




Demystifying Resolver Work: A Formal Perspective on DNS Resolver Performance Vulnerabilities

Liwen Xu, Doctoral Researcher, ETH Zürich (liwen.xu@inf.ethz.ch)

Joint work with: Huayi Duan (HKUST-GZ), Zechao Cai, and Adrian Perrig (ETH Zürich)

Classes of DoS Attacks Targeting Resolvers


Traditional Volumetric DDoS



Many queries
Large traffic volume

Classes of DoS Attacks Targeting Resolvers

Traditional Volumetric DDoS



Many queries
Large traffic volume


Resolver Work Amplification



A few queries
Disproportionate work

Classes of DoS Attacks Targeting Resolvers

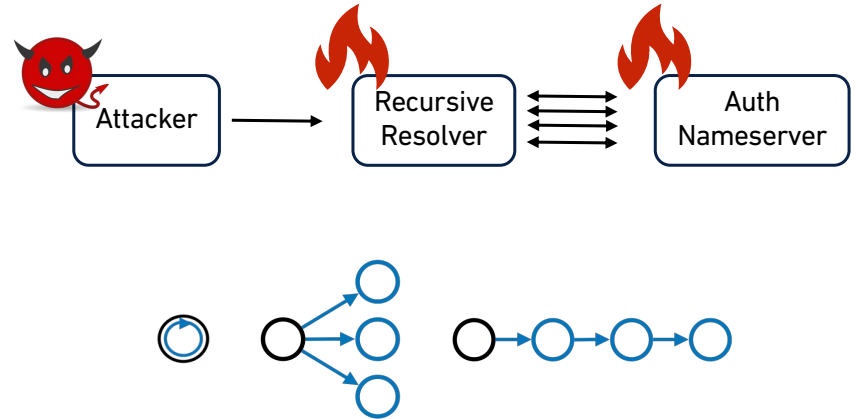
Traditional Volumetric DDoS



Many queries
Large traffic volume

Resolver Work Amplification

(aka, Self-Amplification @ DNS-OARC'43)



A few queries
Disproportionate work

The Problem That Never Got Answered Explicitly To The Full

The recommended priorities for the resolver designer are:

1. **Bound the amount of work** (packets sent, parallel processes started) so that a request can't get into an infinite loop or start off a chain reaction of requests or queries with other implementations **EVEN IF SOMEONE HAS INCORRECTLY CONFIGURED SOME DATA.**

– RFC1034

The Problem That Never Got Answered Explicitly To The Full

The recommended priorities for the resolver designer are:

1. **Bound the amount of work** (packets sent, parallel processes started) so that a request can't get into an infinite loop or start off a chain reaction of requests or queries with other implementations **EVEN IF SOMEONE HAS INCORRECTLY CONFIGURED SOME DATA.**

– RFC1034

What RFC 1034 leaves undefined

- Precise definition of "work"
- Way to measure it
- Guidance on the bound

The Problem That Never Got Answered Explicitly To The Full

The recommended priorities for the resolver designer are:

1. **Bound the amount of work** (packets sent, parallel processes started) so that a request can't get into an infinite loop or start off a chain reaction of requests or queries with other implementations **EVEN IF SOMEONE HAS INCORRECTLY CONFIGURED SOME DATA.**


