

Reachability Analysis of the Domain Name System

DHRUV NEVATIA, ETH Zurich, Switzerland

SI LIU, ETH Zurich, Switzerland

DAVID BASIN, ETH Zurich, Switzerland

The high complexity of DNS poses unique challenges for ensuring its security and reliability. Despite continuous advances in DNS testing, monitoring, and verification, protocol-level defects still give rise to numerous bugs and attacks. In this paper, we provide the first decision procedure for the DNS verification problem, establishing its complexity as 2ExpTime , which was previously unknown.

We begin by formalizing the semantics of DNS as a system of recursive communicating processes extended with timers and an infinite message alphabet. We provide an algebraic abstraction of the alphabet with finitely many equivalence classes, using the subclass of semigroups that recognize positive prefix-testable languages. We then introduce a novel generalization of bisimulation for labelled transition systems, weaker than strong bisimulation, to show that our abstraction is sound and complete. Finally, using this abstraction, we reduce the DNS verification problem to the verification problem for pushdown systems. To show the expressiveness of our framework, we model two of the most prominent attack vectors on DNS, namely amplification attacks and rewrite blackholing.

CCS Concepts: • **Security and privacy** → **Logic and verification**; **Denial-of-service attacks**; • **Networks** → **Application layer protocols**; **Formal specifications**; • **Theory of computation** → **Automata over infinite objects**; **Regular languages**.

Additional Key Words and Phrases: DNS, Formal Semantics, Bisimulation, Reachability Analysis

ACM Reference Format:

Dhruv Nevatia, Si Liu, and David Basin. 2025. Reachability Analysis of the Domain Name System. *Proc. ACM Program. Lang.* 9, POPL, Article 62 (January 2025), 32 pages. <https://doi.org/10.1145/3704898>

1 Introduction

The Domain Name System (DNS) is a central component of the Internet's infrastructure. It translates human-readable domain names, such as `www.sigplan.org`, into machine-recognizable IP addresses via *name resolution*, thereby simplifying Internet navigation and resource access for users. This seemingly simple translation service, handled by *recursive resolvers*, is underpinned by an intricate, hierarchical, globally distributed database. Each organization, like Google or Cloudflare, provides name resolution for its portion or *zone* of the entire DNS namespace. Operators within each organization manage that zone through, often manually configured, *zone files*. These zone files, stored on *authoritative nameservers*, map domain names to IP addresses as well as other types of DNS records, specifying further actions to be taken such as query rewriting or delegation.

The IETF (Internet Engineering Task Force) [28] publishes *requests for comments* (RFCs), which are technical documents describing the Internet's technical foundations. To date, hundreds of RFCs (e.g., [20, 39, 40, 48]) have been published that define and guide the design, implementation, and operation of DNS. The sophisticated nature and large scale of DNS pose unique challenges for developers

Authors' Contact Information: [Dhruv Nevatia](mailto:Dhruv.Nevatia@inf.ethz.ch), ETH Zurich, Zurich, Switzerland, dhruv.nevatia@inf.ethz.ch; [Si Liu](mailto:Si.Liu@inf.ethz.ch), ETH Zurich, Zurich, Switzerland, si.liu@inf.ethz.ch; [David Basin](mailto:David.Basin@inf.ethz.ch), ETH Zurich, Zurich, Switzerland, basin@inf.ethz.ch.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2025 Copyright held by the owner/author(s).

ACM 2475-1421/2025/1-ART62

<https://doi.org/10.1145/3704898>

and operators, who have invested considerable efforts to ensure its functional correctness, security, and availability. These challenges often stem from name resolution failures, which result mainly from misconfigurations or attacks, and have historically led to large-scale outages [52, 56, 57]. Of particular concern are *denial of service* (DoS) vectors, which have been frequently identified over the past decade [7, 35, 37]. In particular, *amplification attacks*, which exploit the DNS protocol to significantly increase the query load on the victim (resolvers or authoritative nameservers), have seen a surge in recent years [6, 19, 36, 41, 58].

Despite continuous advances in DNS testing and monitoring [14, 27, 29, 32, 51, 59], numerous issues, like those just mentioned, still arise from RFCs or protocol-level defects in extensively tested production resolvers and nameservers. This highlights the necessity for a principled, proactive way to mitigate such problems, preferably at an early design stage. Two recent formal efforts [30, 36] have attempted to address this issue. GRoot [30] is the first static verifier for DNS configuration errors, building on the protocol-level DNS resolution semantics. It has been successfully applied to identify misconfigurations in DNS zone files, prior to their deployment. DNSMaude [36] is a formal framework, implemented in the Maude [17] language, for both the qualitative (e.g., functional correctness) and quantitative (e.g., amplification) analysis of DNS protocols. It is based on a semantics for the entire end-to-end name resolution, covering essential features such as resolver caching and recursive subqueries, which GRoot abstracts away. *Henceforth, when we refer to DNS, we consider this more comprehensive semantics.* DNSMaude’s accompanying analyzer has also discovered multiple attacks on DNS with large amplification effects.

Research Gaps. Both works have provided partial solutions to the *DNS verification* (DNSVERIF) problem where, given the DNS zone files \mathcal{Z} , one asks whether a property ϕ of interest holds for some behavior of DNS (see Section 5 for its formal definition). Intuitively, the property ϕ captures the *bad* behaviors exhibiting attacks. As every behavior of DNS starts with an initial query from the resolver, on behalf of a client, an immediate challenge in this problem is to explore the infinite space of queries exhaustively. A priori, this was infeasible.

GRoot tackles this problem by introducing an equivalence relation on the query space, with finitely many equivalence classes (ECs), such that queries within an EC are resolved in the same manner and yield the same result. While this equivalence is proven to be sound under a simplified DNS semantics, it is unsound with respect to the more realistic semantics involving resolver caching (formalized in DNSMaude). Moreover, GRoot must be supplied with a pre-determined bound on the length of runs in DNS, making it incomplete by design. Additionally, the number of ECs increases explosively with query rewrites, to the order of n^{255} in the size of zone files, due to advanced DNS features like `DNAME` rewriting [48] (see also Section 2.1), which can quickly inflate GRoot’s verification time. Finally, GRoot is inherently incapable of specifying behavioural vulnerabilities like amplification DoS attacks within its semantic framework.

DNSMaude also works with GRoot’s ECs but supports specifying common behavioural vulnerabilities. However, it still struggles with attack discovery. Since it uses GRoot’s ECs, it may fail to explore certain bad behaviours that it believes are equivalent to those that have already been explored. It may then incorrectly conclude that there is an absence of bad behaviours, rendering its search procedure incomplete. Moreover, DNSMaude’s search procedure is also incomplete by design; it uses bounded, explicit-state, linear temporal logic (LTL) model checking [17] for functional correctness properties, searching for bugs or potential attacks. DNSMaude also employs statistical model checking (SMC) [49] for quantitative properties like amplification. SMC verifies a property probabilistically, up to a given confidence level, by running Monte Carlo simulations. As

a result, the search for attacks is again bounded, for example, in the number of simulations, and attacks may be missed.¹

Overall, neither of these two efforts provides a *decision procedure* for the `DNSVERIFY` problem. Moreover, an upper bound on the complexity of this problem still remains unknown [31].

Our Approach. In this paper, we aim to address the above challenges. As the first step, we must faithfully model the semantics of DNS. We consider the well-studied model of a system of recursive communicating processes [25, 26], and extend it with timers. We call this extension, a *system of recursive communicating processes with timers* (trCPS). We model the DNS semantics as an instance of a trCPS. Intuitively, the timer keeps track of the age of a resource record in a resolver’s cache. As resolvers are *stateful* (due to referrals or query rewriting), we utilize stacks to model how they track the resolution process, where each stack entry corresponds to a subquery being resolved. Nameservers just behave as *labelled transition systems* (LTSs) as they are *stateless*.

DNS is Eager. It is well-known that reachability is undecidable for networks of LTSs that contain cycles in the underlying topology, assuming a finite message alphabet and perfect channels [13]. However, this problem becomes decidable with *lossy communications* [5, 22]. Since DNS typically operates over UDP (User Datagram Protocol) [40], which is inherently lossy, we can use lossy channels. Nevertheless, the problem once again becomes undecidable if a cycle contains a process that is a pushdown system (PDS) [8, Thm. 1]. Unfortunately, DNS falls within this class, as a resolver is modeled as a PDS with timers and communicates bidirectionally with nameservers (sending queries and receiving answers). Heussner et al. [25, 26] show that the reachability problem for networks of PDSs over *perfect* channels and a *finite* message alphabet, is decidable for eager runs and pointed network topologies. We show that all runs of DNS are eager by design, and that the underlying topology is pointed. This makes DNS fall within this decidable class of systems.

Congruence. Since DNS operates over both (i) an infinite query space and (ii) infinitely many timer valuations, the eagerness of DNS by itself is insufficient to obtain decidability. This is because an infinite query space also necessitates an infinite message alphabet. We propose a novel *congruence* on the domain space with finitely many equivalence classes, derived using syntactic congruences for carefully constructed regular languages [42, 46], from the given zone file configuration \mathcal{Z} . We then show that the domains in any two DNS configurations only need to be syntactically congruent to each other with respect to certain *positive prefix-testable languages* (see Section 4.1). Moreover, these languages are parametric only in the names and values appearing in \mathcal{Z} ’s resource records (see Section 3.1), in order to exhibit equivalent behaviours with respect to the DNS semantics. Notably, owing to the nature of these languages, the equivalence is robust to an unbounded number of `DNAME` rewrites (unlike `GRoot`), enabling an unbounded analysis. This effectively yields an abstraction with an *equivalent* finite message alphabet for the system. Finally, the set of timer valuations are abstracted using a standard region construction.

Kernel Bisimulation. The domain congruence induces a homomorphism on the semantics of DNS, allowing us to reduce it to a recursive CPS with finitely many control states. However, to show that the reduction is both sound and complete, we introduce a novel, generalized notion of bisimulation, called *generalized kernel bisimulation*. A crucial property of this bisimulation is the pointwise equivalence of runs from equivalent states *up to the equivalence of transitions*. We prove that the homomorphism induces a generalized kernel bisimulation on the DNS semantics such that any run in the abstract semantics corresponds to a class of equivalent runs in the original

¹We have experimented with `DNSMaude`’s simulation mode, which is used to discover attacks as presented in [36]. We observe that around 30% of the runs do not achieve the maximum amplification factor, thus potentially overlooking attacks.

semantics. We use this to establish both soundness and completeness of our abstract model with respect to DNS behavior.

Reduction to Reachability. As we prove that DNS is an eager trCPS over a pointed network topology, the reduced abstract DNS is also an eager recursive CPS over the same topology. We then reduce the `DNSVERIFY` problem to an instance of the reachability problem for PDSs, which is known to be decidable [12]. Thus, we obtain the decidability of this problem. In addition, we establish that it is solvable in doubly-exponential time in the size of the zone file configuration \mathcal{Z} .

Contributions. Overall, we make the following contributions.

- We establish the first decision procedure for the `DNSVERIFY` problem (Theorem 4.14). Moreover, we show that our algorithm has an upper-bound of 2ExpTime complexity (Lemma 5.3), which was previously unknown.
- At the technical level, we initially propose trCPS, a model for networks of recursive communicating processes with timers, as the underlying formal model for DNS (Section 3). But it comprises an infinite message alphabet due to the infinite query space. Owing to the monoidal nature of queries, we devise an algebraic abstraction for them (Section 4.3), with finitely many equivalence classes, using a special class of semigroups (Lemma 4.5). We then introduce a novel generalization of bisimulation for LTSs, weaker than strong bisimulation [38] but incomparable to weak bisimulation, to show that our abstraction is sound and complete (Section 4.2). Both these tools may also be of independent theoretical interest.
- We show how our framework can be applied to instantiate two of the most prominent attacks on DNS, namely *amplification attacks* and *rewrite blackholing* (Section 6). This demonstrates the versatility of our approach.

2 Preliminaries

2.1 DNS in a Nutshell

The Domain Name System (DNS) is a decentralized database that translates human-readable domain names into IP addresses. For example, when a user visits the website `www.sigplan.org`, the web browser sends a DNS request to retrieve the IP address of the server hosting the site.

Namespace. The DNS namespace is organized hierarchically in a tree structure, with a root domain at the top, followed by top-level domains (TLDs), second-level domains (SLDs), and further subdomains. Figure 1 shows a small portion of this namespace, including the TLDs `org` and `com` as well as SLDs such as `sigplan`. A domain name such as `www.sigplan.org` uniquely identifies a path in this tree, with the associated labels concatenated bottom up, using the ‘.’ character; the root node has an empty label, which is typically omitted. The entire namespace is divided into *zones* or subtrees, each managed by authoritative *nameservers*.

Resource Records. A DNS *resource record* (RR) stores specific information about a domain name, including the owner name (where this RR resides), record type, time to live (TTL), and value. In general, multiple RRs with the same owner name and type can exist, provided they have different values. Common types of RRs include `A`, `NS`, `SOA`, `TXT`, `CNAME`, and `DNAME`, each serving a distinct purpose. For example, an `A` record maps a domain name to an IPv4 address. An `NS` record specifies the authoritative nameservers for a domain. A `CNAME` record maps an alias domain name, e.g., `example.net`, to its canonical name, e.g., `www.example.com`; any request for `example.net` will be redirected to `www.example.com`. A `DNAME` record provides redirection for an *entire subtree* of the domain name space to another domain. For instance, a `DNAME` record at `legacy.example.com`




Fig. 1. A small portion of the DNS namespace. Dashed lines mark the zone cuts.

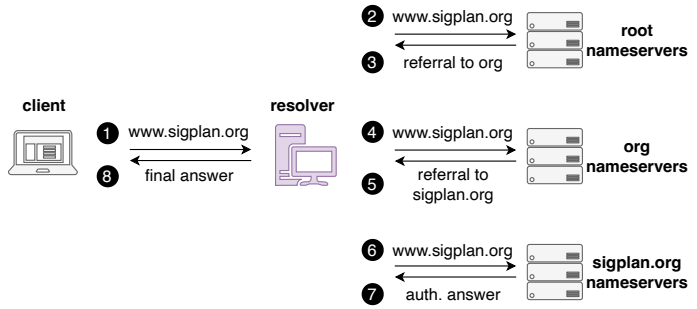


Fig. 2. Name resolution for the type A query `www.sigplan.org`, with an empty cache at the DNS resolver.

with value `new.example.com` would cause a query for any `.legacy.example.com` to be restarted for any `.new.example.com`.

A *wildcard* record is a special type (\star) of DNS record that matches requests for non-existent domain names. This simplifies DNS management by allowing a single record to cover multiple potential queries. For instance, suppose there is an A record with `1.2.3.4` for $\star.example.org$. Queries for `one.example.org` or `any.example.org` will all resolve to `1.2.3.4`, provided that no other records exist for these specific names.

Name Resolution. A client such as a web browser initiates a DNS request by directing the query to a *recursive resolver* that handles the resolution process on the client’s behalf. If the recursive resolver has a cached response from a prior query, it can promptly reply to the client. Otherwise, it queries a hierarchy of authoritative nameservers to retrieve the answer.

Example 2.1 (Name Resolution). Figure 2 depicts a DNS name resolution process. A client initiates a type A query for the name `www.sigplan.org` (1). Assuming an empty cache, the recursive resolver forwards this query to one of the root nameservers (2). Here, the root nameserver cannot answer the query as it is only authoritative for the root zone. Hence, it refers the resolver to the authoritative nameservers for the `org` zone by attaching the associated `ns` records and IP addresses in the referral (3). The resolver then asks one of the `org` nameservers (4). As the nameserver is not authoritative for the full domain, it further refers the resolver to the authoritative nameservers for `sigplan.org` (5). The resolver finally receives the result, the IP address for `www.sigplan.org` (6 and 7), which it then sends back to the client (8).

Note that, owing to *caching*, if the same query needs to be resolved again, the resolver can answer it directly without contacting the nameservers, provided the records’ TTL has not expired. Likewise, for a query of another name within an already cached zone, the resolver only needs to reach out to the relevant nameservers rather than starting from the root. Note also that when a resolver has multiple nameservers to choose from, the selection is typically based on factors like response time and availability. DNS RFCs do not prescribe a specific nameserver selection algorithm, leaving this to the developers’ discretion.

Subqueries. During DNS resolution, the recursive resolver may need to perform additional queries to collect all the necessary information to resolve a client query. For instance, suppose that a user wants to visit `example.org` and the authoritative nameserver for the `example.org` zone is `sub.query.com`, the resolver must first resolve `sub.query.com`.² This process begins by querying

²This is a randomly chosen name just for illustrative purposes. As long as the authoritative nameserver’s name differs from `example.org`, the resolver must then make additional subqueries to resolve it.

the com TLD, followed by query.com (assuming the authoritative nameserver for query.com provides the IP address for sub.query.com). These two additional queries are called *subqueries*, as they are initiated by the resolver to resolve the original query.

