated Software Engineering, 2018. pp. 225-235. DOI: 10.1145/3238147.3238159.

[6]. Li Y., Sun Y., Xu Z., Cao J., Li Y., Li R., Chen H., Cheung S.-C., Liu Y., Xiao Y. RegexScalpel: Regular expression denial of service (ReDoS) defense by Localize-and-Fix. In Proc. of the 31st USENIX Security Symposium (USENIX Security 22), 2022. pp. 4183-4200.

[7]. Li Y., Chen Z., Cao J., Xu Z., Peng *Q*., Chen H., Chen L., Cheung S. ReDoSHunter: A combined static and dynamic approach for regular expression DoS detection. In Proc. of the 30th USENIX Security Symposium (USENIX Security 21), 2021. pp. 3847-3864.

[8]. Google. Official public repository of RE2 library. Available at: *https://github.com/google/re2*, accessed 01.07.2023.

[9]. Bruggemann-Klein A. and Wood D. One-unambiguous regular languages. Information and Computation, vol. 140, no. 2, 1998. pp. 229-253. DOI: 10.1006/inco.1997.2688.

[10]. Freydenberger D. D., Schmid M. L. Deterministic regular expressions with back-references. Journal of Computer and System Sciences, vol. 105, 2019. pp. 1-39. DOI: 10.1016/j.jcss.2019.04.001.

[11]. Allauzen C., Mohri M. A unified construction of the Glushkov, Follow, and Antimirov automata. In Proc. of the Mathematical Foundations of Computer Science, 2006. pp. 110-121. DOI: 10.1007/11821069_10.

[12]. Eric Pin J. Mathematical foundations of automata theory. Available at: https://www.irif.fr/~jep/PDF/MPRI/MPRI.pdf, accessed 01.07.2023.

[13]. Weber A., Seidl H. On the degree of ambiguity of finite automata. Theoretical Computer Science, vol. 88, no. 2, 1991. pp. 325-349. DOI: 10.1016/0304-3975(91)90381-B.

[14]. Allauzen C., Mohri M., Rastogi A. General algorithms for testing the ambiguity of finite automata. In Proc. of the Developments in Language Theory, 2008. pp. 108-120. DOI: 10.1007/978-3-540-85780-8_8.

[15]. Bruggemann-Klein A. Regular expressions into finite automata. Theoretical Computer Science, vol. 120, no. 2, 1993. pp. 197-213. DOI: 10.1016/0304-3975(93)90287-4.

[16]. Almeida M., Moreira N., Reis R. On the performance of automata minimization algorithms. In Proc. of the 4th Conference on Computability in Europe, 2008. pp. 3-14.

[17]. Weideman N. Regex static analyzer. Available at: *https://github.com/NicolaasWeideman/RegexStaticAnalysis*, accessed 01.07.2023.

[18]. Shen Y., Jiang Y., Xu C., Yu P., Ma X., Lu J. Rescue. Available at: *https://github.com/2bdenny/ReScue*, accessed 01.07.2023.

[19]. Liu Y., Zhang M., Meng W. Revealer. Available at: *https://github.com/cuhkseclab/Revealer*, accessed 01.07.2023.

[20]. Gelade W., Neven F. Succinctness of the Complement and Intersection of Regular Expressions. In Proc. of the 25th International Symposium on Theoretical Aspects of Computer Science, 2008. pp. 325-336. DOI: 10.4230/LIPIcs.STACS.2008.1354.

[21]. Birget J., Margolis S. W., Meakin J. C., Weil P. Pspace-complete problems for subgroups of free groups and inverse finite automata. Theoretical Computer Science, vol. 242, no. 1-2, 2000. pp. 247-281. DOI: 10.1016/S0304-3975(98)00225-4.

Информация об авторах / Information about authors

Антонина Николаевна НЕПЕЙВОДА – научный сотрудник Института Программных Систем РАН. Сфера научных интересов: теория формальных языков, программная семантика, математическая логика и функциональное программирование.

Antonina Nikolaevna NEPEIVODA – researcher in the Program Systems Institute of RAS. Research interests: formal language theory, program semantics, mathematical logic, and functional programming.

Юлия Андреевна БЕЛИКОВА – студент Московского государственного технического университета им. Н. Э. Баумана. Сфера научных интересов: теория формальных языков, искусственный интеллект, анализ данных.

Yulia Andreevna BELIKOVA – student of the Bauman Moscow State Technical University. Research interests: formal language theory, artificial intelligence, data analysis.

Кирилл Константинович ШЕВЧЕНКО – студент Московского государственного технического университета им. Н. Э. Баумана. Сфера научных интересов: анализ данных и машинное обучение.

Kirill Konstantinovich SHEVCHENKO – student of the Bauman Moscow State Technical University. Research interests: data science and machine learning.

Михаил Романович ТЕРЮХА – студент Московского государственного технического университета им. Н. Э. Баумана. Сфера научных интересов: интернет вещей и распределённые вычисления.

Mikhail Romanovich TERIUKHA – student of the Bauman Moscow State Technical University. Research interests: internet of things and distributed systems.

Данила Павлович КНЯЗИХИН – студент Московского государственного технического университета им. Н. Э. Баумана. Сфера научных интересов: абстрактная алгебра.

Danila Pavlovich KNYAZIHIN – student of the Bauman Moscow State Technical University. Research interests: abstract algebra.

Александр Дмитриевич ДЕЛЬМАН – студент Московского государственного технического университета им. Н. Э. Баумана. Сфера научных интересов: конструирование компиляторов и облачные вычисления.

Aleksandr Dmitrievich DELMAN – student of the Bauman Moscow State Technical University. Research interests: compiler design and cloud computing.

Анна Сергеевна ТЕРЕНТЬЕВА – студент Московского государственного технического университета им. Н. Э. Баумана. Сфера научных интересов: конструирование компиляторов.

Anna Sergeevna TERENTYEVA – student of the Bauman Moscow State Technical University. Research interests: compiler design and optimisation.