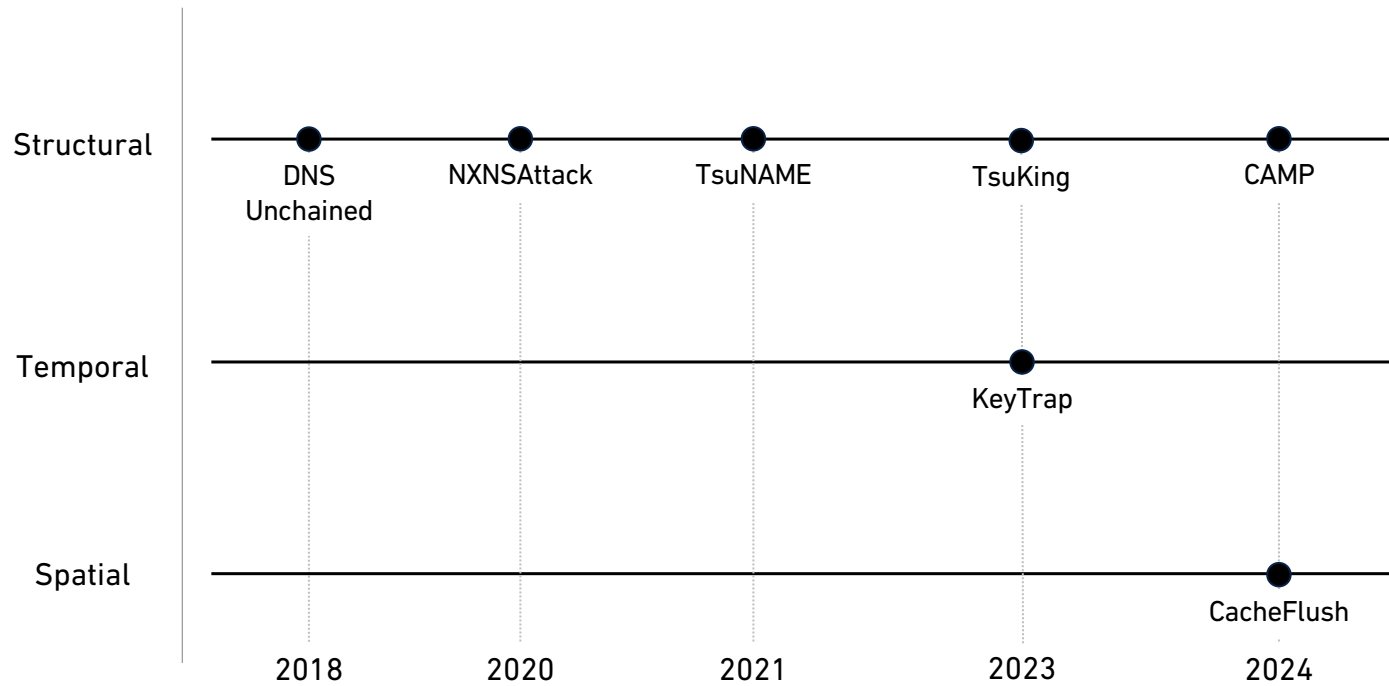
– RFC1034

What RFC 1034 leaves undefined

- Precise definition of "work"
- Way to measure it
- Guidance on the bound

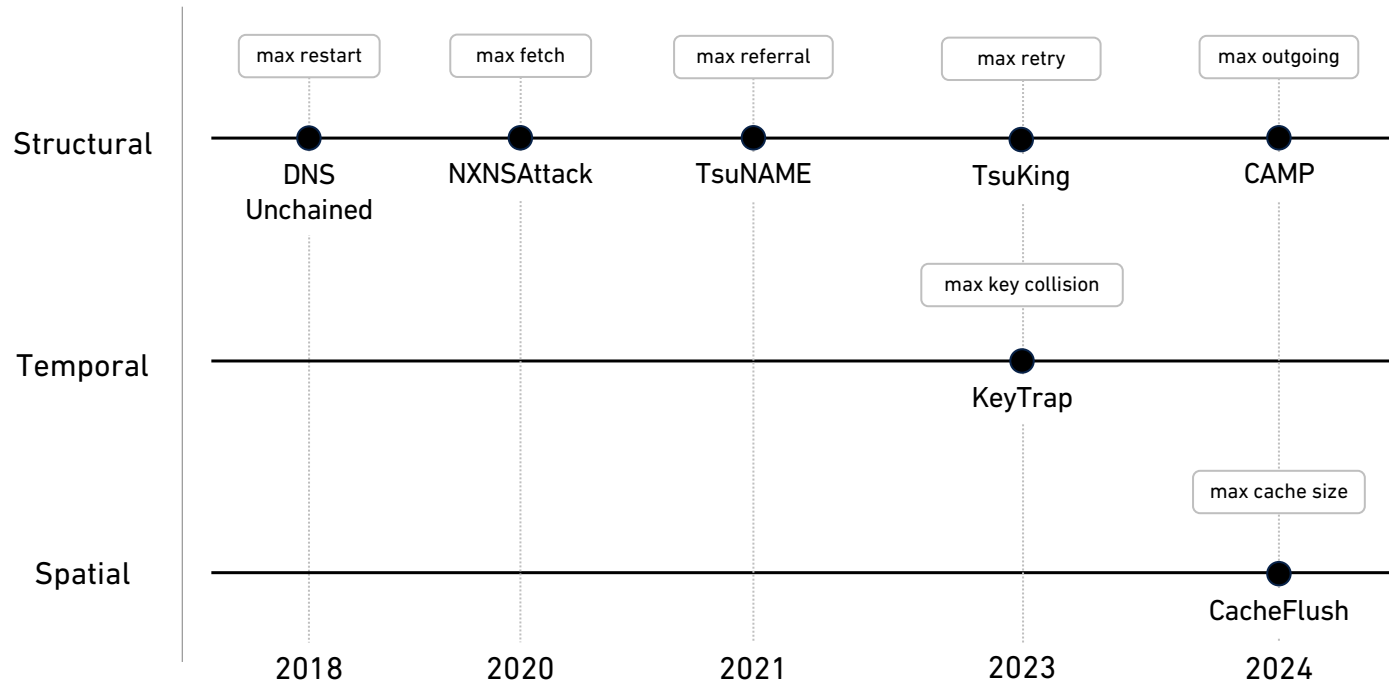
What is the maximum amount
of work a resolver can
perform for a single query?

We're Caught In A Break-and-Fix Cycle




* Taxonomy based on the primary source of introduced complexity

We're Caught In A Break-and-Fix Cycle





* Taxonomy based on the primary source of introduced complexity



New attack vector

↓
New patch(es)