2.2 Labelled Transition Systems

Given an alphabet Σ , a *string* $str \in A^*$ is a sequence of labels $l_0 \dots l_k$, where l_0 is the first label and l_k is the last label. The length of the sequence, $k + 1$, is denoted by $|str|$. For $i < j < |str|$, $str[i \dots j]$ denotes the sequence $l_i \dots l_j$ and ϵ denotes the *empty* sequence. Given two strings str and str' , we write $str \preceq str'$ ($str \prec str'$) to denote that str is a (proper) prefix of str' . For notational convenience, we sometimes drop the explicit binary operator for concatenation.

Definition 2.2. A *labelled transition system* (LTS) [26] $\mathcal{A} = \langle S, S_{init}, A, \text{TRANS} \rangle$ is given by a set of *states* S , a set of initial states $S_{init} \subseteq S$, an alphabet A of actions, and a set of transition rules TRANS of the form $s \xrightarrow{a} s'$, where $s, s' \in S$ and $a \in A$. As usual, we define an *automaton* $\mathcal{F} = \langle S, S_{init}, A, \text{TRANS}, S_{fin} \rangle$ as the LTS \mathcal{A} along with a set of *final* states $S_{fin} \subseteq S$. The automaton \mathcal{F} is called a *finite automaton* (FA) if the sets S and A are finite.

A *run* in the LTS \mathcal{A} is a finite sequence $\rho = s_0 \xrightarrow{a_1} s_1 \dots s_{n-1} \xrightarrow{a_n} s_n$, for $s_i \in S$ such that for every $1 \leq i \leq n$, $s_{i-1} \xrightarrow{a_i} s_i \in \text{TRANS}$. We say that ρ is a run from s_0 to s_n . The length of ρ is n and is denoted by $|\rho|$. The *trace* of the run ρ is the sequence of actions $a_1 \dots a_n$. Two runs are *trace equivalent* if they have the same traces. Given the FA \mathcal{F} , a run in \mathcal{A} is *accepting* if it starts in an initial state and ends in a final state.

A state $s \in S$ is *reachable* in \mathcal{A} if there exists a run of \mathcal{A} from an initial state s_0 to s . The *LTS verification problem* asks whether, given an LTS $\mathcal{A} = \langle S, S_{init}, A, \text{TRANS} \rangle$, and an automaton \mathcal{F}' over strings in A^* , there exists a run ρ in \mathcal{A} from an initial state, and an accepting, trace equivalent run ρ' in \mathcal{F}' . We call the corresponding trace a *witness trace*. The problem can be analogously defined for any system with an LTS semantics.

3 A Formal Model for DNS

In this section, we define a formal model for DNS and provide the semantics for all its syntactic components, viz. authoritative nameservers and recursive resolvers. In Section 3.2, we first define the meta-model of recursive communicating processes with timers, which is a timed extension of the commonly known model of recursive communicating processes. In Sections 3.3–3.5, we use our meta-model to comprehensively capture complex features such as recursive subqueries, time-to-live (TTL) records, channel communication, non-determinism, etc. In particular, authoritative nameservers are modelled as LTS components and the recursive resolver is modelled as a tPDS component, allowing it to use timers, for keeping track of the age of cache elements, and a stack, to simulate recursive subquery resolution. As shown in Section 4, we provide a reduction of this model to a recursive CPS.

3.1 Definitions and Notations

Consider a monoid $\mathcal{M} = \langle M, \cdot, 1 \rangle$, where $M = \langle G \rangle$ is generated by the subset $G \subseteq M$, with the binary operation (\cdot) , and the identity element 1 . Given any two subsets $B, C \subseteq M$, we define the sets, $B \cdot C := \{x \cdot y \mid x \in B, y \in C\}$, and $m \cdot B := \{m \cdot x \mid x \in B\}$, for an arbitrary $m \in M$.

3.1.1 Domain names and IP addresses. An IP address is a string $ip \in \mathbb{IP}$, where \mathbb{IP} is the set of all IP addresses, disjoint from M . A *domain name* over \mathcal{M} is a string $d \in M$ used to point to an IP address. The identity element 1 represents the root domain name. Henceforth, as a convention, we write our domain names in the *reverse order*, starting with the labels higher up in the namespace

hierarchy. For example, we will write `com.example.www` instead of `www.example.com`. For better readability, the set of all domain names over \mathcal{M} is denoted by $\text{DOMAIN}_{\mathcal{M}} = M$.

3.1.2 Resource records. A *resource record* $rr = \langle d, t, \sigma, v, b \rangle$ is a tuple consisting of a *record name* $d \in \text{DOMAIN}_{\mathcal{M}} \cup \text{DOMAIN}_{\mathcal{M}\star}$ where $\text{DOMAIN}_{\mathcal{M}\star}$ denotes the set of domain names in wildcard records over the isomorphic monoid $\mathcal{M}\star = \langle \mathcal{M}\star = \langle G\star, \otimes, \mathbf{1}\star \rangle$ where $G\star = \{g\star \mid g \in G\}$, such that $m\star \otimes m'\star = (m \cdot m')\star$; the record type $t \in \text{TYPE} = \{A, NS, MX, DNAME, CNAME, SOA, \dots\}$; a time-to-live value $\sigma \in \mathbb{N}$; the value of the record $v \in \text{IP} \cup \text{DOMAIN}_{\mathcal{M}}$; and a bit value $b \in \{0, 1\}$ that denotes whether the record has been synthesized from a DNAME, CNAME, or WILDCARD record. We shall write $\text{dn}(rr)$, $\text{ty}(rr)$, $\text{ttl}(rr)$, $\text{val}(rr)$, and $\text{syn}(rr)$, to denote the fields d, t, σ, v , and b , in rr , respectively. The class of all resource records is denoted by $\text{RECORD}_{\mathcal{M}}$.

A *zone* is a set of resource records. The set of all zones is denoted by $\text{ZONE}_{\mathcal{M}} = \mathcal{P}(\text{RECORD}_{\mathcal{M}})$. Additionally, we assume all zones are *well-formed*, as described in [30, Appendix A]. We shall write $\text{dn}(z)$ to denote the domain name of a well-formed zone.

A *zone file configuration* (ZFC) is a pair, $\langle \mathcal{Z}, \Theta \rangle$, consisting of a function $\mathcal{Z} : \mathcal{S} \rightarrow \mathcal{P}(\text{ZONE}_{\mathcal{M}})$ from a finite set $\mathcal{S} \subseteq \text{IP}$ of IPs to the powerset of $\text{ZONE}_{\mathcal{M}}$, along with a subset $\Theta \subseteq \mathcal{S}$. Intuitively, \mathcal{S} denotes a set of nameservers and Θ denotes the set of root nameservers in the network; $\mathcal{Z}(s)$ denotes the set of zones for which s is authoritative. The set of all ZFCs is denoted by $\text{ZFC}_{\mathcal{M}}$.

3.1.3 Queries. A DNS *query* $q = \langle d, t \rangle$ is a pair consisting of a domain name $d \in \text{DOMAIN}_{\mathcal{M}}$ and the type of query $t \in \text{TYPE}$. For example, a query expecting a DNAME rewrite for the domain name `com.example` would be $\langle \text{com.example}, \text{DNAME} \rangle$. The set of all queries is denoted by $\text{QUERY}_{\mathcal{M}} = (\text{DOMAIN}_{\mathcal{M}} \times \text{TYPE}) \uplus \{\perp\}$, where \perp denotes a non-existent query. We shall write $\text{dn}(q)$ and $\text{ty}(q)$ to denote the fields d and t .

3.1.4 Answers. Every DNS *answer* $a = \langle l, \text{rec} \rangle$ is a pair consisting of a label $l \in \{\text{ANS}, \text{ANSQ}, \text{REF}, \text{REC}, \text{NXDOMAIN}, \text{REFUSED}\}$ denoting the type of answer (i.e., an answer ANS, a rewrite ANSQ, a delegation REF, a recursive subquery request REC, an indication for a non-existent domain name query NXDOMAIN, etc.) and a set of resource records $\text{rec} \in \mathcal{P}(\text{RECORD}_{\mathcal{M}})$ that contains information for further evaluation. The set of all answers is denoted by $\text{ANSWER}_{\mathcal{M}}$.

3.1.5 Server List. A *server list* $sl \in M^*$, is a string of domain names where each domain name is unique. That is, $sl[i] \neq sl[j]$, for $0 \leq i < j \leq |sl|$. The set of all server lists is denoted by $\text{SLIST}_{\mathcal{M}}$. We additionally define a left action, \odot , of M on $\text{SLIST}_{\mathcal{M}}$ such that for any $m \in M$,

$$m \odot sl = \begin{cases} (m)(sl), & \text{if } m \text{ does not match any position in } sl \\ (m)(sl'), & \text{otherwise,} \end{cases}$$

where sl' is the string derived from sl by dropping the (only) position labelled by m . Given a string $\mathbf{m} = m_1 \dots m_n \in M^*$, we denote by $\mathbf{m} \odot sl$, the right associative composition of actions $m_1 \odot (\dots \odot (m_n \odot sl))$.

Figure 3 describes some notation we use when specifying DNS behavior, along with the intended semantics. RANK adopts the product ordering \triangleleft over the order $\text{false} \leq \text{true}$ of truth values.

Example 3.1. Consider the free monoid of strings $\mathcal{L} = \langle \text{LABEL}^*, (\cdot), \epsilon \rangle$, generated by an infinite set of labels LABEL under concatenation. The strings `com.example.a` and `com.example.b` are examples of domains in $\text{DOMAIN}_{\mathcal{L}}$, if the labels `com`, `example`, `a`, `b` $\in \text{LABEL}$. The empty string ϵ represents the root domain.

The expression $(\text{com}^{-1}.\text{com.example.a})$ denotes the singleton set $\{\text{example.a}\}$; $\text{prec}(\text{com.example.a})$ is the set of prefixes $\{\epsilon, \text{com}, \text{com.example}, \text{com.example.a}\}$;

$d_1^{-1} \cdot d_2 := \{d \mid d \in M, d_1 \cdot d = d_2\}$	domain inverse relation
$prec(d) := \{d_1 \mid d \in d_1 \cdot M\}$	domain prefix set
$d_1 \simeq_j d_2 := prec(d_1) \cap prec(d_2) \geq j - 1$	domain prefix match
$max_{\simeq}(d_1, d_2) := prec(d_1) \cap prec(d_2) - 1$	maximal prefix match
$d_1 < d_2 := prec(d_1) \subset prec(d_2)$	domain partial order
$dn(rr) \ni_* d := d_* \in dn(rr) \otimes G_* \wedge dn(rr) \in M_*$	domain wildcard match
$WILDCARD(rr) := dn(rr) \in M_*$	wildcard record
$MATCH(rr, d) := dn(rr) \ni_* d \vee (d \in dn(rr) \cdot M \wedge dn(rr) \in M)$	record matches query
$RANK(rr, \langle d, t \rangle, z) := \langle MATCH(rr, d), ty(rr) = NS \wedge dn(rr) \neq dn(z), max_{\simeq}(dn(rr), d), WILDCARD(rr) \rangle$	resource record rank
$rr_1 <_{q,z} rr_2 := RANK(rr_1, q, z) \triangleleft RANK(rr_2, q, z)$	resource record order

Fig. 3. Some DNS notation.

$max_{\simeq}(\text{com.example.a}, \text{com.example.b})$ is 2, the size of the set $\{\text{com}, \text{com.example}\}$, of non-trivial common prefixes; consequently, $(\text{com.example.a} \simeq_1 \text{com.example.b})$ holds; however, $(\text{com.example.a} \leq \text{com.example.b})$ does not hold as $\text{com.example.a} \notin prec(\text{com.example.b})$.

The next lemma follows immediately from our definitions.

LEMMA 3.2. *The sets $DOMAIN_{\mathcal{M}}$, $RECORD_{\mathcal{M}}$, $ZONE_{\mathcal{M}}$, $ZFC_{\mathcal{M}}$, $QUERY_{\mathcal{M}}$, $ANSWER_{\mathcal{M}}$, and $SLIST_{\mathcal{M}}$ are all functorial³ in the monoid \mathcal{M} .*

Notably, IP addresses remain unchanged with the change in the underlying monoid \mathcal{M} .

3.2 Recursive Communicating Processes with Timers

We will now introduce the notion of recursive communicating processes with timers, extending the model of recursive communicating processes (rCPS) [25, 26], where finitely many processes communicate with each other over perfect channels, dictated by an underlying topology. Therefore, we first define the notion of a network topology, which is a directed graph whose vertices represent processes and whose edges represent channels.

Definition 3.3. A network topology \mathcal{T} is a 3-tuple $\langle P, CH, ch \rangle$, where P is a finite set of processes, CH is a finite set of perfect channel identifiers, and $ch : CH \rightarrow P \times P$ is a partial map from channel identifiers to the processes it connects.

We write ch_c to denote the channel from the *source* process $i \in P$ to the *destination* process $j \in P$, if $ch(c) = \langle i, j \rangle$. Moreover, we assume that $ch^{-1}(\langle i, i \rangle)$ is always empty i.e., there is no channel from a process to itself.

Consider a network topology $\Gamma = \langle P, CH, ch \rangle$. Given a message alphabet M , we define the set of *channel actions*, $CACTION_{\Gamma, M}^P = \{ch_c!a \mid c \in ch^{-1}\langle p, p' \rangle, p' \in P, a \in M\} \cup \{ch_c?a \mid c \in ch^{-1}\langle p', p \rangle, p' \in P, a \in M\}$ for every $p \in P$. The operations of the form $ch_c?a$ and $ch_c!a$ denote the actions of *receiving* and *sending* the message a on channel ch_c respectively.

Next, we define the model of pushdown systems with timers and a timeless stack, where we follow the model proposed in [18, §2], but use timers [34, 53] instead of clocks.⁴ Unlike clocks,

³Given an expression A that contains B , we say that A is *functorial* in B , if $B \mapsto A$ is a functor.

⁴The restriction to timers reduces the space complexity of the abstract configuration space we introduce in Section 4.4. From a decidability perspective, all the reductions presented in this paper hold for the more expressive class of timed pushdown systems proposed in [18].

timers run *backwards*, starting from an initial value. Moreover, we restrict the timing constraints to only allow *timeout checks*. A *timing constraint* over a set of timers X is a formula generated by the grammar,

$$cc_X ::= x > 0 \mid cc_X \wedge cc_X,$$

where $x \in X$ and timer resets are represented by $[\mathbf{x} \leftarrow \mathbf{k}]$, with $\mathbf{k} \in \mathbb{N}^n$; $\mathbf{x} \in X^n$ is a vector with distinct timer components, denoting that each timer x_i in \mathbf{x} is reset to the value k_i .

A *timer valuation* over X is a map $v : X \rightarrow \mathbb{Q}^{\geq 0} \cup \{\infty\}$, where X is a set of timers and each timer either holds a non-negative rational number (denoted by $\mathbb{Q}^{\geq 0}$) or it has not yet been initialized (denoted by ∞). Given $k \in \mathbb{Q}^{\geq 0}$, we denote by $v - k$ the valuation, where for each $x \in X$,

$$(v - k)(x) = \begin{cases} v(x) - k, & \text{if } v(x) \geq k \\ 0, & \text{if } v(x) < k \\ \infty, & \text{if } v(x) = \infty. \end{cases}$$

Given a vector $k \in \mathbb{N}^n$, we denote by $v[\mathbf{x} \leftarrow \mathbf{k}]$ the updated timer valuation such that $v[\mathbf{x} \leftarrow \mathbf{k}](x) = k_i$ if $x = x_i$ for some i ; otherwise, $v[\mathbf{x} \leftarrow \mathbf{k}](x) = v(x)$.⁵

Definition 3.4. A *pushdown system with timers* (tPDS) is a 6-tuple $\mathcal{P} = \langle S, S_{init}, A, \Gamma, X, \text{TRANS} \rangle$, where S is a finite set of states, $S_{init} \subseteq S$ is a finite set of initial states, A is a finite input alphabet, Γ is a finite stack alphabet, X is a finite set of timers, and TRANS is a set of transition rules. The transition rules are of the form $s \xrightarrow[\text{[x} \leftarrow \mathbf{k}], cc]{\langle \alpha, a, \alpha' \rangle} s'$, with $s, s' \in S$ states, $a \in A \cup \{\epsilon\}$ an input letter, and $\mathbf{x} \in X^n$ an n -dimensional vector of timers, being reset to \mathbf{k} , for some n ; $cc \in cc_X$ is a timer constraint; $\alpha \in \Gamma^*$ represents the finite lookup of the top of the stack, and $\alpha' \in \Gamma \cup \epsilon$ represents the element to be pushed (or none thereof, if $\alpha' = \epsilon$). The *size* of \mathcal{P} is the cumulative size of all its components.