We're Caught In A Break-and-Fix Cycle

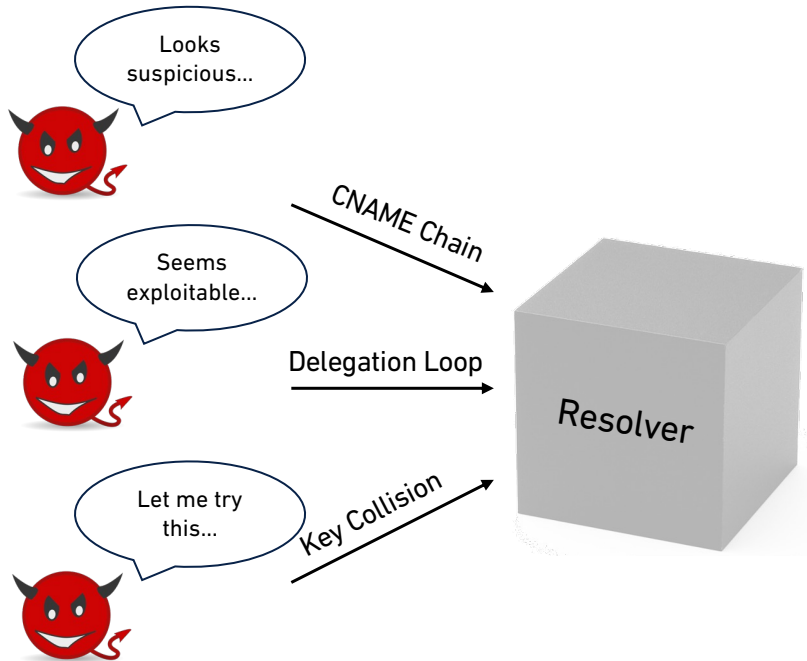


New attack vector
↓
New patch(es)

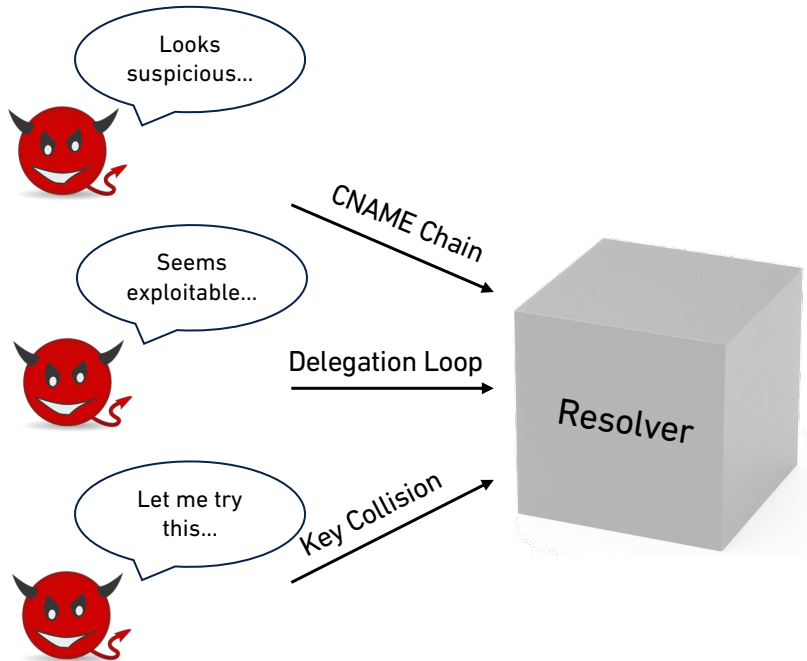
* Taxonomy based on the primary source of introduced complexity

Are there more? How many more?

From Intuition to Formalization

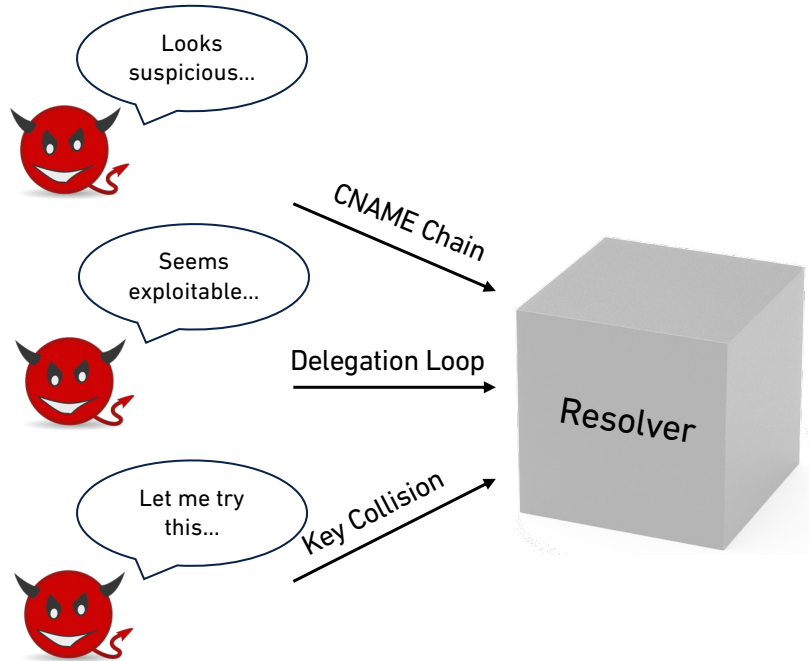


From Intuition to Formalization



Let's go back to where it all started.

From Intuition to Formalization



Let's go back to where it all started.

The top level algorithm has four steps:

1. See if the answer is in local information, and if so return it to the client.
2. Find the best servers to ask.
3. Send them queries until one returns a response.
4. Analyze the response, either:
 - a. if the response answers the question or contains a name error, cache the data as well as returning it back to the client.
 - b. if the response contains a better delegation to other servers, cache the delegation information, and go to step 2.
 - c. if the response shows a CNAME and that is not the answer itself, cache the CNAME, change the SNAME to the canonical name in the CNAME RR and go to step 1.
 - d. if the response shows a servers failure or other bizarre contents, delete the server from the SLIST and go back to step 3.