A *configuration* C of the tPDS is a 3-tuple $\langle s, v, \gamma \rangle$, where $s \in S$, v is a *timer valuation* over X , and $\gamma \in \Gamma^{\dagger}$ is the stack valuation with the stack top being the leftmost label in γ , and the special symbol $\dagger \notin \Gamma$ denoting the bottom of the stack. The set of all configurations is the set $\text{CONFIG}(\mathcal{P}) = S \times (\mathbb{Q}^{\geq 0} \cup \{\infty\})^X \times \Gamma^{\dagger}$. The set of initial configurations is $\text{CONFIG}(\mathcal{P})_{init} = S_{init} \times \{\infty\}^X \times \{\dagger\}$. A configuration in $\{s\} \times (\mathbb{Q}^{\geq 0} \cup \{\infty\})^X \times \Gamma^{\dagger}$, for $s \in S$, is called an *s-configuration*.

The semantics for the tPDS is the LTS $\langle \mathcal{P} \rangle = \langle \text{CONFIG}(\mathcal{P}), \text{CONFIG}(\mathcal{P})_{init}, A', \text{TRANS}' \rangle$, where $A' = ((\Gamma^* \cup \Gamma^{\dagger}) \times A \times (\Gamma \uplus \{\epsilon\})) \times \{[\mathbf{x} \leftarrow \mathbf{k}] \mid \mathbf{x} \in X^n, \mathbf{k} \in \mathbb{N}^n, n \in \mathbb{N}\} \times cc_X$ and TRANS' contains two kinds of transitions:

- (i) *timed transitions* labelled by $t \in \mathbb{Q}^{\geq 0}$, such that $\langle s, v, \gamma \rangle \xrightarrow{t} \langle s, v - t, \gamma \rangle$, and
- (ii) *discrete transitions*, such that $\langle s, v, \gamma \rangle \xrightarrow[\text{[x} \leftarrow \mathbf{k}], cc]{\langle \alpha, a, \alpha' \rangle} \langle s', v', \gamma' \rangle \in \text{TRANS}'$, whenever $v \models cc$, $v' =$

$$v[\mathbf{x} \leftarrow \mathbf{k}], s \xrightarrow[\text{[x} \leftarrow \mathbf{k}], cc]{\langle \alpha, a, \alpha' \rangle} s' \in \text{TRANS}, \text{ and we have,}$$

$$\gamma' = \begin{cases} \alpha' \gamma, & \text{if } \alpha' \neq \epsilon, \alpha = \epsilon \\ \gamma, & \text{if } \alpha' = \epsilon, \alpha = \epsilon \\ \gamma[1 \dots |\gamma|], & \text{if } \alpha' = \epsilon, \alpha \neq \epsilon, \text{ and } \gamma = \alpha \gamma'', \text{ for some } \gamma'' \in \Gamma^{\dagger} \\ \alpha' \cdot \gamma[1 \dots |\gamma|], & \text{if } \alpha' \neq \epsilon, \alpha \neq \epsilon, \text{ and } \gamma = \alpha \gamma'', \text{ for some } \gamma'' \in \Gamma^{\dagger} \end{cases}$$

⁵Note that the semantics of a timer is less expressive than that of a clock. However, since most implementations of DNS use the notion of timer-based timeouts in their states, using timers greatly simplifies further reductions in this paper.

Intuitively, each stack operation checks whether a finite length of the top of the stack matches α , and performs either a push operation, *nop* (i.e., let the stack remain the same), pop operation, or a *pop and push* operation.

The runs of $\langle \mathcal{P} \rangle$ are slightly different from the definition for LTS given in Section 2.2. A run in $\langle \mathcal{P} \rangle$ is a finite *alternating* sequence of the form $\rho = C_0 \xrightarrow[\text{[x}_1 \leftarrow \text{k}_1], \text{cc}_1]{\langle \alpha_1, a_1, \alpha'_1 \rangle} C_1 \xrightarrow{t_1} C'_1 \xrightarrow[\text{[x}_2 \leftarrow \text{k}_2], \text{cc}_2]{\langle \alpha_2, a_2, \alpha'_2 \rangle} C_2 \dots C'_{n-1} \xrightarrow[\text{[x}_n \leftarrow \text{k}_n], \text{cc}_n]{\langle \alpha_n, a_n, \alpha'_n \rangle} C_n$, where $t_j \in \mathbb{Q}^{\geq 0}$ and for every $1 \leq i \leq n$, $C'_{i-1} \xrightarrow[\text{[x}_n \leftarrow \text{k}_n], \text{cc}_i]{\langle \alpha_i, a_i, \alpha'_i \rangle} C_i \in \text{TRANS}'$. The trace of ρ is the sequence $a_1 \dots a_n \in A^*$. If the set of timers X is empty, \mathcal{P} behaves as a standard *untimed* pushdown system (PDS).

3.2.1 Recursive Communicating Processes with Timers. We follow the definition of an rCPS proposed in [26, Definition 1.7], extending it with timers.

A *system of recursive communicating processes with timers* (trCPS) $\mathcal{R} = \langle T, M, (\mathcal{P}^p)_{p \in P} \rangle$, is given by a network topology $T = \langle P, \text{CH}, \text{ch} \rangle$, a *message* alphabet M , and for each process $p \in P$, a tPDS $\mathcal{P}^p = \langle S^p, S^p_{\text{init}}, A^p, \Gamma^p, X^p, \text{TRANS}^p \rangle$ such that $\text{CACTION}^p_{T, M} \subseteq A^p$. We write $s = \prod_{p \in P} S^p$ for the set of *global states*. The topology T is *pointed* with respect to \mathcal{R} if exactly one process in P corresponds to a tPDS and all the others correspond to an LTS with finitely many states (see Figure 4).

The semantics of trCPS is defined in terms of a global LTS $\langle \mathcal{R} \rangle = \langle \text{CONFIG}, \text{CONFIG}_{\text{init}}, A, \text{TRANS} \rangle$, where $\text{CONFIG} = (\text{CONFIG}(\mathcal{P}^p))_{p \in P} \times (M^*)^{\text{CH}}$ is the set of *global configurations*, $\text{CONFIG}_{\text{init}} = (\text{CONFIG}(\mathcal{P}^p)_{\text{init}})_{p \in P} \times \{\epsilon\}^{\text{CH}}$ is the set of initial configurations, and $A = \biguplus_{p \in P} A^p$ is the set of actions. Given a *global configuration* $c \in \text{CONFIG}$, we write $c^p \in \text{CONFIG}(\mathcal{P}^p)$ to denote the *local configuration* in c that corresponds to the process $p \in P$, and c^c to denote the contents of channel ch_c such that $c^c = c^{c'}$ if $\text{ch}_c = \text{ch}_{c'}$. A configuration in c is called an *s-configuration* when, for every $p \in P$, c^p is an s^p -configuration, for some $s = \langle s^p \rangle_{p \in P} \in s$.




Fig. 4. A pointed topology. The red process is the only tPDS in the network.

The discrete transitions in TRANS of the form $c \xrightarrow[\text{[x} \leftarrow \text{k}], \text{cc}]{\langle \alpha, a, \alpha' \rangle} c'$, for $a \in A^p$, satisfy the following:

- (i) $c^p \xrightarrow[\text{[x} \leftarrow \text{k}], \text{cc}]{\langle \alpha, a, \alpha' \rangle} c'^p \in \text{TRANS}'^p$ and $c^q = c'^q$ for all $q \in P \setminus p$,
- (ii) if $a \in A^p \setminus \text{CACTION}^p_{T, M}$, $c^c = c'^c$ for $c \in \text{CH}$,
- (iii) if $a = \text{ch}_c!m$ then $c^{c''} = c'^{c''} \cdot m$ for each $c'' \in \text{ch}^{-1}(\text{ch}_c)$ and $c'^{c''} = c''$ for $c'' \in \text{CH} \setminus \text{ch}^{-1}(\text{ch}_c)$,
- (iv) if $a = \text{ch}_c?m$ then $m \cdot c'^{c''} = c''$ for each $c'' \in \text{ch}^{-1}(\text{ch}_c)$ and $c'^{c''} = c''$ for $c'' \in \text{CH} \setminus \text{ch}^{-1}(\text{ch}_c)$.

The timed transitions in TRANS' of the form $c \xrightarrow{t} c'$ satisfy $c^p \xrightarrow{t} c'^p \in \text{TRANS}'^p$ for all $p \in P$. We will refer to runs in $\langle \mathcal{R} \rangle$ as runs in \mathcal{R} .

Observe that, if none of the component processes in a trCPS uses timing operations, our model essentially becomes a rCPS [26, Def. 1.7]. Similarly, without timing and stack operations, the model becomes a CPS [26, Def. 1.2].

3.2.2 Eager Runs. The notion of eager runs was first introduced by Heussner et. al. [25, 26], in the context of rCPS. We will extend this notion to trCPS.

Definition 3.5. A run $\rho = C_0 \xrightarrow[\text{[x}_1 \leftarrow \text{k}_1], \text{cc}_1]{\langle \alpha_1, a_1, \alpha'_1 \rangle} C_1 \xrightarrow{t_1} C'_1 \xrightarrow[\text{[x}_2 \leftarrow \text{k}_2], \text{cc}_2]{\langle \alpha_2, a_2, \alpha'_2 \rangle} C_2 \dots C'_{n-1} \xrightarrow[\text{[x}_n \leftarrow \text{k}_n], \text{cc}_n]{\langle \alpha_n, a_n, \alpha'_n \rangle} C_n$ in a trCPS is *eager* if for all $1 \leq i \leq n$, if a_i is of the form $\text{ch}_c?m$ then $i > 1$ and a_{i-1} is the corresponding write action $\text{ch}_c!m$ in ρ .

Intuitively, *eagerness* formalizes that, in a run, send actions are immediately preceded by their matching receive actions. A configuration $c \in \text{CONFIG}$ is *eager-reachable* in a trCPS \mathcal{R} if there exists an eager run from an initial configuration in $\text{CONFIG}_{\text{init}}$ to c . We say that a trCPS \mathcal{R} is *eager* when every run in \mathcal{R} is eager. The *state eager-reachability problem* for \mathcal{R} asks whether, for a given global state $s \in \mathcal{S}$, there exists an eager run from an initial configuration to an s -configuration.

Heussner et. al. [25, Prop. 2.4] characterized a subclass of topologies for which the state eager-reachability problem for rCPS, satisfying certain additional restrictions on their communication behaviour, is decidable. As it happens, the class of rCPS over pointed topologies satisfies these conditions; in particular, this fact is subsumed by the case (ii) in their proof, where they reduce the search for an eager run in the rCPS to the search for a run in a (untimed) pushdown system. The reduced pushdown system even preserves the set of traces of all eager runs in the original rCPS. In the case of pointed topologies, such a reduction leads to a pushdown system of size polynomial in the size of each of the processes.

LEMMA 3.6. *Given any rCPS \mathcal{R} , over a pointed topology, it is PolyTime-reducible to an untimed pushdown system \mathcal{P} with the set of states as the set of global states in \mathcal{R} , such that, for every eager run in \mathcal{R} , from an s -configuration to an s' -configuration, there exists a run in \mathcal{P} , from an s -configuration to an s' -configuration, with the same trace.*

3.3 DNS Semantics

Given an arbitrary zone file configuration $\langle \mathcal{Z}, \Theta \rangle \in \text{ZFC}_{\mathcal{M}}$, we model the corresponding instance of DNS as the trCPS $\text{DNS}_{\mathcal{Z}} = \langle \mathbb{T}_{\mathcal{Z}}, \mathbb{M}_{\mathcal{Z}}, (\mathcal{P}^p)_{p \in P_{\mathcal{Z}}} \rangle$, where $\mathbb{T}_{\mathcal{Z}} = \langle P_{\mathcal{Z}}, \text{CH}_{\mathcal{Z}}, \text{ch}_{\mathcal{Z}} \rangle$, and $\mathbb{M}_{\mathcal{Z}} = \text{QUERY}_{\mathcal{M}} \uplus \text{ANSWER}_{\mathcal{M}}$ is the disjoint union of the set of all queries and answers, such that the following hold:

- (i) $P_{\mathcal{Z}} = \{r\} \uplus \mathcal{S} \subseteq \text{IP}$,
- (ii) $\text{CH}_{\mathcal{Z}} = \mathcal{S} \times \mathcal{S} \uplus \{r\} \times \mathcal{S} \uplus \mathcal{S} \times \{r\}$,
- (iii) $\text{ch}_{\mathcal{Z}}$ is the identity map,
- (iv) \mathcal{P}^p is a tPDS of the form Res_p , if $p = r$, and an LTS of the form NS_p , otherwise.

Intuitively, Res_r denotes the resolver, identified by its IP address; there are no communication channels between the resolver to itself. This follows, by convention [30, 36], from a strict restriction on our model, where we typically do not allow a nameserver to also serve as a resolver in the network. However, note that this is a logical separation, as in reality this is not the case; resolvers may indeed themselves behave as nameservers [39, 40] in a network; NS_p is the formal model of a nameserver. We define these models below, in Sections 3.4 and 3.5. Note that $\mathbb{T}_{\mathcal{Z}}$ is a pointed topology with only one tPDS process \mathcal{P}^r . We shall crucially use this fact in Section 5 to obtain the decidability of the $\text{DNS}_{\text{VERIFY}}$ problem.

Unless otherwise stated, we will implicitly assume that the underlying monoid \mathcal{M} is the free monoid \mathcal{L} , as described in Example 3.1. For notational convenience, we shall thus drop the corresponding subscript \mathcal{L} from all our terms. For the rest of the paper, we also arbitrarily fix an instance of DNS with the zone file configuration $\langle \mathcal{Z} : \mathcal{S} \rightarrow \mathcal{P}(\text{ZONE}), \Theta \rangle$, denoting it as DNS , by dropping the subscript \mathcal{Z} . If clear from the context, we shall also use \mathcal{Z} interchangeably to denote both, the set of all zones in the image of the configuration, and the set of all records in those zones.

3.4 Authoritative Nameserver

An *authoritative nameserver*, NS_s , identified by its IP address, $s \in \mathcal{S}$, is an LTS of the form $\langle \text{NSCONFIG}_s, \{\{init, \perp, \perp\}\}, \text{SACTION}_s, \text{TRANS}_s \rangle$, where NSCONFIG_s is the set of tuples $\{init, lookup, exact, wildcard\} \times \text{QUERY} \times \mathcal{R}$, $\text{SACTION}_s = \text{CACTION}_{\mathbb{T}, \mathcal{M}}^s \cup \{in\}$ with the special action label in for internal actions, and $\text{TRANS}_s \subseteq \text{NSCONFIG}_s \times \text{SACTION} \times \text{NSCONFIG}_s$.

Given a configuration $p = \langle status, q, r \rangle$,

$$\begin{aligned}
\mathcal{N}(s, q) &= \max_{\text{dn}} \{z \in \mathcal{Z}(s) \mid \text{dn}(z) \leq \text{dn}(q)\} & \mathcal{M}(s, q) &= \bigcup_{z \in \mathcal{N}(s, q)} \{rr \in \max_{< q, z} z\} \\
\mathcal{T}(s, q) &= \{rr \in \mathcal{M}(s, q) \mid \text{ty}(rr) = \text{ty}(q)\} & \tau(s, q) &= \{\text{ty}(rr) \mid rr \in \mathcal{M}(s, q)\} \\
\mathcal{T}_t(s, q) &= \{rr \in \mathcal{M}(s, q) \mid \text{ty}(rr) = t\} & \mathcal{G}(s, q) &= \mathcal{T}_{\text{NS}}(s, q) \cup \mathcal{G}_A(s, q) \\
\mathcal{G}_A(s, q) &= \{rr \in \bigcup \mathcal{Z}(s) \mid \exists rr' \in \mathcal{T}_{\text{NS}}(s, q). \text{val}(rr') = \text{dn}(rr) \wedge \text{ty}(rr) = \wedge\}
\end{aligned}$$

Fig. 5. Some auxiliary operations.

- (1) *status* is *init* when the nameserver does not have pending query requests, *lookup* when it must start the lookup process for q , *exact* when $\text{dn}(q)$ finds an exact match with the domain name field of some resource record in a zone that NS_s is authoritative for, and *wildcard* when $\text{dn}(q)$ finds a wildcard match with the domain name field of some resource record in such a zone.
- (2) q is the current query.
- (3) r is the identity of the resolver that the nameserver communicates with.

Additionally, we define auxiliary operations in Figure 5. They all take the server identity s and an input query q as input. $\mathcal{N}(s, q)$ obtains the zones whose domain names have a maximal match with q . $\mathcal{M}(s, q)$ accumulates the resource records with maximal RANK in the zones in $\mathcal{N}(s, q)$. $\mathcal{T}_t(s, q)$ gathers the resource records in $\mathcal{M}(s, q)$ of type t . $\mathcal{T}(s, q)$ is a special case of $\mathcal{T}_t(s, q)$ when $t = \text{ty}(q)$. $\tau(s, q)$ gathers the types of the resource records in $\mathcal{M}(s, q)$. $\mathcal{G}(s, q)$ yields the NS records in $\mathcal{M}(s, q)$ that have an associated \wedge record. Finally, $\mathcal{G}_A(s, q)$ denotes the set of associated \wedge records. Since the operators are parameterized only by a DNS query and the choice of nameserver, they can be computed from the current state information in NS_s .