– RFC1034, §5.3.3

Resolution as a Transition System

Algorithm 1: Recursive Resolution Algorithm

```
Input: A query  $Q = (SNAME, STYPE)$  from a client
Output: Resolved answer  $A$  or an error response  $\perp$ 
/* Step 1: Check local cache */
1 if  $\exists A = (SNAME, STYPE, RRs) \in CACHE$  then
2   return  $A$  to client ;
3   terminate ;
/* Step 2: Find the best authoritative servers */
4  $SLIST \leftarrow \text{FindBestServers}(Q)$  ;
/* Step 3: Send queries and wait for a response */
5 repeat
6   foreach  $s \in SLIST$  do
7     Send query  $Q$  to server  $s$  ;
8     Collect all responses in  $\mathcal{R}$  ;
9   until  $\exists R \in \mathcal{R}$  such that  $R \neq \emptyset$  ;
/* Step 4: Case A */
10 if  $A = (SNAME, STYPE, RRs) \in R$  or  $R = \perp$  (Name Error) then
11    $CACHE \leftarrow CACHE \cup \{A\}$  ;
12   return  $R$  to client ;
13   terminate ;
/* Step 4: Case B */
14 else if  $NS: D = (SNAME, S) \in R$  and  $D$  is valid then
15    $CACHE \leftarrow CACHE \cup \{D\}$  ;
16    $SLIST \leftarrow SLIST \cup \{S\}$  ;
17   goto step 2 ;
/* Step 4: Case C */
18 else if  $CNAME: C = (SNAME \rightarrow SNAME') \in R$  and  $STYPE \neq CNAME$  then
19    $CACHE \leftarrow CACHE \cup \{C\}$  ;
20    $Q \leftarrow (SNAME', STYPE)$  ;
21   goto step 1 ;
/* Step 4: Case D */
22 else if  $R$  indicates a server failure or an invalid response then
23    $SLIST \leftarrow SLIST \setminus \{s\}$  ;
24   goto step 3
25   or
26   return SERVFAIL to client;
27   terminate ;
```

Resolution as a Transition System

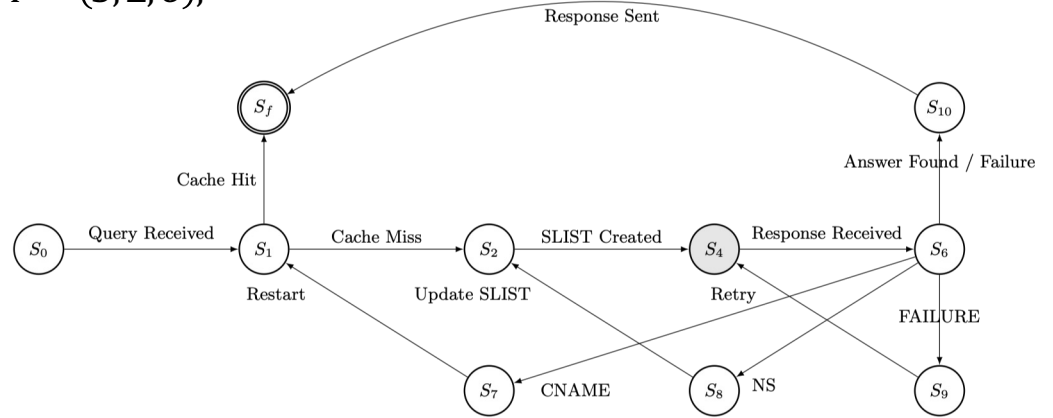
Algorithm 1: Recursive Resolution Algorithm

```

Input: A query  $Q = (SNAME, STYPE)$  from a client
Output: Resolved answer  $A$  or an error response  $\perp$ 
/* Step 1: Check local cache */
1 if  $\exists A = (SNAME, STYPE, RRs) \in CACHE$  then */
2   return  $A$  to client ;
3   terminate ;
/* Step 2: Find the best authoritative servers */
4  $SLIST \leftarrow \text{FindBestServers}(Q)$  ;
/* Step 3: Send queries and wait for a response */
5 repeat */
6   foreach  $s \in SLIST$  do
7     Send query  $Q$  to server  $s$  ;
8   Collect all responses in  $\mathcal{R}$  ;
9 until  $\exists R \in \mathcal{R}$  such that  $R \neq \emptyset$  ;
/* Step 4: Case A */
10 if  $A = (SNAME, STYPE, RRs) \in R$  or  $R = \perp$  (Name Error) then */
11    $CACHE \leftarrow CACHE \cup \{A\}$  ;
12   return  $R$  to client ;
13   terminate ;
/* Step 4: Case B */
14 else if  $NS: D = (SNAME, S) \in R$  and  $D$  is valid then */
15    $CACHE \leftarrow CACHE \cup \{D\}$  ;
16    $SLIST \leftarrow SLIST \cup \{S\}$  ;
17   goto step 2 ;
/* Step 4: Case C */
18 else if  $CNAME: C = (SNAME \rightarrow SNAME') \in R$  and  $STYPE \neq CNAME$  then */
19    $CACHE \leftarrow CACHE \cup \{C\}$  ;
20    $Q \leftarrow (SNAME', STYPE)$  ;
21   goto step 1 ;
/* Step 4: Case D */
22 else if  $R$  indicates a server failure or an invalid response then */
23    $SLIST \leftarrow SLIST \setminus \{s\}$  ;
24   goto step 3
25   or
26   return SERVFAIL to client;
27   terminate ;