3.4.1 Transitions. The transitions model the local steps taken by the nameserver during query resolution. In Figure 6, we illustrate a transition in this set.

$$\frac{\text{dn}(\mathcal{M}(s, q)) < \text{dn}(q), \text{DNAME} \in \tau(s, q), \text{dn} \in \text{DOMAIN}, \mathcal{T}_{\text{DNAME}}(s, q) = \{rr\}, rr' = \langle \text{dn}(q), \text{CNAME}, \text{ttl}(rr), \text{val}(rr)\text{dn}(rr)^{-1}\text{dn}(q), 1 \rangle}{\langle \text{lookup}, q, r \rangle \xrightarrow{c_{(\text{dn}, s), r}!(\text{ANSQ}, \{rr, rr'\})}} \langle \text{init}, \perp, \perp \rangle}$$

Fig. 6. A transition in NS_s . The premise of the transition specifies the condition when the transition is taken. The action label is above the arrow.

In this transition, the nameserver NS_s is about to start the lookup process, initiated by the status *lookup*, for the current query q , requested by the resolver r . The premise component $\text{dn}(\mathcal{M}(s, q)) < \text{dn}(q)$ states that $\text{dn}(q)$ is an extension of the domain name of every resource record in $\mathcal{M}(s, q)$, where $\mathcal{M}(s, q)$ represents the set of the resource records that the nameserver may potentially send back in the response to the query q ; the components $\text{DNAME} \in \tau(s, q)$ and $\mathcal{T}_{\text{DNAME}}(s, q) = \{rr\}$ state that the record rr in $\mathcal{M}(s, q)$ is a DNAME record. In its response, the nameserver must send back the record rr and its CNAME derivative synthesized for q in order to prompt a rewrite of the query. The premise component $rr' = \langle \text{dn}(q), \text{CNAME}, \text{ttl}(rr), \text{val}(rr)\text{dn}(rr)^{-1}\text{dn}(q), 1 \rangle$ represents the synthesized record as denoted by the fact that $\text{syn}(rr') = 1$, and $\text{val}(rr')$ is the rewritten domain name. With this, the nameserver then sends the answer $\langle \text{ANSQ}, \{rr, rr'\} \rangle$, which prompts a rewrite (ANSQ) using the records rr and rr' , and moves back to its initial configuration.

For the most part, the transitions correspond to the DNS lookup semantics defined in [30, Figure 3]. Figure 7 demonstrates two other nameserver transitions that showcase two tests that are run by the nameserver on the incoming query name. In particular, notice their premises: $\text{dn}(\mathcal{M}(s, q)) = \text{dn}(q)$ checks whether the query name can be exactly matched with the name of a resource record in $\mathcal{M}(s, q)$. If so, NS_s transitions from the *lookup* status to *exact*. In the second rule, $\text{dn}(\mathcal{M}(s, q)) \ni_{\star} \text{dn}(q)$ checks whether there is a wildcard record name that pattern-matches with the current query. If so, the system transitions to *wildcard* status.




Fig. 7. Two other transitions.

3.5 Resolver


The *resolver* Res_r is a tPDS of the form $\langle \text{RESSTAT}_r, \{\text{init}_r\}, \text{REACTION}_r, \text{QUERY}, X_{\mathcal{Z}}, \text{TRANS}_r \rangle$, where, $\text{REACTION}_r = \text{CACTION}_{T, M}^r \cup \{\text{in}\}$ and $X_{\mathcal{Z}} = \{x_{rr}\}_{rr \in \mathcal{Z}}$ is the set containing a unique timer for each record in \mathcal{Z} . Next, we describe the sets RESSTAT and $\text{TRANS}_r \subseteq \text{RESSTAT}_r \times \text{REACTION} \times \text{RESSTAT}_r$, of states and transitions, respectively.

The set $\{\text{init}, \text{query}, \text{wait}\} \times \text{SLIST}$, is the set of control states in Res_r , and $\text{init}_r = \langle \text{init}, \epsilon \rangle$ is the initial control state. Intuitively, given a configuration $\langle \langle \text{status}, \text{sl} \rangle, \text{cache}, \text{st} \rangle$ of the tPDS Res_r :

- (i) *status* is *init* when the resolver awaits a query request, *query* when it has an active query to resolve, and *wait* when the resolver awaits a response for the active query under resolution.
- (ii) *sl* maintains a list of domain names, which correspond to nameservers, available to the resolver for query resolution, in a decreasing order of *priority*. This means, if $d = \text{sl}[i]$ for some $0 \leq i \leq |\text{sl}| - 1$, then the resolver understands that d is the domain name of a nameserver in the network that may potentially be able to answer its query; moreover, it expects that the IP address of the nameserver itself is stored in its cache. The action \odot denotes the addition of a nameserver reference to *sl*, while potentially assigning it the maximum priority, if it is added to the left of *sl*. In practice, the resolver only queries the leftmost domain name in *sl*. If for some reason, the leftmost domain name gets removed, the next domain name becomes available for querying.
- (iii) *st* $\in \text{QUERY}^*$ denotes the current stack contents of the configuration, recording the query hierarchy in case of subquery resolution. The stack top denotes the present subquery under resolution and the query at the stack bottom denotes the original query to be resolved.
- (iv) *cache* denotes the resolver's current cache contents. If the valuation for the timer x_{rr} , associated with the record rr , satisfies the constraint $x_{rr} > 0$, i.e., $\text{cache}(x_{rr}) > 0$, it means that the record rr is still alive in the cache of the resolver and can be used for query resolution. Equivalently, we will also write $rr \in \text{cache}$ to denote this fact.

3.5.1 Transitions. In Figure 8, we illustrate some transitions in TRANS_r . The action sequence labels are above the arrow and the clock constraints below. To simplify the presentation, we have concatenated the action labels to denote the execution of a sequence of actions. Below, we describe these transitions in more detail.

Transition ①. Before this transition, the resolver status is *query*, indicating that it is about to resolve the query q_1 at the top of the stack. The premise component $\text{dn}(q_1) \leftarrow_A rr$, for $rr \in \text{cache}$, states that q_1 can be resolved with the \wedge record rr in the resolver's cache; this means that the resolver can immediately resolve q_1 and thus pop the top of the stack. Note that the resolver checks that the record rr has not yet expired, with the clock constraint $x_{rr} > 0$. Since the action component $q_1 \cdot q_2$ checks that the rest of the stack is non-empty, the resolver can then revert to the resolution of the next stack entry (which is the previous subquery) q_2 . It does this by sending the query q_2 to

Fig. 8. Some transitions in Res_r .

the nameserver with the IP address $\text{val}(rr)$, which is assumed to have the domain name $\text{dn}(q_1)$ (using the premise). After the transition, the resolver moves to the status *wait* as it waits for a response from the network, reflected by the state $\langle \text{wait}, sl \rangle$.

Intuitively, this transition is taken when the resolver is about to successfully resolve the latest *subquery* at the top of the stack required to resolve the penultimate (sub)query and seemingly moves one step closer to the resolution of the original query (i.e., reduces the stack depth by one).

Transition ②. Prior to this transition, the resolver configuration is similar to the previous one. However, this time the premise implies that the query (or subquery) q cannot be resolved or rewritten using the cache. To resolve any further, the resolver picks the leftmost domain $ns = sl[\min_{sl[i] < \text{dn}(q)} i]$ in sl which is a prefix of the current query name. The next step is to derive the IP for ns from the cache. Consequently, the resolver pushes (ns, Λ) onto the query stack in order to find the IP address of the nameserver with the domain name ns .

Transition ③. In this transition, the resolver is waiting for a response from a nameserver and receives a referral response accompanied by a set of resource records R . The resolver then adds the records in R to its cache, instantiating a timer for each element and setting the timer to the corresponding time-to-live. Additionally, the resolver filters out the NS records from R , represented by $\text{Val}(R)$, and for each such record, say rr , makes the server list entry $\text{val}(rr)$, dictated by the priority order in $list$; $\text{Lin}(\text{Val}(R))$ computes the set of all possible priorities. Since the resolution of the original query is not yet complete, the resolver updates its status to *query*.

Note that each of the above transitions only checks for the existence of a record in the resolver's cache. Hence, one can force the transition to execute without any elapse of time, effectively making the cache elements seemingly persistent. However, DNS behaves statefully upon the *absence* of records as well. For the sake of clarity and ease of presentation, the set of all resolver transitions has been moved to our technical report [43, Appendix A.2].

Remark 1. Real-world nameservers are often caching nameservers [39, 40], i.e., they are equipped with limited resolution capabilities (they are unable to perform subquery resolution) and caching capabilities (to store answers from resolved queries). Resolvers are sometimes configured to cache negative answers with the NXDOMAIN label. In this paper, we do not assume such systems to simplify our presentation. However, this is not a restriction of our model as our results hold even otherwise.

3.6 DNS is Eager

We conclude this section by showing that the formal model DNS is eager.

LEMMA 3.7. *DNS is eager.*

PROOF SKETCH. The proof consists of two steps. First, we show that every reachable configuration in DNS contains exactly one process capable of either sending or receiving a message. Second, we show that once a process sends (or receives) a message, it transitions to a state where it can only receive (or send) messages. Now, since the runs in DNS are interleavings of local runs of processes, every transition is executed by *exactly* one process. Hence, any send action must be immediately followed by a receive action.

The first step follows by induction on the length of the runs. Observe that the initial configuration only allows the resolver process to send a message, thus covering the base case. The second step can be shown with a case-by-case analysis of the transitions (see [43, Appendix A] for the full set of transitions). \square

Hence, considering an arbitrary run in DNS is equivalent to considering an eager run. We will use the above lemma, together with Lemma 3.6, to reduce the DNSVERIF problem to the LTS verification problem for a pushdown system. In Section 4 we prove that DNS can in fact be reduced to an rCPS with an *equivalent* set of traces.

4 Property-Agnostic Abstractions

We have so far established that DNS can be modelled as a trCPS. The set of its global control states $(\text{NSCONFIG}_s)_{s \in \mathcal{S}} \times (\text{RESSTAT}_r)_{r \in \mathcal{R}}$, as well as the message alphabet \mathcal{M} , are infinite. A priori, this prevents us from using any reachability results from Section 3.2.2. In particular, this is caused by three main factors: (1) an infinite query space, (2) clock valuations, and (3) an unbounded stack.

In this section, we consider the first two factors and develop ways to bound them, consequently reducing the DNS to an rCPS. Broadly speaking, we will use finite monoids as an abstraction for the query space. To abstract time, we will use a variant of the region abstraction method proposed in [10]. In the next section, we will formally define the DNSVERIF problem and derive a decision procedure for it.

Let us revisit some basic concepts from algebraic language theory. We first fix a finite alphabet Σ .

Definition 4.1 (Syntactic Congruence). Given a regular language $L \subseteq \Sigma^*$, the *syntactic congruence* with respect to L is the coarsest congruence relation \sim_L on Σ^* such that,

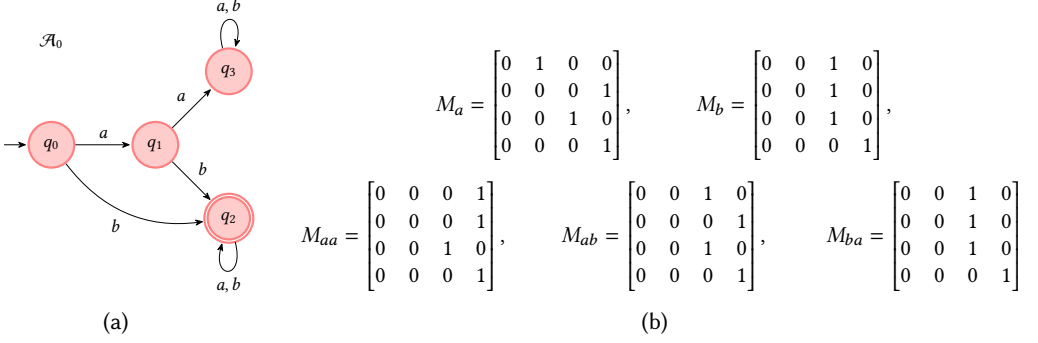
$$u \sim_L v \iff \forall x, y \in \Sigma^*. (xuy \in L \iff xvy \in L).$$

Observe that, if $u_1 \sim_L v_1$ and $u_2 \sim_L v_2$, then $u_1 \cdot u_2 \sim_L v_1 \cdot v_2$. Consequently, the quotient structure $\langle L / \sim_L, \cdot \rangle$ forms a finite monoid [42, 46], called the *syntactic monoid* of L . We denote it by $\text{SYN}(L)$. This induces a canonical surjective homomorphism from the free monoid Σ^* to $\text{SYN}(L)$, called the *syntactic homomorphism* with respect to L , denoted by η_L . The syntactic monoid of a language can be effectively constructed [33] as the transition monoid of its minimal automaton [1].

4.1 Positive Prefix-Testable Languages

A regular language L is *prefix-testable* [9, 54], if membership of L can be tested by inspecting prefixes up to some length, namely if L is a finite Boolean combination of languages of the form $u\Sigma^*$, for a finite word u . We will use the *positive* fragment of prefix-testable languages to construct a sound and complete abstraction for the set of domain names DOMAIN . A *positive prefix-testable* (PPT) language is a finite combination of languages of the form $u\Sigma^*$, using only the union and intersection operations.⁶

⁶The class of positive prefix-testable languages is an old player in algebraic language theory, and is a simple example of what is known as a positive variety of regular languages [45]. Syntactic congruence for such languages corresponds to the class of ordered monoids. However, unordered monoids suffice for our purposes.

Fig. 9. Construction of the syntactic monoid for L_0 .

Example 4.2. Given $\Sigma = \{a, b\}$, the language $L_0 = ab\Sigma^* \cup b\Sigma^*$ is a PPT language. Figure 9a shows the minimal automaton \mathcal{A}_0 for the language; Figure 9b shows the elements in the syntactic monoid $\text{SYN}(L_0)$, i.e., the transition monoid of \mathcal{A}_0 . This can be constructed by first computing the adjacency matrices M_a and M_b that correspond to the a - and b -labelled edges in the automaton's underlying graph. We then close them under matrix multiplication, such that the transition matrix for the word w is denoted by M_w . For instance, the multiplication $M_{aa} \cdot M_{ab}$ yields the matrix M_{aa} again. Since $M_{bb} = M_b \cdot M_b = M_b$, it has not been included in the set.

Notice that in the above example, the transition matrix for any word of length ≥ 3 is identical to the transition matrix for some word of length ≤ 2 . Below we formally state this property.

LEMMA 4.3. [9, §3.7.2] *Given a PPT language $L = \bigcup_{j \in J} w_j \Sigma^*$, for a finite set of indices J , there exists a natural number $k \in \mathbb{N}$, such that the following hold:*

- (1) *the set $\eta_L(\Sigma^k)$ comprises of all the idempotents in $\text{SYN}(L)$, and*
- (2) *all the idempotents in $\text{SYN}(L)$ are left-zeroes*

Moreover, $k = \max_{j \in J} |w_j|$ can be computed.

Intuitively, this means that the transitive closure of the transition relation, from any state in the corresponding automaton, is dictated by the *first* k positions of the input word.

Example 4.4. In the transition monoid $\text{SYN}(L_0)$ from Example 4.2 (see also Figure 9), the idempotents are the elements $M_b = M_{bb}$, M_{aa} , M_{ab} , and M_{ba} . These idempotents also act as left-zeroes. Indeed, from any state, after reading the first two letters of the input word, the automaton reaches one of the sink states q_2 and q_3 .

We will now prove a crucial property specific to a subclass of PPT languages that forms the cornerstone of our domain abstraction in the following sections.

LEMMA 4.5. *Consider the PPT language L of the form $p \cdot \Sigma^+$ or $p \cdot \Sigma^{\geq 2}$, for some $p \in \Sigma^*$. If $u \preceq p$ is a prefix of p , then $u \sim_L u' \implies u = u'$. Moreover, if L is of the form $p \cdot \Sigma^{\geq 2}$, and $u \in p \cdot \Sigma$, then $u \sim_L u' \implies u' \in p \cdot \Sigma$.*

PROOF. For the first part, assume that L is of the form $p \cdot \Sigma^+$; the proof for the other case is similar. Since $u \preceq p$, for some letter a , the string $uu^{-1}pa \in L$. Moreover, since $u \sim_L u'$, $u'(u^{-1}pa) \in L$. We now consider the size of u' and do a case-by-case analysis:

$|u'| > |p|$. It must be that $p \prec u'$, and hence $u' \in L$. But, since $u \sim_L v$ we have $u \in L$, which implies that $p \prec u$. This leads to a contradiction.

$|u'| \leq |p|$. This implies that $u' \preceq p$. Now we have $u \preceq p$, $u' \preceq p$ and $u \sim_L u'$. If $u = u'$, we are done. Otherwise, either $u \prec u'$ or $u' \prec u$. Without loss of generality, assume $u \prec u'$. Then there exists w s.t. $u'w \in p\Sigma \subseteq p\Sigma^+$. Since $u \sim_L u'$, we must have that $uw \in L$. But since $u \prec u'$, $|u| < |u'|$, which implies that $|uw| < |u'w| = |p| + 1$. This leads to a contradiction.

For the second part, since $u \in p \cdot \Sigma$, there exists $a \in \Sigma$ such that $u = p \cdot a$. Moreover, there exists $x \in \Sigma^*$ such that $p \cdot a \cdot x \in L$. Since $u \sim_L u'$, $u' \cdot x \in L$. We now consider the size of u' and do a case-by-case analysis:

$|u'| > |p| + 1$. There must exist $b, c \in \Sigma$ and $y \in \Sigma^*$ such that $u' = p \cdot b \cdot c \cdot y$. Using Lemma 4.3, we know that $\eta_L(p \cdot b \cdot c)$ is a left-zero in $\text{SYN}(L)$. Hence, $u' \sim_L p \cdot b \cdot c$. By transitivity, we thus obtain $p \cdot a \sim_L p \cdot b \cdot c$. But clearly, for the empty string ϵ , $p \cdot a \cdot \epsilon \notin L$ but $p \cdot b \cdot c \cdot \epsilon \in L$. This leads to a contradiction.

$|u'| < |p| + 1$. This means that $u' \preceq p$. From the first part, we then know that $u = u'$, but by assumption $|u| = |p| + 1$. This leads to a contradiction.

Thus, it must be that $|u'| = |p| + 1$, and hence, $u' \in p \cdot \Sigma$. \square

The above lemma shows that for a language L of the form $p \cdot \Sigma^+$ (and $p \cdot \Sigma^{\geq 2}$), the set $\eta_L^{-1}(\eta_L(u))$ is a singleton, for $u \preceq p$, i.e., all prefixes of p are syntactically congruent only to themselves, with respect to L . Moreover, for $L = p \cdot \Sigma^+$ (or $p \cdot \Sigma^{\geq 2}$) and $u \preceq p$, $\eta_L^{-1}(\text{prec}(\eta_L(u)))$ is the set of prefixes of u .

Before we define our abstraction on the query space, we will first introduce a more generalized notion of bisimulation for systems with LTS semantics. We will then show that our abstraction induces this generalized notion of bisimulation on DNS.

4.2 A Generalized Bisimulation

Recall the definition of an LTS from Section 3.2. Here we define a novel generalization of bisimulation on LTSes, called a *generalized kernel bisimulation*. In contrast to standard notions of bisimulation, a generalized kernel bisimulation consists of equivalence relations on both states *and* transitions such that two equivalent states are *bisimilar* up to their equivalent outgoing transitions. In order to formalize this, we adjust Glushkov's definition of Mealy automata homomorphisms [23, p. 3] to define homomorphisms for LTS.

Definition 4.6 (LTS Homomorphism). An LTS homomorphism of an LTS $\mathcal{A}^1 = \langle S^1, S_{init}^1, A^1, \text{TRANS}^1 \rangle$ onto an LTS $\mathcal{A}^2 = \langle S^2, S_{init}^2, A^2, \text{TRANS}^2 \rangle$ is a product map $H = \langle h_{state}, h_{act} \rangle$ of the surjective maps $h_{state} : S^1 \twoheadrightarrow S^2$ and $h_{act} : A^1 \twoheadrightarrow A^2$, such that the following formulae are valid for any $s^1 \in S^1$ and $a^1 \in A^1$,

$$\begin{aligned} \forall s'^1. (s^1 \xrightarrow{a^1} s'^1 \in \text{TRANS}^1 &\implies h_{state}(s^1) \xrightarrow{h_{act}(a^1)} h_{state}(s'^1) \in \text{TRANS}^2) \\ s^1 \in S_{init}^1 &\implies h_{state}(s^1) \in S_{init}^2 \end{aligned}$$

We write $H : \mathcal{A}^1 \rightarrow \mathcal{A}^2$ to denote the homomorphism. The set of all homomorphisms is closed under composition. The identity homomorphism is the identity map $id_{\mathcal{A}^1} = \langle id_{S^1}, id_{A^1} \rangle$, where id_{S^1} and id_{A^1} are the identity maps on S^1 and A^1 respectively.

Given a run $\rho = s_0 \xrightarrow{a_1} s_1 \dots s_{n-1} \xrightarrow{a_n} s_n$ in \mathcal{A}^1 , we write $H(\rho)$ to denote its pointwise homomorphic image $h_{state}(s_0) \xrightarrow{h_{act}(a_1)} h_{state}(s_1) \dots h_{state}(s_{n-1}) \xrightarrow{h_{act}(a_n)} h_{state}(s_n)$, which is a run in \mathcal{A}^2 .

Essentially, a homomorphism from an LTS \mathcal{A}^1 to \mathcal{A}^2 induces an equivalence relation on the set of states and actions. Indeed, we say two states (actions) are equivalent if they have the same image with respect to the function h_{state} (h_{act}). Alternatively, s^1 (a^1) is equivalent to s'^1 (a'^1) iff

$\langle s^1, s'^1 \rangle (\langle a^1, a'^1 \rangle)$ is in the *kernel* of the homomorphism h_{state} (h_{act}), i.e., $\langle s^1, s'^1 \rangle \in \ker(h_{state})$ ($\langle a^1, a'^1 \rangle \in \ker(h_{act})$), where $\ker(f)$ denotes the kernel of a function f . Consequently, such an equivalence induces a congruence on the structure \mathcal{A}^1 . We now define a generalized kernel bisimulation on an LTS.

Definition 4.7 (Generalized Kernel Bisimulation). Given an LTS $\mathcal{A}^1 = \langle S^1, S^1_{init}, A^1, \text{TRANS}^1 \rangle$, a homomorphism H of \mathcal{A}^1 to $\mathcal{A}^2 = \langle S^2, S^2_{init}, A^2, \text{TRANS}^2 \rangle$ is called a *generalized kernel bisimulation* on \mathcal{A}^1 if and only if, for every tuple $\langle p, p' \rangle \in \ker(h_{state})$ and all actions $a \in A^1$, if $p \xrightarrow{a} q \in \text{TRANS}^1$, then $p' \xrightarrow{a} q' \in \text{TRANS}^1$, such that $\langle a, a' \rangle \in \ker(h_{act})$ and $\langle q, q' \rangle \in \ker(h_{state})$.

Given two states p and p' in S^1 , we say p is *kernel bisimilar* to p' with respect to H , written $p \sim_H p'$, if H is a generalized kernel bisimulation on \mathcal{A}^1 , such that $\langle p, p' \rangle \in \ker(h_{state})$.

If the generalized kernel bisimulation H is such that $h_{act} = id_{A^1}$ is the identity map on A^1 , then $\ker(h_{state})$ is a strong bisimulation [43, Appendix B]. Conversely, for any strong bisimulation $R \subseteq S^1 \times S^1$ on \mathcal{A}^1 , we define the map $H^R = \langle h^R_{state}, id_{A^1} \rangle$, where $h^R_{state} : S^1 \rightarrow S^1/R$ is the quotient map with respect to the equivalence relation R . The map H^R is indeed a homomorphism of \mathcal{A}^1 onto the LTS $\mathcal{A}'^2 = \langle S^1/R, S^1_{init}/R, A^1, \text{TRANS}'^1 \rangle$ where $\text{TRANS}'^1 = \{h^R_{state}(p) \xrightarrow{a} h^R_{state}(q) \mid p \xrightarrow{a} q \in \text{TRANS}^1\}$. Moreover, for any $a \in A^1$,

$$\begin{aligned} p \sim_{HR} p' &\Leftrightarrow \langle p, p' \rangle \in R \\ &\Rightarrow (p \xrightarrow{a} q \in \text{TRANS}^1 \Rightarrow \exists q'. p' \xrightarrow{a} q' \in \text{TRANS}^1 \wedge \langle q, q' \rangle \in R) \\ &\Leftrightarrow (p \xrightarrow{a} q \in \text{TRANS}^1 \Rightarrow \exists q'. p' \xrightarrow{a} q' \in \text{TRANS}^1 \\ &\quad \wedge \langle a, a \rangle \in \ker(id_{A^1}) \wedge q \sim_{HR} q'). \end{aligned}$$

Hence, generalized kernel bisimulation is weaker than strong bisimulation. Below, we state as a corollary, that any two states that are generalized kernel bisimilar have equivalent sets of traces starting from them, upto their homomorphic image.

LEMMA 4.8. *Consider an LTS $\mathcal{A}^1 = \langle S^1, S^1_{init}, A^1, \text{TRANS}^1 \rangle$ and a kernel bisimulation $H : \mathcal{A}^1 \rightarrow \mathcal{A}^2$. Then for any run ρ in \mathcal{A}^2 from a state s'^2 , there exists a run ρ' in \mathcal{A}^1 from s^1 , for each $s^1 \in h^{-1}_{state}(s'^2)$, such that $H(\rho') = \rho$.*

PROOF. The proof follows by induction, on the length of runs in \mathcal{A}^2 . The base case for runs of length 1 is immediate from the definition of generalized kernel bisimulation. \square

Note that kernel bisimulations have already been introduced and well-studied in the context of category theory [24, 50]. Moreover, it is well-known that they coincide with strong bisimulations in the context of LTSes. Thus, generalized kernel bisimulations essentially generalize kernel bisimulations in the context of LTSes. Hereafter, for brevity, we will refer to generalized kernel bisimulations as just kernel bisimulations.

4.3 Syntactic Domain Congruence

Recall that every query consists of a domain name and a query type. Since there are finitely many query types, the unbounded nature of `QUERY` arises from the unbounded⁷ cardinality of `DOMAIN`. This is because every domain name consists of an unbounded sequence of labels, where each label

⁷In practice, it is bounded by hardware or pre-defined limits. However, the limits are extremely large. By current standards, the total number of possible queries is $\sum_{i=0}^{253} 38^i$, where 253 is the maximum length of a domain name, and each position is labelled by one of 38 valid characters. The total number of possible labels is $\sum_{i=0}^{253} 37^i$, as a label is a domain without any periods (`.`).

comes from the set LABEL (Example 3.1). In this section, we define a more robust notion of a syntactic congruence over DOMAIN with finitely many equivalence classes that preserves the semantics of DNS. In particular, the congruence is preserved not just under concatenation but also domain name rewrites. We shall show that every congruence class is a PPT language over LABEL.

As usual, we fix the zone-file configuration \mathcal{Z} for all our computations. Since every domain name is also a string, henceforth, we call a set of domain names a language. To begin with, we build some intuition on how one may abstract the domain names to a finite set.

4.3.1 Basic Intuition. As a first approach, we will try to define an equivalence relation on DOMAIN with finitely many equivalence classes, such that any two global configurations in DNS that are equivalent upto their query components, are kernel bisimilar. To do that, we first identify some languages by observing the behaviour of the system,

$$\begin{aligned} L_{s,t,\times} &= \{d \mid d \in \text{DOMAIN}, \text{dn}(\mathcal{M}(s, \langle d, t \rangle)) \times d\} \\ L_{t,rr} &= \{d \mid d \in \text{DOMAIN}, d \leftarrow_t rr\} \end{aligned}$$

for $s \in \mathcal{S}, t \in \text{TYPE}, \times \in \{=, \ni, <\},$ and $rr \in \mathcal{Z}$. Any two domain names d and d' with identical membership relations with the above languages are defined to be equivalent, i.e., d and d' are equivalent if for each of the above languages, say $L, d \in L$ iff $d' \in L$. Intuitively, two global configurations in DNS where the corresponding occurrence of domain names are equivalent (we exclude the unsynthesized resource records from this comparison) according to the above relation will either both satisfy the premise for a transition or neither, for example, observe the premises of transitions in Figures 6, 7, and 8. Since the above set of languages is finite, there are finitely many equivalence classes. Readers may check that this also induces a kernel bisimilarity on the global configuration space of DNS.

Note that this also suffices to ensure kernel bisimilarity in the case of DNAME rewrites. Indeed, by definition, for two equivalent domain names d and d' , and a DNAME record rr , either they both contain $\text{dn}(rr)$ as a prefix or neither. This means, either both names can undergo the rewrite operation or neither. After the rewrite, both the resulting domain names will have the same prefix up to all the domain names of the records in \mathcal{Z} . However, a priori it is not clear whether rewrites also preserve the equivalence of the domain names, as the membership in the languages of the form $L_{s,t,\times}$ may now differ. Moreover, a more fundamental concern is that there is no obvious way to effectively represent the above languages.

4.3.2 Formal Definition. We will now define a syntactic congruence on domain names that slightly refines the one sketched above. To begin with, let us denote the set of all labels occurring in the domain names in the records of \mathcal{Z} by $\Sigma_{\mathcal{Z}} \subseteq \text{LABEL}$. Additionally, we symbolically denote the elements in the complement set $(\text{LABEL} \setminus \Sigma_{\mathcal{Z}})$ by the special label *other*. This finally gives us the finite alphabet, $\Sigma_{\text{DNS}} = \Sigma_{\mathcal{Z}} \cup \{\text{other}\}$.

Definition 4.9 (Syntactic Domain Congruence). For every record $rr \in \mathcal{Z}$, consider the PPT languages $L_{rr,\text{dn}}$ and $L_{rr,\text{val}}$, such that

$$L_{rr,\text{dn}} = \begin{cases} \text{dn}(rr). \Sigma_{\text{DNS}}^{\geq 2}, & \text{if } \text{dn}(rr) \in (\text{LABEL}^*)_{\star} \\ \text{dn}(rr). \Sigma_{\text{DNS}}^+, & \text{if } \text{dn}(rr) \in (\text{LABEL}^*) \end{cases} \quad L_{rr,\text{val}} = \begin{cases} \text{val}(rr). \Sigma_{\text{DNS}}^+, & \text{if } \text{ty}(rr) = \text{NS} \\ \Sigma_{\text{DNS}}^*, & \text{if } \text{ty}(rr) \neq \text{NS}. \end{cases}$$

The *syntactic domain congruence* with respect to \mathcal{Z} is the congruence relation $\sim_{\mathcal{Z}}$ on Σ_{DNS}^* s.t.

$$u \sim_{\mathcal{Z}} v \iff \forall rr \in \mathcal{Z}. (\eta_{L_{rr,\text{dn}}}(u) = \eta_{L_{rr,\text{dn}}}(v) \wedge \eta_{L_{rr,\text{val}}}(u) = \eta_{L_{rr,\text{val}}}(v)).$$

We then say that the domains u and v are *syntactically congruent* (with respect to \mathcal{Z}).

Equivalently, $u \sim_{\mathcal{Z}} v$ iff $\eta_{\mathcal{Z}}(u) = \eta_{\mathcal{Z}}(v)$, where $\eta_{\mathcal{Z}}(w) = \langle \eta_{L_{rr,j}}(w) \rangle_{rr \in \mathcal{Z}, j \in \{\text{dn}, \text{val}\}}$. The co-domain of $\eta_{\mathcal{Z}}$ is the product monoid $\prod_{rr \in \mathcal{Z}, j \in \{\text{dn}, \text{val}\}} \text{SYN}(L_{rr,j})$, denoted by $\text{SYN}_{\mathcal{Z}}$.

Intuitively, the idea is to define a congruence that is sufficiently refined such that two domain names satisfy the conditions required to trigger a transition in DNS if and only if they are congruent to each other; the languages must be carefully chosen to make this happen. Below we show that this is indeed the case. In particular, we show that the syntactic domain congruence refines the monadic relations, defined as languages, from our discussion above.

LEMMA 4.10. *Consider two domain names d and d' such that $d \sim_{\mathcal{Z}} d'$, then the following hold:*

- (1) $\forall rr \in \mathcal{Z}. (\text{dn}(rr) \times d \iff \text{dn}(rr) \times d')$, where $\times \in \{=, \ni_{\star}, <\}$.
- (2) $\forall t \in \text{TYPE.RANK}(rr, \langle d, t \rangle, z) = \text{RANK}(rr, \langle d', t \rangle, z)$, where z is an arbitrary zone in \mathcal{Z} and rr is an arbitrary record in z .
- (3) $\forall rr \in \mathcal{Z}. \text{ty}(rr) = \text{NS} \implies (\text{val}(rr) = d \iff \text{val}(rr) = d')$.

PROOF. (1) Without loss of generality, we shall do a case-by-case analysis of the domain name d , generalizing $rr \in \mathcal{Z}$,

($\times \equiv =$). Trivially, $d = \text{dn}(rr)$ implies $d \preceq \text{dn}(rr)$. Since $d \sim_{\mathcal{Z}} d'$, then $d \sim_{L_{rr,\text{dn}}} d'$. Using Lemma 4.5, we thus have $d = d'$, which implies $d' = \text{dn}(rr)$.