```

$$T = (S, \Sigma, \delta),$$



① Resolver Top Level Algorithm

Answer $\delta(S_6, \sigma_{ANSWER}) \rightarrow (S_{10})$

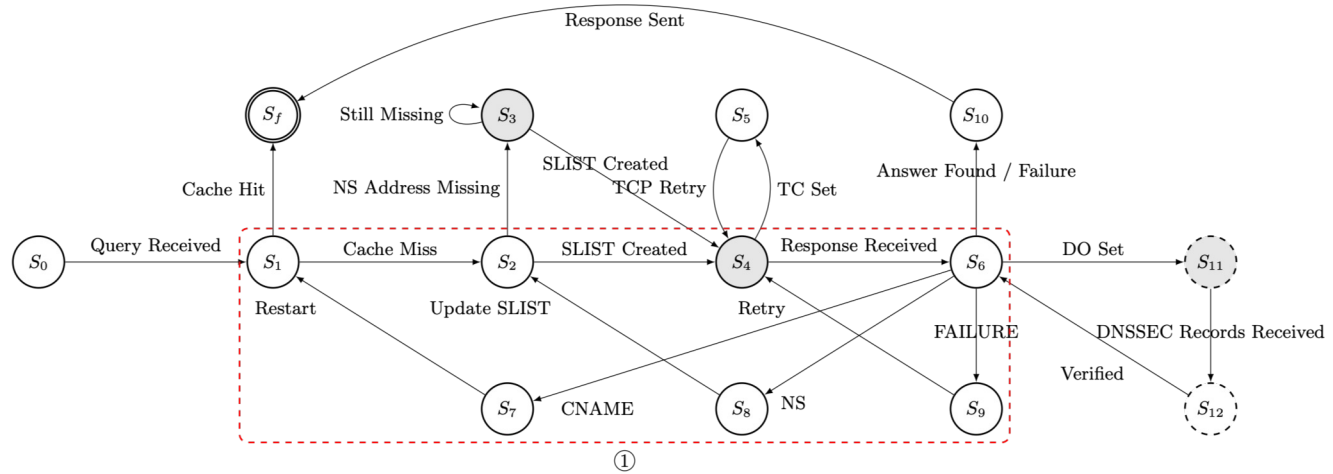
Referral $\delta(S_8, \sigma_{NS}) \rightarrow (S_2)$

Rewrite $\delta(S_7, \sigma_{CNAME}) \rightarrow (S_1)$

Failover $\delta(S_9, \sigma_{FAIL}) \rightarrow (S_4)$

State Transition

$$T = (S, \Sigma, \delta),$$



① Resolver Top Level Algorithm

Answer

$$\delta(S_6, \sigma_{ANSWER}) \rightarrow (S_{10})$$

Rewrite

$$\delta(S_7, \sigma_{CNAME}) \rightarrow (S_1)$$

Referral

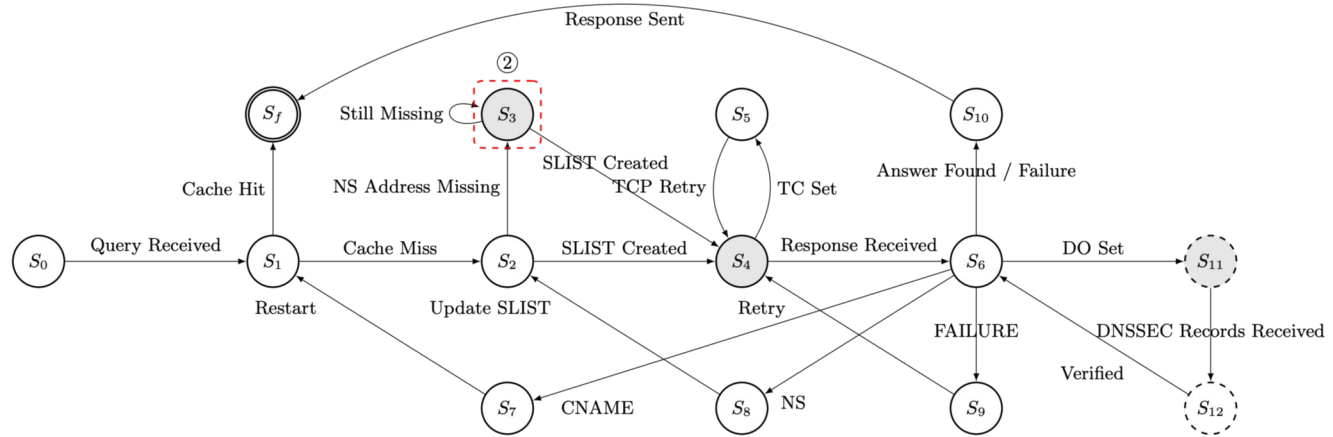
$$\delta(S_8, \sigma_{NS}) \rightarrow (S_2)$$