($\times \equiv \ni_{\star}$). Recalling the definition of \ni_{\star} from Figure 3, we know that $d_{\star} \in \text{dn}(rr) \otimes (\Sigma_{\text{DNS}})_{\star}$ and $\text{dn}(rr) \in (\text{LABEL}^*)_{\star}$. By definition, $L_{rr,\text{dn}} = \text{dn}(rr) \cdot \Sigma_{\text{DNS}}^{\geq 2}$. Since $d \sim_{L_{rr,\text{dn}}} d'$, using Lemma 4.5, we again have that $d' \in \text{dn}(rr) \cdot (\Sigma_{\text{DNS}})$, or equivalently, $d'_{\star} \in \text{dn}(rr) \otimes (\Sigma_{\text{DNS}})_{\star}$ (recall the definition of \mathcal{L}_{\star} from Section 3.1).

($\times \equiv <$). Recalling the definition of $<$ from Figure 3, $\text{dn}(rr) \subset \text{prec}(d)$. Since the underlying monoid is a free monoid, this is equivalent to $\text{dn}(rr) \prec d$. Using Lemma 4.3, $\eta_{L_{rr,\text{dn}}}(d)$ is thus a left-zero in $\eta_{L_{rr,\text{dn}}}$, which implies that $\eta_{L_{rr,\text{dn}}}(d')$ is also a left-zero in that semigroup. Hence, using Lemma 4.5, $\text{dn}(rr) \prec d'$, or equivalently, $\text{dn}(rr) \subset \text{prec}(d')$.

(2) It suffices to prove that $\text{MATCH}(rr, \text{dn}(rr)) = \text{MATCH}(rr, d) \iff \text{MATCH}(rr, \text{dn}(rr)) = \text{MATCH}(rr, d')$ and $\text{max}_{\succeq}(\text{dn}(rr), d) = \text{max}_{\succeq}(\text{dn}(rr), d')$. The first formula follows from the arguments in (1).

To prove the second formula, we must prove that $|\text{prec}(\text{dn}(rr)) \cap \text{prec}(d)| = |\text{prec}(\text{dn}(rr)) \cap \text{prec}(d')|$. Since we are working with the free monoid \mathcal{L} , $\text{prec}(\text{dn}(rr)) \cap \text{prec}(d)$ essentially contains the prefix-closed set of the largest common prefix of $\text{dn}(rr)$ and d . We will show that any common prefix of $\text{dn}(rr)$ and d must be a prefix of d' .

Suppose d_1 is a common prefix of $\text{dn}(rr)$ and d . Then there must exist an $m \in \text{SYN}(L_{rr,\text{dn}})$ such that, $\eta_{L_{rr,\text{dn}}}(d_1) \cdot m = \eta_{L_{rr,\text{dn}}}(d) = \eta_{L_{rr,\text{dn}}}(d')$. Hence, there is a prefix d_2 , of d' , in $\eta_{L_{rr,\text{dn}}}^{-1}(d_1)$. However, since d_1 is a prefix of $\text{dn}(rr)$, it follows from Lemma 4.5, that $d_1 = d_2$. Thus, d_1 is a prefix of d' as well.

(3) This follows immediately from the definitions and Lemma 4.5 using similar arguments as before. \square

The following theorem is a direct corollary of the above lemma,

THEOREM 4.11. *Consider two domain names d and d' such that $d \sim_{\mathcal{Z}} d'$, then for every $s \in \mathcal{S}$, $t \in \text{TYPE}$, $\times \in \{=, \ni_{\star}, <\}$ and $rr \in \mathcal{Z}$, it holds that*

- (1) $d \in L_{s,t,\times} \iff d' \in L_{s,t,\times}$
- (2) $d \in L_{t,rr} \iff d' \in L_{t,rr}$.

Below, we state that the syntactic domain congruence also commutes with DNAME rewrites.

LEMMA 4.12. Consider a *DNAME* resource record $\langle pre, \text{DNAME}, \sigma, post, b \rangle$, and two syntactically congruent domain names d and d' . Then $(post.(pre)^{-1}.d) \sim_{\mathcal{Z}} (post.(pre)^{-1}.d')$.

PROOF. It suffices to show that $pre \preceq d$ iff $pre \preceq d'$. Since, $pre = \text{dn}(rr)$, for some resource record rr , this follows immediately from Lemma 4.10 (1). \square

The congruence can be naturally extended to a (non-symbolic) domain name $d \in \text{DOMAIN}$, where any label in $(\text{LABEL} \setminus \Sigma_{\mathcal{Z}})$ is mapped to $\eta_{\mathcal{Z}}$ (*other*). Additionally, the root domain ϵ is mapped to $\eta_{\mathcal{Z}}(\epsilon)$. Thus, we obtain a monoid homomorphism of the free monoid \mathcal{L} . We can effectively construct the equivalence classes of the entire set of domain names occurring in \mathcal{Z} , as shown in the following lemma.

LEMMA 4.13. Assume that \mathcal{Z} has n resource records and that the maximum length of the domain name of a resource record in \mathcal{Z} is c . Then the class of abstract domains can be computed in $O(n^{c+1}c^5)$ time.

PROOF. Each resource record in the zone file configuration corresponds to at most two PPT languages, $L_{rr, \text{dn}}$ and $L_{rr, \text{val}}$. There are at most $O(n)$ such languages. Due to the nature of these languages, one can show that to compute the corresponding automata we require $O(|\text{dn}(rr)|)$ time. The adjacency matrix for each letter in the alphabet Σ_{DNS} can be initiated in $O(|\Sigma_{\text{DNS}}| |\text{dn}(rr)|^2)$ time. To compute adjacency matrices, we thus need $O(|\Sigma_{\text{DNS}}| |\text{dn}(rr)|^2)$ time.

To compute the set of abstract domains, we take the diagonal product over all the languages, for each label, and then compute all concatenations up to the length c . Since there are at most $O(n)$ matrices for each label (since there are at most that many languages), we will compute $|\Sigma_{\text{DNS}}|^c nc$ matrix multiplications in all. Each multiplication takes $O(c^3)$ time because the dimensions of a matrix are at most $c \times c$. Therefore, to compute all the domains we need $O(|\Sigma_{\text{DNS}}|^c nc^4)$ time.

However, since $|\Sigma_{\text{DNS}}| \leq nc$, we thus have $O(|\Sigma_{\text{DNS}}|^c nc^4) = O(n^{c+1}c^5)$. \square

4.4 Abstract DNS Model

Recalling Example 3.1, we have defined an *abstraction* on $\text{DOMAIN}_{\mathcal{L}}$ with respect to $\mathcal{Z}_{\mathcal{L}}$, induced by a congruence on Σ_{DNS}^* such that the quotient structure forms a monoid. The abstraction can be computed using the monoid homomorphism $\eta_{\mathcal{Z}}$.

Using Lemma 3.2, the homomorphism $\eta_{\mathcal{Z}}$ can be functorially mapped to the map $Z : \text{ZFC}_{\mathcal{L}} \rightarrow \text{ZFC}_{\text{SYN}_{\mathcal{Z}}}$. We also observe that, for a zone file configuration $\mathcal{Z}' \in \text{ZFC}_{\mathcal{M}}$ over the monoid \mathcal{M} , the semantics of the DNS instance $\langle \text{DNS}_{\mathcal{Z}'} \rangle$ is functorial in \mathcal{M} . This is a straightforward extension of Lemma 3.2, following the constructions described in Sections 3.4 and 3.5. Now, the arbitrarily fixed instance DNS , used in Sections 3.4 and 3.5, is defined over the monoid \mathcal{L} . Hence, the homomorphism $\eta_{\mathcal{Z}} : \mathcal{L} \rightarrow \text{SYN}_{\mathcal{Z}}$ can be functorially mapped to an LTS homomorphism, $H_{\eta_{\mathcal{Z}}} : \langle \text{DNS}_{\mathcal{Z}} \rangle \rightarrow \langle \text{DNS}_{\mathcal{Z}(\mathcal{Z})} \rangle$. We call $\text{DNS}_{\mathcal{Z}(\mathcal{Z})}$ the *abstraction* of the concrete model $\text{DNS}_{\mathcal{Z}}$.

Figure 10 shows the *abstract* nameserver transitions in $\text{DNS}_{\mathcal{Z}(\mathcal{Z})}$ from Figures 6 and 7. Intuitively, in all the states and transitions, every occurrence of a domain name is replaced by its abstraction.

We conclude this section with one of the main contributions of this paper, establishing both the soundness and completeness of our abstraction.

THEOREM 4.14 (SOUNDNESS AND COMPLETENESS). For every run ρ in $\langle \text{DNS}_{\mathcal{Z}(\mathcal{Z})} \rangle$, there exists a run ρ' in $\langle \text{DNS}_{\mathcal{Z}} \rangle$ such that $H_{\eta_{\mathcal{Z}}}(\rho') = \rho$. The converse holds as well.

PROOF. It suffices to show that the LTS homomorphism $H_{\eta_{\mathcal{Z}}}$ is a kernel bisimulation. The statement then follows from Lemma 4.8. \square

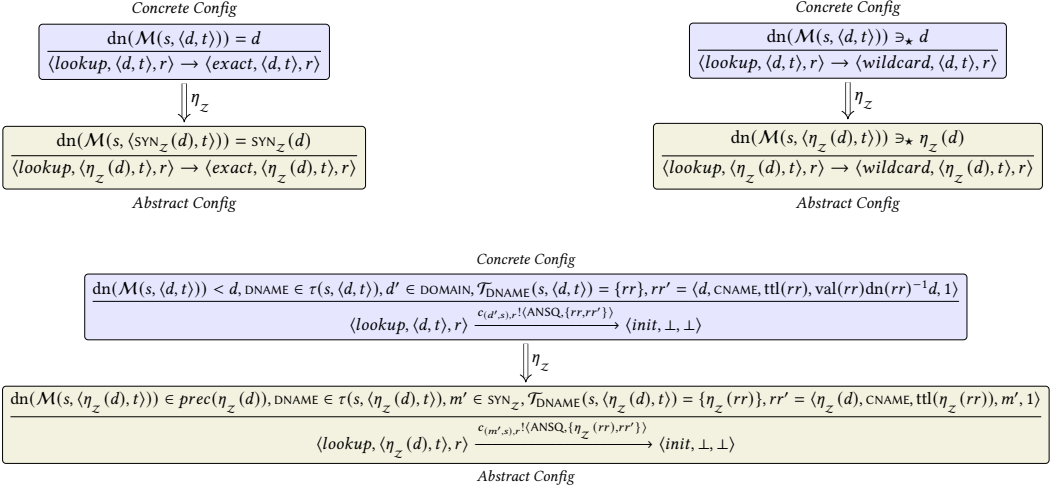


Fig. 10. Some abstract NS_s transitions; the concrete configurations are written before the arrows and the abstract ones after. For the abstract configuration in the bottom-most transition, we assume $m' \in \text{val}(\eta_z(rr))\text{dn}(\eta_z(rr))^{-1}\eta_z(d)$.

This theorem shows that the set of runs in the abstraction $\text{DNS}_{Z(\mathcal{Z})}$ is neither an over-approximation nor an under-approximation, but captures the exact set of runs in DNS_Z , up to the homomorphic image with respect to H_{η_z} . Using Lemma 3.7, we further have the following corollary.

COROLLARY 4.15. $\text{DNS}_{Z(\mathcal{Z})}$ is eager.

With this, we have obtained a sound and complete abstraction of DNS in the form of a trCPS with finitely many global states. In the rest of this paper, runs refer to the runs in the abstraction $\text{DNS}_{Z(\mathcal{Z})}$, unless otherwise stated. As the next step, we reduce the semantics of $\text{DNS}_{Z(\mathcal{Z})}$ to its *untimed model*.

Abstracting Time. The real-time semantics of the model makes the global configuration space infinite. Since the timer semantics in our model only involve resetting the timer ($[x \leftarrow k]$), by checking the expiration of the timer ($x > 0$) and devolution of time ($\overset{t}{\leftarrow}$), this immediately allows us to employ a standard region construction [10] to our model in order to obtain an untimed symbolic reduction of $\text{DNS}_{Z(\mathcal{Z})}$.

As the reachability relation of configurations is preserved under region construction, Theorem 4.14 also holds for the reduced untimed model. By abuse of notation, we refer to this model as $\text{DNS}_{Z(\mathcal{Z})}$.

5 The DNS Verification Problem

In this section, we study the decidability of the DNS verification problem for DNS_Z . Borrowing notation from Section 3.3, we formally define our problem DNSVERIF as follows,

Input. A ZFC $\langle \mathcal{Z}, \Theta \rangle \in \text{ZFC}_{\mathcal{Z}}$ over the free monoid \mathcal{L} , and a $\text{FO}(<)$ property ϕ over strings in \mathbb{A}^* , given as the corresponding minimal automaton \mathcal{F}^ϕ , where $\mathbb{A} = \text{ACTION}_r \uplus \biguplus_{s \in \mathcal{S}} \text{SACTION}_s$.

Output. The decision of the verification problem for $\langle \text{DNS}_Z \rangle$, and $\mathcal{F}^{\eta_Z(\phi)}$ over $\eta_Z(A)^*$, along with a witness trace, where $\mathcal{F}^{\eta_Z(\phi)}$ is an FA obtained from \mathcal{F}^ϕ by mapping every transition label $a \in A$ to $\eta_Z(a)$. We denote the corresponding problem instance by $\text{DNSVERIF}(Z, \Theta, \phi)$.

Note that the DNSVERIF problem does not target the correctness of the DNS protocol itself, but rather an instantiation of the protocol for a given zone configuration. The problem aims to reflect the situation where network administrators already maintain a zone file configuration and would like to find vulnerabilities in the resulting DNS behaviour. Intuitively, we want to capture security vulnerabilities as traces of runs in DNS_Z characterized by $\text{FO}(\prec)$ -definable languages. We will look at some common attack vectors in Section 6.

5.1 The Verification Approach

Recall from Section 3.3 that the untimed model $\text{DNS}_{Z(Z)}$ is an rCPS over the pointed topology T_Z . Using Lemma 3.6 and Corollary 4.15, we can construct, in polynomial time, a pushdown system \mathcal{P}_Z , such that for any run in $\text{DNS}_{Z(Z)}$ from an s -configuration to an s' -configuration, there is a trace equivalent run in \mathcal{P}_Z , also from an s -configuration to an s' -configuration. It then immediately follows from Theorem 4.14 that:

LEMMA 5.1. $\text{VERIF}(Z, \Theta, \phi)$ reduces to the verification problem for the PDS \mathcal{P}_Z and the FA $\mathcal{F}^{\eta_Z(\phi)}$.

It is well-known [21, Alg. 3] that the verification problem for a PDS and an FA is in PTime and PSpace respectively. Moreover, it is possible to provide a witness trace that satisfies the said property in doubly-exponential time [47, Section 3.1.4]. Since \mathcal{P}_Z is exponential in the size of the input ZFC (due to the exponential blow-up of the local state space of the resolver Res_r),⁸ we have the following result:

LEMMA 5.2. The set of reachable configurations in \mathcal{P}_Z can be computed in exponential time in the size of the input $\langle Z, \Theta \rangle$.

We define the *product* of $\mathcal{F}^{\eta_Z(\phi)}$ with \mathcal{P}_Z as the standard product PDS $\mathcal{F}^{\eta_Z(\phi)} \times \mathcal{P}_Z$, such that $(p_1, p_2) \xrightarrow{\langle \alpha, a, \alpha' \rangle} (q_1, q_2)$ is a transition in $\mathcal{F}^{\eta_Z(\phi)} \times \mathcal{P}_Z$ if and only if $p_1 \xrightarrow{a} q_1$ is a transition in $\mathcal{F}^{\eta_Z(\phi)}$ and $p_2 \xrightarrow{\langle \alpha, a, \alpha' \rangle} q_2$ is a transition in \mathcal{P}_Z . The set of initial states is the product of the set of initial states in $\mathcal{F}^{\eta_Z(\phi)}$ and \mathcal{P}_Z . We then call a state $\langle p_1, p_2 \rangle$ in $\mathcal{F}^{\eta_Z(\phi)} \times \mathcal{P}_Z$ *bad* if p_1 is a final state in $\mathcal{F}^{\eta_Z(\phi)}$. It is straightforward to see that the verification problem for \mathcal{P}_Z and $\mathcal{F}^{\eta_Z(\phi)}$ is equivalent to checking the membership of an s -configuration, for a bad state s , in the set of reachable configurations of $\mathcal{F}^{\eta_Z(\phi)} \times \mathcal{P}_Z$. Thus, we have the complexity result,

LEMMA 5.3. $\text{DNSVERIF}(Z, \Theta, \phi)$ is 2ExpTime. The complexity reduces to ExpTime if the output of the problem does not require a witness trace.

Figure 11 provides an overview of our decision procedure. Since our abstraction is effectively defined as a monoid homomorphism (η_Z), one can deduce a concrete trace by mapping back every configuration in the abstract output (trace) to a canonical concrete configuration. In particular, it suffices to replace every occurrence of a congruence class, under η_Z , in a configuration by its canonical element. Since every congruence class can be effectively represented by a regular language, the canonical element can be chosen as one of the shortest domain names in the language.⁹ Most importantly, the soundness of this procedure follows from the completeness of our abstraction.

⁸Note that this does not change even after the abstraction as the blow-up appears when we enumerate the caches and server lists in the resolver states.

⁹This choice can be made deterministic by defining an appropriate lexicographic ordering.




Fig. 11. Overview of our decision procedure.

6 Applications

In this section, we show some applications of our results in the context of a DoS analysis of DNS. We present the examples informally, showing how instances of security properties can be characterized as FO($<$)-properties, and translating the security analysis to an instance of the `DNSVERIF` problem. We further demonstrate how we can compute the abstractions of the inputs of the `DNSVERIF` instance using syntactic domain congruence.

6.1 Amplification Attacks

An *amplification attack* is a type of DoS attack where an attacker exploits DNS to overwhelm a target, such as resolvers and nameservers, with a large volume of traffic, effectively disrupting its ability to function. For example, in the recently discovered NXNSAttack (NoneXistent Name Server Attack) [6], a small DNS query triggers the DNS resolver to issue a much larger response, potentially achieving a packet amplification factor (PAF) of over a thousand times the original query size.

Let us first have a closer look at amplification attacks. Informally, given a ZFC $\langle \mathcal{Z}, \Theta \rangle$, a target process p , and an amplification factor $paf \in \mathbb{N}$, the corresponding instance of DNS exhibits an amplification attack if there exists some query that, during the query resolution process, causes the target process p (either a nameserver or a resolver) in the network to receive more than paf messages. Typically, the target process refers to a nameserver in the network which, for a sufficiently large paf , is then unable to respond to legitimate query requests. In order to discover amplification attacks, one must exhaustively analyze all possible DNS behaviours, *for every query*, and check if there exists a behaviour where a target nameserver receives more than paf messages.

Borrowing notation from earlier sections, we can formally characterize the set of traces that exhibit an amplification attack as the formula,¹⁰

$$\phi_{amp}(paf, p) = \exists x_1, \dots, x_{paf} \in \mathbb{N}. \left(\bigwedge_{i \neq j \in \{1, \dots, paf\}} x_i \neq x_j \right) \wedge \left(\bigwedge_{i \in \{1, \dots, paf\}} \bigvee_{j \in P_{\mathcal{Z}}} (\exists msg \in M_{\mathcal{Z}}. (\text{ch}_{(j,p)}?msg)(x_i)) \right)$$

The above formula captures the property that there exist paf distinct positions in a trace, such that every one of these positions corresponds to a receive action for the target process p . Indeed, any run in $\text{DNS}_{\mathcal{Z}}$ whose trace satisfies the above property exhibits a receive action for the target component \mathcal{P}^p at least paf times. Since every run in $\text{DNS}_{\mathcal{Z}}$ begins with a query request, a run that satisfies $\phi_{amp}(paf, p)$ corresponds to an amplification attack.

Example 6.1. Figure 12a illustrates how the NXNSAttack is mounted using a simple example. This example involves two attacker components, namely a client and an authoritative nameserver at `attacker.net`, and two innocent components, namely the recursive resolver and an authoritative

¹⁰We use many-sorted quantification in the formula, in order to quantify over the space of all messages in $M_{\mathcal{Z}}$, which is an infinite set. The alternative would be an infinite disjunction over all possibilities of msg .

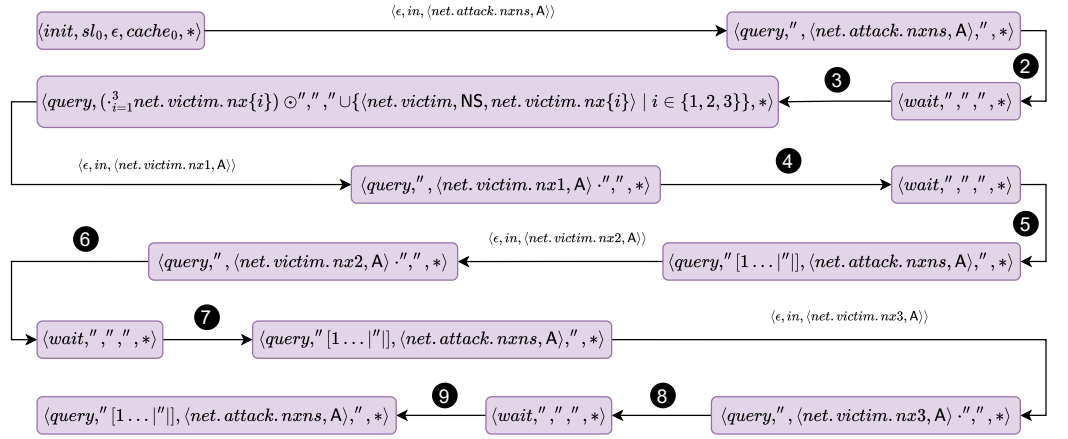
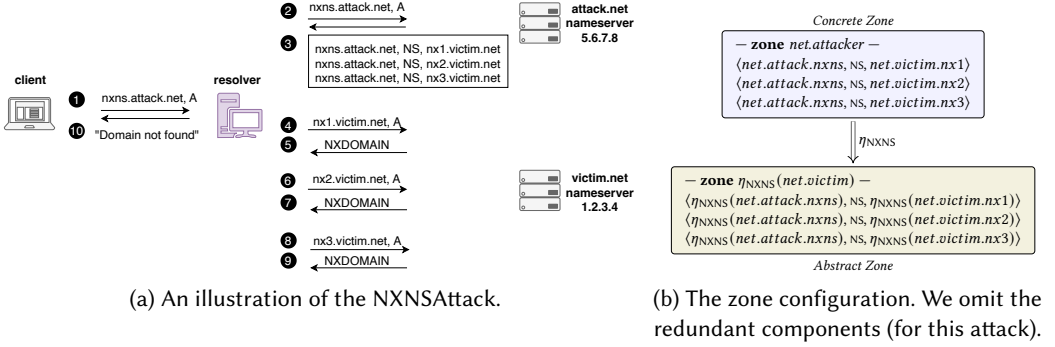


Fig. 12. NXNSAttack.

nameserver `victim.net`. The client issues a request for the subdomain `nxns.attack.net` that is authorized by the attacker's authoritative nameserver (1). The resolver then attempts to resolve the subdomain (2). In response, the attacker's nameserver `attack.net` sends an NS referral response with multiple nameserver names (three in this case) but without their IP addresses (3). This forces the resolver to issue additional subqueries to resolve each of these names (4–9), thus amplifying the DNS traffic.

The ZFC can be represented by $\langle \text{NXNS}, \{5.6.7.8, 1.2.3.4\} \rangle$, where the domain of NXNS is the set $\{5.6.7.8, 1.2.3.4\}$, such that 5.6.7.8 is the IP of the attacker's nameserver and 1.2.3.4 that of the victim. The function NXNS essentially¹¹ maps 1.2.3.4 to the empty set and 5.6.7.8 to the concrete zone

¹¹For simplicity, we only consider the records that are relevant for the attack.

illustrated in Figure 12b. The malicious zone contains three NS records that point to the nameservers with domain names $net.victim.nx\{i\}$, for $i \in \{1, 2, 3\}$. But, observe that none of them have an accompanying record of type A that points those names to an IP address. The IP addresses 1.2.3.4 and 5.6.7.8 are recorded as the root nameserver IP addresses.

We now describe the run of the resolver in DNS_{NXNS} . The resolver starts from the initial configuration, and takes the internal transition with the label $\langle \epsilon, in, \langle net.attacker.nxns, A \rangle \rangle$ to model that it received the query $\langle net.attacker.nxns, A \rangle$ from the client. When it receives the query, it moves to the *query* state ready to start resolution. The current stack configuration is $\langle net.attack.nxns, A \rangle$, which denotes that the initial query is currently under the resolution. It then sends the query out to the leftmost nameserver in the server list that contains the domain name of the query in its zone, namely *attacker.net*, and then waits for an answer by moving to the state *wait*. Once it receives the response with a set of NS records from $NS_{5.6.7.8}$, it adds them to its cache, along with the corresponding domain name of the nameserver in the value of each record (in the answer) to its server list. It moves back to the *query* state to decide the next step for the resolution. Since every NS record in the answer was glueless, the resolver proceeds with the resolution of $\langle net.victim.nx1, A \rangle$ in order to find the glue. Consequently, it pushes the query on its stack while staying in the *query* state. Since the query name is in the zone *victim.net*, it sends the query to $NS_{1.2.3.4}$. However, since the zone file for *victim.net* is empty, the resolver gets $\langle NXDOMAIN, \emptyset \rangle$ (6) as the answer, thus deleting $net.victim.nx1$ from its server list and moving back to the *query* state. This repeats two more times for the other two entries in the server list. It is then straightforward to see that this behaviour leads to $NS_{1.2.3.4}$ receiving three queries. In fact, this behavioural pattern can be easily extended such that $NS_{1.2.3.4}$ receives n queries for any $n \in \mathbb{N}$, by adding more records of the form $\langle net.attack.nxns, NS, net.victim.nx\{i\} \rangle$ to the zone *attacker.net*. Figure 12c illustrates the formal run in the process Res_r .

The PPT languages required to compute the abstraction are $L_0 = net \cdot attack \cdot nxns \cdot \Sigma_{DNS}^+$, and $L_i = net \cdot victim \cdot nx\{i\} \cdot \Sigma_{DNS}^+$, where $\Sigma_{DNS} = \{net, attack, nxns, 1, 2, 3, other\}$ and $i \in \{1, 2, 3\}$. We compute the monoid $SYN_{NXNS} = \prod_{0 \leq j \leq 3} \langle SYN(L_j) \rangle$, denoting the syntactic domain congruence classes and the corresponding homomorphism η_{NXNS} . Fixing the packet amplification factor to be 3 and the target process to be 1.2.3.4, the formula $\phi_{amp}(3, 1.2.3.4)$ can then be rewritten as

$$\phi'_{amp}(3, 1.2.3.4) = \exists x_1, \dots, x_3 \in \mathbb{N}. \left(\bigwedge_{i \neq j \in \{1, \dots, 3\}} x_i \neq x_j \right) \wedge \left(\bigwedge_{i \in \{1, \dots, 3\}} \bigvee_{j \in \{r, 5.6.7.8\}} \bigvee_{\substack{m \in SYN_{NXNS} \\ t \in TYPE}} (ch_{(j, 1.2.3.4)}? \langle m, t \rangle)(x_i) \right),$$

which no longer contains any quantifications other than the positions in the trace. The zone can be translated as shown in Figure 12b.

Unbounded Analysis. Note that, if there are multiple nameservers, unlike the trivial example given above, there is no limit to the length of runs which can ensure the presence or absence of an amplification attack. Indeed, a resolver may communicate with the same benign nameserver owning multiple zones before establishing communication with the target nameserver. Hence, a complete search for amplification attacks necessitates the unbounded reachability analysis of DNS.

6.2 Rewrite Blackholing

Rewrite blackholing occurs when a query is eventually redirected to a non-existent domain name. This can happen when the resolver, for example, follows some CNAME records, but ultimately finds that the final “canonical” name does not exist. While rewrite blackholing can be useful for blocking malicious traffic or preventing attacks, it can also be exploited to blackhole legitimate services or domains, thus causing a DoS for valid queries.

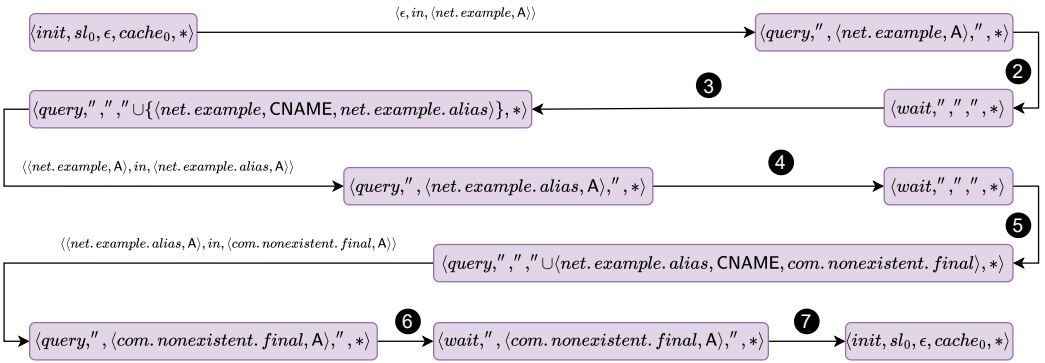
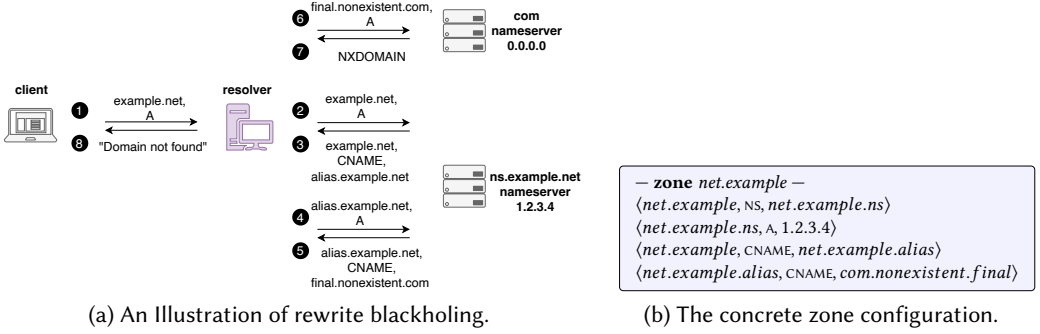


Fig. 13. Rewrite blackholing.

We now show that this attack vector can also be specified in our framework. Informally, given a ZFC $\langle \mathcal{Z}, \Theta \rangle$, the corresponding instance of DNS exhibits rewrite blackholing if there exists some query that, during its resolution process, causes the resolver to eventually receive an answer with the label `NXDOMAIN` along with a non-empty set of records.

We can again formally characterize the set of traces exhibiting rewrite blackholing as the formula

$$\phi_{rb} = \exists x \in \mathbb{N}. \exists recs \in \mathcal{P}(\text{RECORD}_L). \left(recs \neq \emptyset \wedge \bigvee_{j \in P_Z} ((ch_{(j,p)}?(\text{NXDOMAIN}, recs))(x_j)) \right).$$

Example 6.2. As illustrated in Figure 13a, a DNS query for `example.net` returns a `CNAME` pointing to `alias.example.net` (3). The resolver then queries `alias.example.net`, which returns another `CNAME` pointing to `final.nonexistent.com` (5). Since `final.nonexistent.com` does not exist, the resolver eventually receives an `NXDOMAIN` response, indicating that the domain cannot be resolved (7). In this rewrite blackholing scenario, two authoritative nameservers are involved: one for handling the resolution of `example.net` and `alias.example.net` (as both are within the same domain zone `example.net`), and the other for `nonexistent.com`. The concrete zone file for the `example.net` domain is shown in Figure 13b and the corresponding formal run in Res_r is illustrated in Figure 13c.

7 Related Work

Our framework, along with its underlying formalism, analysis approach, and application domain, is closely related to the following three lines of earlier research.

Formal Modeling and Analysis of DNS. Kakarla et al. [30] formalize the DNS resolution semantics, abstracting away critical resolver-side logic including caching and recursive queries. Building upon this semantics, they also provide the first static analyzer GRoot for identifying DNS misconfiguration errors. Liu et al. [36] complement GRoot’s semantics by providing the first formal semantics of end-to-end name resolution. The accompanying tool DNSMaude offers a collection of formal analyses for both qualitative (e.g., rewrite blackholing) and quantitative properties (e.g., amplification). We have already discussed their limitations in Section 1.

Testing, Monitoring, and Verification of DNS Implementations. Recent years have seen significant advances in DNS testing tools [14, 29, 32, 59]. ResolFuzz [14] is a differential tester, searching for semantic bugs such as inaccuracies in resolver responses across DNS resolver implementations. ResolverFuzz [59] fuzzes resolvers via short query-response sequences. It also leverages differential testing and can identify resolver vulnerabilities like cache poisoning bugs. Unlike pure fuzzing, SCALE [32] guides its test generation by symbolically executing GRoot’s semantics. It focuses on detecting RFC compliance errors in authoritative nameserver implementations rather than recursive resolvers. While also deriving exhaustive test cases using symbolic execution, Eywa [29] leverages large language models to construct intended DNS behaviors from natural language sources like RFCs. All these testers, along with excellent DNS monitors [27, 51], can only show the presence of bugs or vulnerabilities, not their absence, in DNS implementations. Moreover, in contrast to our framework, they do not encompass the full name resolution process involving both recursive resolvers and nameservers.