Failover

$$\delta(S_9, \sigma_{FAIL}) \rightarrow (S_4)$$

State Transition

$$T = (S, \Sigma, \delta),$$



① Resolver Top Level Algorithm

Answer

$$\delta(S_6, \sigma_{ANSWER}) \rightarrow (S_{10})$$

Rewrite

$$\delta(S_7, \sigma_{CNAME}) \rightarrow (S_1)$$

Referral

$$\delta(S_8, \sigma_{NS}) \rightarrow (S_2)$$

Failover

$$\delta(S_9, \sigma_{FAIL}) \rightarrow (S_4)$$

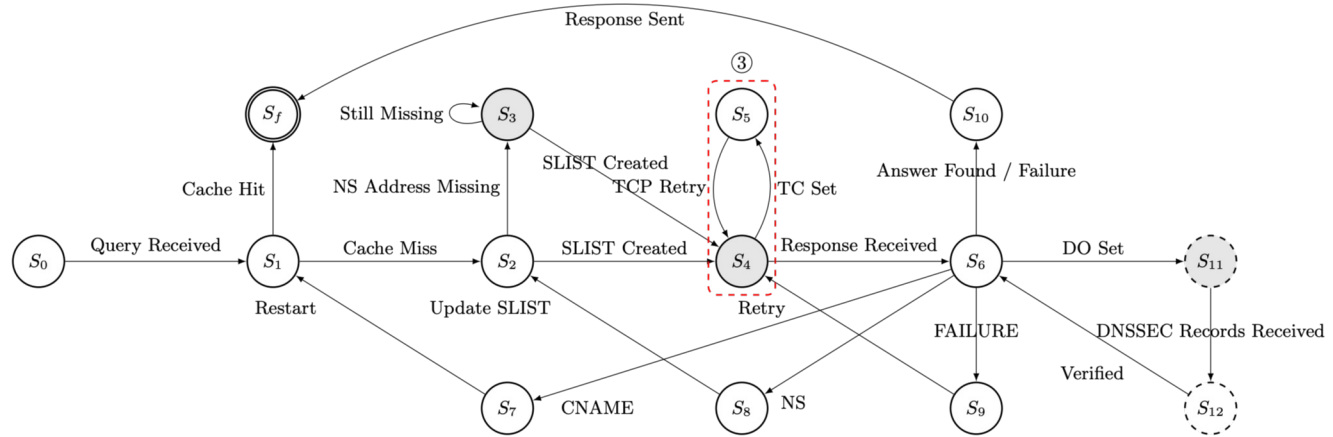
② Missing Nameserver Addresses

$$\delta(S_2, \sigma_{MISS}) \rightarrow (S_3)$$

$$\delta(S_3, \sigma_{MISS}) \rightarrow (S_3)$$

State Transition

$$T = (S, \Sigma, \delta),$$



① Resolver Top Level Algorithm

Answer

$$\delta(S_6, \sigma_{ANSWER}) \rightarrow (S_{10})$$

Rewrite

$$\delta(S_7, \sigma_{CNAME}) \rightarrow (S_1)$$

Referral

$$\delta(S_8, \sigma_{NS}) \rightarrow (S_2)$$

Failover

$$\delta(S_9, \sigma_{FAIL}) \rightarrow (S_4)$$

② Missing Nameserver Addresses

$$\delta(S_2, \sigma_{MISS}) \rightarrow (S_3)$$

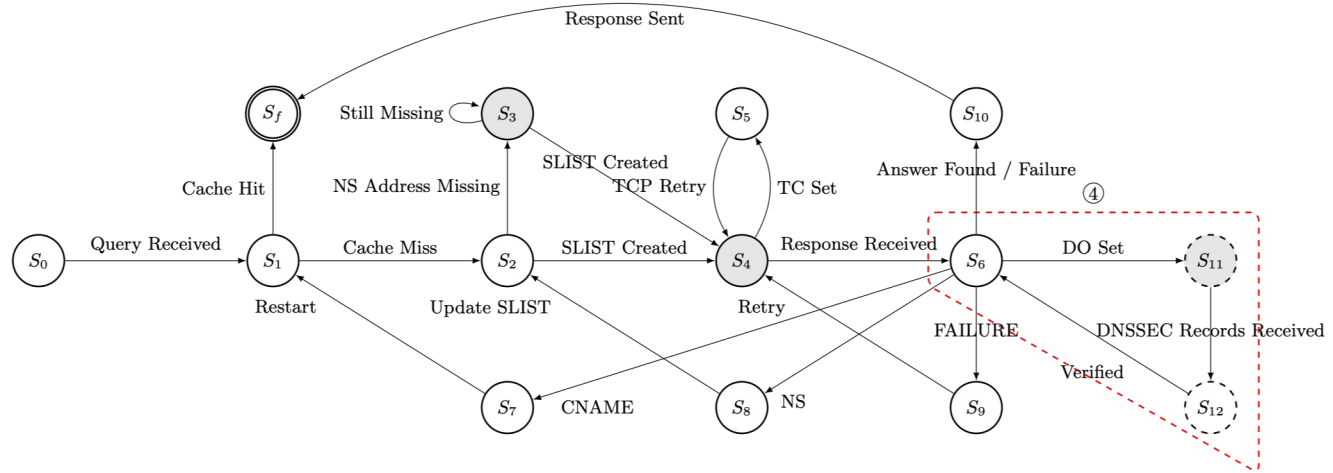
$$\delta(S_3, \sigma_{MISS}) \rightarrow (S_3)$$

③ TCP Retransmission

$$\delta(S_4, \sigma_{TC}) \rightarrow (S_5)$$

State Transition

$$T = (S, \Sigma, \delta),$$



① Resolver Top Level Algorithm

Answer

$$\delta(S_6, \sigma_{ANSWER}) \rightarrow (S_{10})$$

Rewrite

$$\delta(S_7, \sigma_{CNAME}) \rightarrow (S_1)$$

Referral

$$\delta(S_8, \sigma_{NS}) \rightarrow (S_2)$$

Failover

$$\delta(S_9, \sigma_{FAIL}) \rightarrow (S_4)$$

② Missing Nameserver Addresses

$$\delta(S_2, \sigma_{MISS}) \rightarrow (S_3)$$

$$\delta(S_3, \sigma_{MISS}) \rightarrow (S_3)$$

③ TCP Retransmission

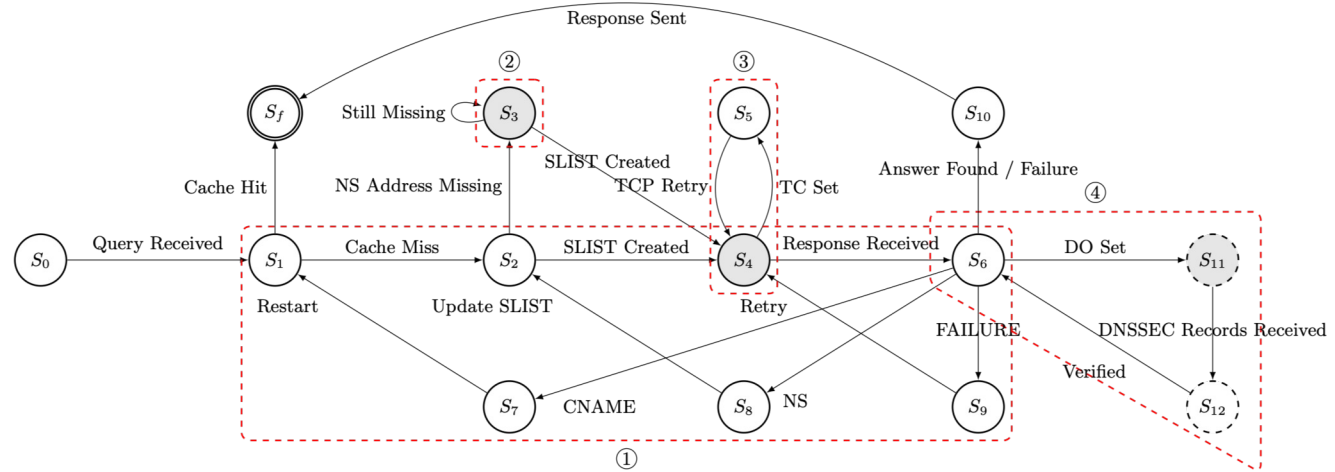
$$\delta(S_4, \sigma_{TC}) \rightarrow (S_5)$$

④ DNSSEC Validation

$$\delta(S_6, \sigma_{DNSSEC}) \rightarrow (S_{11})$$

State Transition

$$T = (S, \Sigma, \delta),$$



① Resolver Top Level Algorithm

Answer

$$\delta(S_6, \sigma_{ANSWER}) \rightarrow (S_{10})$$

Rewrite

$$\delta(S_7, \sigma_{CNAME}) \rightarrow (S_1)$$

Referral

$$\delta(S_8, \sigma_{NS}) \rightarrow (S_2)$$

Failover

$$\delta(S_9, \sigma_{FAILURE}) \rightarrow (S_4)$$

② Missing Nameserver Addresses

$$\delta(S_2, \sigma_{MISS}) \rightarrow (S_3)$$

$$\delta(S_3, \sigma_{MISS}) \rightarrow (S_3)$$

③ TCP Retransmission

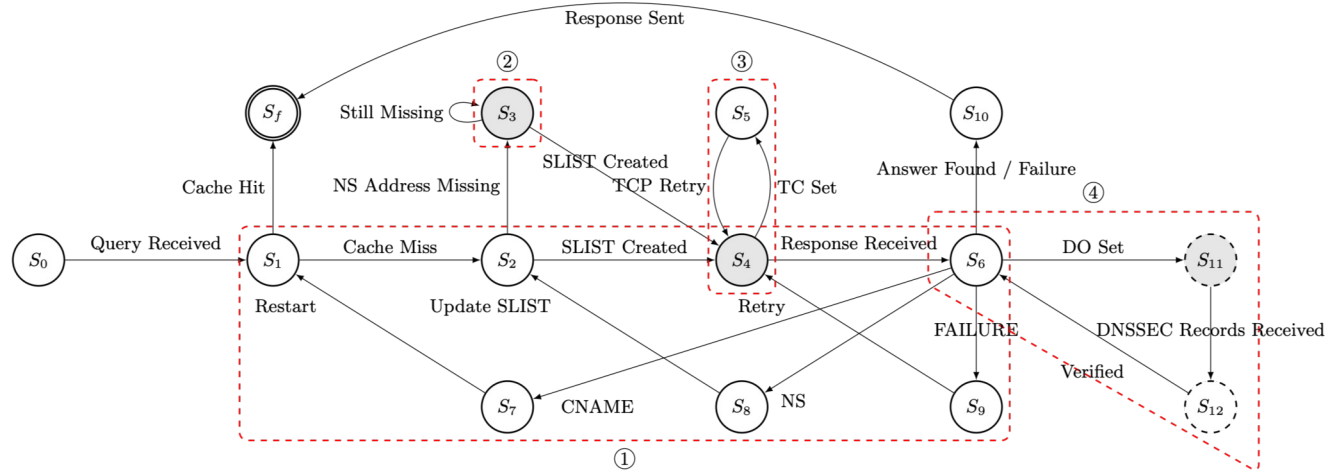
$$\delta(S_4, \sigma_{TC}) \rightarrow (S_5)$$