IRONSIDES [15] is a nameserver implementation provably secure against DNS vulnerabilities like single-packet DoS attacks. However, it may still be susceptible to application-layer issues such as amplification attacks, in contrast to our framework. DNS-V [60] is a verification framework for an in-house DNS authoritative engine in Alibaba Cloud. A symbolic-execution-based approach is used to facilitate the adoption of layered verification, where source code within each decomposed layer of DNS can be verified independently with respect to correctness properties. Quantitative properties like amplification are out of its scope. Wang et al. [55] scale up the verification of DNS configuration by parallelizing the analysis of zone-file-specific query processing behaviour and symbolically combining the analysis results. Unlike our framework, the use of formal methods in none of these works is related to the DNS semantics.

All these testing, monitoring, and verification efforts focus on DNS implementations. Given that many DNS vulnerabilities have stemmed from RFCs or protocol-level defects in heavily tested production resolvers and nameservers, protocol-level verification at an early design stage is highly desirable. Our framework takes a step forward towards this goal.

Verification Problem for Infinite-State Systems. Verification of infinite-state systems is a well-studied problem and is still an active area of research. Abdulla & Jonsson [4, 5] and Cécé et al. [16] independently show that the verification problem for CPSes communicating over lossy channels is decidable in non-primitive recursive time, for arbitrary topologies. Abdulla et al. [2] later show that the decidability holds even when the communicating processes are timed and the messages contain time-stamps. Heussner et al. [25, 26] study this problem for rCPSes, however, over perfect channels. They restrict their attention to eager runs and show that the problem is decidable in exponential-time for a restricted class of topologies that they call *non-confluent*. Prior to this work, Atig et al. [11] show that the problem is decidable even without the restriction of eagerness, but

instead requires the underlying topology to be acyclic. In the context of PDSes, Bouajjani et al. [12] show that the problem is decidable for PDS, using a *saturation method* that computes the set of all reachable computations. Abdulla et al. [3] and Clemente and Lasota [18] extend the decidability to timed PDS with a timed stack. In contrast, our work handles a timed extension of rCPS with an infinite message alphabet. In addition, we focus on models pertaining to the DNS semantics.

Kakarla et al. [31] study the expressive complexity of DNS. They show that the verification of DNS zone files, a simpler problem than `DNSVERIF`, is likely to take at least cubic time in the number of records, effectively giving a lower bound. In this work, we provide an upper bound for this problem (Lemma 5.3).

8 Limitations and Conclusion

A formal framework with a decision procedure for the `DNSVERIF` problem is essential to achieving a secure, reliable DNS infrastructure. Aside from being of theoretical interest in the context of infinite-state systems, this work provides a foundation for the advancement of automated verification tools for *and* eliminates the costly (Break and Fix)* cycle with DNS security analysis.

We have demonstrated our framework for two applications. However, in principle, we can model all the static vulnerabilities considered by GRoot [30, Table 1], including rewrite blackholing (Section 6.2), rewrite loop, and answer inconsistency. This is immediate as every property defined in GRoot is an LTL-property, which is equivalent to $FO(<)$ over traces. Additionally, our framework is capable of modeling all the attacks considered by DNSMaude [36, Table 4], including those covered by GRoot and amplification attacks (Section 6.1). However, our current model also presents several limitations.

Multiple Resolvers. We are currently limited to having a single resolver in our model. A promising direction for future work would be to consider multiple concurrent resolvers in the system. While our abstraction on the query space would still be sound and complete, the underlying topology would then have two pushdown systems (the resolvers) communicating with each other over a lossy bag channel, which violates the confluence criterion from [25, 26].

Dynamic Updates. We have adopted the semantics considered by DNSMaude, which subsumes that of GRoot. In particular, our framework currently does not support the query type `IXFR`, meant for incremental zone transfers. This type is responsible for updating the DNS records stored by a nameserver. Under some conditions, such queries force two-way reliable (TCP) communication between nameservers (which are finite-state systems themselves). The resulting network topology now allows a non-lossy cycle and thus yields the undecidability of the corresponding state-reachability problem, which in turn implies the undecidability of the `DNSVERIF` problem. More details on this kind of behaviour can be found in [44].

Further model-specific investigation is needed to remove the above limitations, which is outside the scope of this paper. Finally, the $2ExpTime$ complexity upper bound of our procedure is rather pessimistic. Hence, empirical results are necessary to prove the practicality of our approach. At this point, our contribution is primarily theoretical, and significant engineering work will be required to design and implement these techniques. Conducting an experimental evaluation of our approach remains a challenging future task.

Acknowledgments

We appreciate the reviewers for their valuable feedback. This research is supported by armasuisse Science and Technology and an ETH Zurich Career Seed Award. We would also like to thank Narayan Kumar and Prakash Saivasan for some useful initial discussions.

References

- [1] 1957. *Fundamental Concepts in the Theory of Systems*. Wright Air Development Center, Air Research and Development Command, United States Air Force.
- [2] Parosh Aziz Abdulla, Mohamed Faouzi Atig, and Jonathan Cederberg. 2012. Timed lossy channel systems. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2012)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [3] Parosh Aziz Abdulla, Mohamed Faouzi Atig, and Jari Stenman. 2012. Dense-timed pushdown automata. In *2012 27th Annual IEEE Symposium on Logic in Computer Science*. IEEE, 35–44.
- [4] Parosh Aziz Abdulla and Bengt Jonsson. 1993. Verifying Programs with Unreliable Channels. In *Proc. LICS '93-8th IEEE Int. Symp. on Logic in Computer Science*. 160–170.
- [5] Parosh Aziz Abdulla and Bengt Jonsson. 1996. Verifying Programs with Unreliable Channels. *Information and Computation* 127, 2 (1996), 91–101.
- [6] Yehuda Afek, Anat Bremner-Barr, and Lior Shafir. 2020. NXNSAttack: Recursive DNS Inefficiencies and Vulnerabilities. In *USENIX Security 2020*. USENIX Association, 631–648.
- [7] Yehuda Afek, Anat Bremner-Barr, and Shani Stajrod. 2023. NRDelegationAttack: Complexity DDoS attack on DNS Recursive Resolvers. In *USENIX Security '23*. USENIX Association, 3187–3204.
- [8] C Aiswarya. 2020. On network topologies and the decidability of reachability problem. In *International Conference on Networked Systems*. Springer, 3–10.
- [9] Jorge Almeida. 1995. *Finite semigroups and universal algebra*. Vol. 3. World Scientific.
- [10] Rajeev Alur and David L Dill. 1994. A theory of timed automata. *Theoretical computer science* 126, 2 (1994), 183–235.
- [11] Mohamed Faouzi Atig, Ahmed Bouajjani, and Tayssir Touili. 2008. On the reachability analysis of acyclic networks of pushdown systems. In *International Conference on Concurrency Theory*. Springer, 356–371.
- [12] Ahmed Bouajjani, Javier Esparza, and Oded Maler. 1997. Reachability analysis of pushdown automata: Application to model-checking. In *CONCUR '97*. Springer, 135–150.
- [13] Daniel Brand and Pitro Zafiropolo. 1983. On communicating finite-state machines. *Journal of the ACM (JACM)* 30, 2 (1983), 323–342.
- [14] Jonas Bushart and Christian Rossow. 2024. ResolFuzz: Differential Fuzzing of DNS Resolvers. In *Computer Security – ESORICS 2023*. Springer Nature Switzerland, 62–80.
- [15] Martin C. Carlisle and Barry S. Fagin. 2012. IRONSIDES: DNS with no single-packet denial of service or remote code execution vulnerabilities. In *GLOBECOM 2012*. IEEE, 839–844.
- [16] Gérard Cécé, Alain Finkel, and S Purushothaman Iyer. 1996. Unreliable channels are easier to verify than perfect channels. *Information and Computation* 124, 1 (1996), 20–31.
- [17] Manuel Clavel, Francisco Durán, Steven Eker, Patrick Lincoln, Narciso Martí-Oliet, José Meseguer, and Carolyn L. Talcott (Eds.). 2007. *All About Maude - A High-Performance Logical Framework, How to Specify, Program and Verify Systems in Rewriting Logic*. LNCS, Vol. 4350. Springer.
- [18] Lorenzo Clemente and Slawomir Lasota. 2015. Timed pushdown automata revisited. In *2015 30th Annual ACM/IEEE Symposium on Logic in Computer Science*. IEEE, 738–749.
- [19] Huayi Duan, Marco Bearzi, Jodok Vieli, David Basin, Adrian Perrig, Si Liu, and Bernhard Tellenbach. 2024. CAMP: Compositional Amplification Attacks against DNS. In *USENIX Security '24*. USENIX Association, 5769–5786.
- [20] Robert Elz and Randy Bush. 1997. Clarifications to the DNS Specification. *RFC* 2181 (1997), 1–14. <https://doi.org/10.17487/RFC2181>
- [21] Javier Esparza, David Hansel, Peter Rossmanith, and Stefan Schwoon. 2000. Efficient algorithms for model checking pushdown systems. In *International Conference on Computer Aided Verification*. Springer, 232–247.
- [22] Alain Finkel and Ph Schnoebelen. 2001. Well-structured transition systems everywhere! *Theoretical Computer Science* 256, 1-2 (2001), 63–92.
- [23] Victor Mikhaylovich Glushkov. 1961. The abstract theory of automata. *Russian Mathematical Surveys* 16, 5 (1961), 1.
- [24] Sergey Goncharov and Dirk Pattinson. 2014. Coalgebraic weak bisimulation from recursive equations over monads. In *International Colloquium on Automata, Languages, and Programming*. Springer, 196–207.
- [25] Alexander Heußner, Jérôme Leroux, Anca Muscholl, and Grégoire Sutre. 2010. Reachability analysis of communicating pushdown systems. In *FoSSaCS '10*. Springer, 267–281.
- [26] Alexander Heußner, Jérôme Leroux, Anca Muscholl, and Grégoire Sutre. 2012. Reachability analysis of communicating pushdown systems. *Logical Methods in Computer Science* 8 (2012).
- [27] Check Host. Accessed in July, 2024. DNS checking. <https://check-host.net/check-dns>.
- [28] IETF. Accessed 2024-10-15. Internet Engineering Task Force. <https://www.ietf.org/>.
- [29] Siva Kesava Reddy Kakarla and Ryan Beckett. 2023. Oracle-based Protocol Testing with Eywa. arXiv:2312.06875 [cs.NI]
- [30] Siva Kesava Reddy Kakarla, Ryan Beckett, Behnaz Arzani, Todd Millstein, and George Varghese. 2020. GRoot: Proactive Verification of DNS Configurations. In *SIGCOMM '20*. ACM, 310–328.

- [31] Siva Kesava Reddy Kakarla, Ryan Beckett, Todd Millstein, and George Varghese. 2021. How Complex is DNS?. In *HotNets '21*. ACM, 116–122.
- [32] Siva Kesava Reddy Kakarla, Ryan Beckett, Todd Millstein, and George Varghese. 2022. SCALE: Automatically Finding RFC Compliance Bugs in DNS Nameservers. In *NSDI '22*. USENIX Association, 307–323.
- [33] Dexter C Kozen. 2012. *Automata and computability*. Springer Science & Business Media.
- [34] James F. Kurose and Keith W. Ross. 2016. *Computer Networking: A Top-Down Approach* (7 ed.). Pearson, Boston, MA.
- [35] X. Li, W. Xu, B. Liu, M. Zhang, Z. Li, J. Zhang, D. Chang, X. Zheng, C. Wang, J. Chen, H. Duan, and Q. Li. 2024. TUDOOR Attack: Systematically Exploring and Exploiting Logic Vulnerabilities in DNS Response Pre-processing with Malformed Packets. In *2024 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 185–185.
- [36] Si Liu, Huayi Duan, Lukas Heimes, Marco Bearzi, Jodok Vieli, David A. Basin, and Adrian Perrig. 2023. A Formal Framework for End-to-End DNS Resolution. In *SIGCOMM '23*. ACM, 932–949.
- [37] Florian Maury. 2015. The “Indefinitely” Delegating Name Servers (iDNS) Attack. <https://indico.dns-oarc.net/event/21/contributions/301/attachments/272/492/slides.pdf>.
- [38] Robin Milner. 1980. *A calculus of communicating systems*. Springer.
- [39] P. Mockapetris. 1987. *DOMAIN NAMES - CONCEPTS AND FACILITIES*. RFC 1034. RFC Editor. <https://doi.org/10.17487/RFC1034>
- [40] P. Mockapetris. 1987. *DOMAIN NAMES - IMPLEMENTATION AND SPECIFICATION*. RFC 1035. RFC Editor. <https://doi.org/10.17487/RFC1035>
- [41] Giovane C. M. Moura, Sebastian Castro, John S. Heidemann, and Wes Hardaker. 2021. TsuNAME: exploiting misconfiguration and vulnerability to DDoS DNS. In *IMC '21*. ACM, 398–418.
- [42] A. Nerode. 1958. Linear Automaton Transformations. *Proc. Amer. Math. Soc.* 9, 4 (1958), 541–544.
- [43] Dhruv Nevatia, Si Liu, and David Basin. 2024. *Reachability Analysis of the Domain Name System*. Technical Report. <https://doi.org/10.48550/arXiv.2411.10188>
- [44] Dr. Masataka Ohta. 1996. Incremental Zone Transfer in DNS. RFC 1995. <https://doi.org/10.17487/RFC1995>
- [45] Jean-Éric Pin. 1995. A variety theorem without complementation. *Russian Mathematics (Izvestija vuzov. Matematika)* 39 (1995), 80–90.
- [46] M. O. Rabin and D. Scott. 1959. Finite Automata and Their Decision Problems. *IBM Journal of Research and Development* 3, 2 (1959), 114–125.
- [47] Thomas Reps, Stefan Schwoon, Somesh Jha, and David Melski. 2005. Weighted pushdown systems and their application to interprocedural dataflow analysis. *Science of Computer Programming* 58, 1-2 (2005), 206–263.
- [48] S. Rose and W. Wijngaards. 2012. *DNAME Redirection in the DNS*. RFC 6672. RFC Editor. <https://doi.org/10.17487/RFC6672>
- [49] Koushik Sen, Mahesh Viswanathan, and Gul Agha. 2005. On Statistical Model Checking of Stochastic Systems. In *CAV 2005 (LNCS, Vol. 3576)*. Springer.
- [50] Sam Staton. 2011. Relating coalgebraic notions of bisimulation. *Logical Methods in Computer Science* 7 (2011).
- [51] ThousandEyes. Accessed in July, 2024. DNS Monitoring Across Your Entire Network. <https://www.thousandeyes.com/solutions/dns-monitoring>.
- [52] Liam Tung. Accessed in July 2024. Azure global outage: Our DNS update mangled domain records, says Microsoft. <https://www.zdnet.com/article/azureglobal-outage-our-dns-update-mangled-domain-records-says-microsoft/>.
- [53] Frits Vaandrager, Masoud Ebrahimi, and Roderick Bloem. 2023. Learning Mealy machines with one timer. *Information and Computation* 295 (2023), 105013.
- [54] Lorijn van Rooijen and Marc Zeitoun. 2013. The separation problem for regular languages by piecewise testable languages. *arXiv preprint arXiv:1303.2143* (2013).
- [55] Yao Wang, Kexin Yu, Ziyi Wang, Kaiqiang Hu, Haizhou Du, Qiao Xiang, Xing Fang, Geng Li, Ruiting Zhou, Linghe Kong, and Jiwu Shu. 2024. Rethinking DNS Configuration Verification with a Distributed Architecture. In *APNet '24*. ACM, 23–30.
- [56] Zack Whittaker. Accessed 2023-02-07. A DNS outage just took down a large chunk of the internet. <https://techcrunch.com/2021/07/22/a-dns-outage-just-took-down-a-good-chunk-of-the-internet/>.
- [57] Wikipedia. Accessed in July 2024. 2021 Facebook outage. https://en.wikipedia.org/wiki/2021_Facebook_outage.
- [58] Wei Xu, Xiang Li, Chaoyi Lu, Baojun Liu, Haixin Duan, Jia Zhang, Jianjun Chen, and Tao Wan. 2023. TsuKing: Coordinating DNS Resolvers and Queries into Potent DoS Amplifiers. In *CCS '23*. ACM, 311–325.
- [59] Qifan Zhang, Xuesong Bai, Xiang Li, Haixin Duan, Qi Li, and Zhou Li. 2024. ResolverFuzz: Automated Discovery of DNS Resolver Vulnerabilities with Query-Response Fuzzing. In *USENIX Security '24*. USENIX Association.
- [60] Naiqian Zheng, Mengqi Liu, Yuxing Xiang, Linjian Song, Dong Li, Feng Han, Nan Wang, Yong Ma, Zhuo Liang, Dennis Cai, Ennan Zhai, Xuanzhe Liu, and Xin Jin. 2023. Automated Verification of an In-Production DNS Authoritative Engine. In *SOSP '23*. ACM, 80–95.

Received 2024-07-11; accepted 2024-11-07