④ DNSSEC Validation

$$\delta(S_6, \sigma_{DNSSEC}) \rightarrow (S_{11})$$

State Transition

$$T = (S, \Sigma, \delta),$$



① Resolver Top Level Algorithm

Answer $\delta(S_6, \sigma_{ANSWER}) \rightarrow (S_{10})$
 Rewrite $\delta(S_7, \sigma_{CNAME}) \rightarrow (S_1)$
 Referral $\delta(S_8, \sigma_{NS}) \rightarrow (S_2)$
 Failover $\delta(S_9, \sigma_{FAIL}) \rightarrow (S_4)$

② Missing Nameserver Addresses

$\delta(S_2, \sigma_{MISS}) \rightarrow (S_3)$
 $\delta(S_3, \sigma_{MISS}) \rightarrow (S_3)$

③ TCP Retransmission

$\delta(S_4, \sigma_{TC}) \rightarrow (S_5)$


④ DNSSEC Validation

$\delta(S_6, \sigma_{DNSSEC}) \rightarrow (S_{11})$

Modeling state transitions makes resolver behavior tractable.


Transition Cost

$$T = (S, \Sigma, \delta, C),$$



Transition Cost

$$T = (S, \Sigma, \delta, C),$$



○ Outgoing Query Processing Cost

Transport Protocol

TCP vs UDP

Response Parsing

- RR Size
- RR Count

Data Transmission


- RR Size

Caching

- RR Size
- RR Count
- Cachability

Transition Cost

$$T = (S, \Sigma, \delta, C),$$



○ Outgoing Query Processing Cost

⊖ DNSSEC Validation Cost

Transport Protocol

TCP vs UDP

Response Parsing

- RR Size
- RR Count

- Algorithm
- Key Size
- RR Count

Data Transmission


- RR Size

Caching

- RR Size
- RR Count
- Cachability

Transition Cost

$$T = (S, \Sigma, \delta, C),$$



○ Outgoing Query Processing Cost

Transport Protocol
- TCP vs UDP

Data Transmission
- RR Size

○ DNSSEC Validation Cost

- Algorithm
- Key Size
- RR Count

○ Fixed Transition Cost


- Constant

Response Parsing
- RR Size
- RR Count

Caching
- RR Size
- RR Count
- Cachability

Transition Cost

$$T = (S, \Sigma, \delta, C),$$



○ Outgoing Query Processing Cost

○ DNSSEC Validation Cost

○ Fixed Transition Cost

Transport Protocol

TCP vs UDP

Response Parsing

- RR Size
- RR Count

- Algorithm
- Key Size
- RR Count

- Constant

Data Transmission

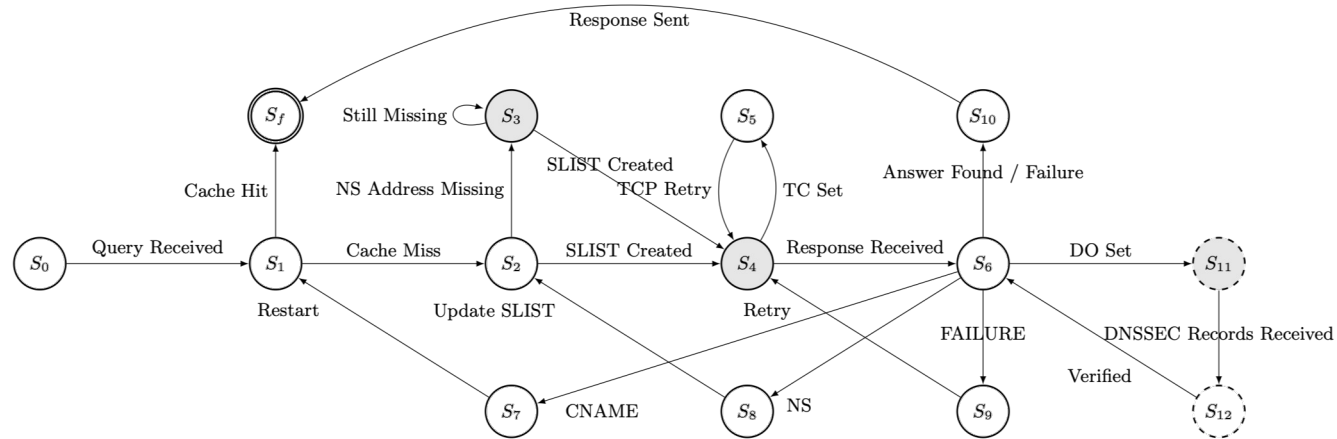
- RR Size

Caching

- RR Size
- RR Count
- Cachability

Annotating transitions with cost makes resolver work quantifiable.

Extended Transition System



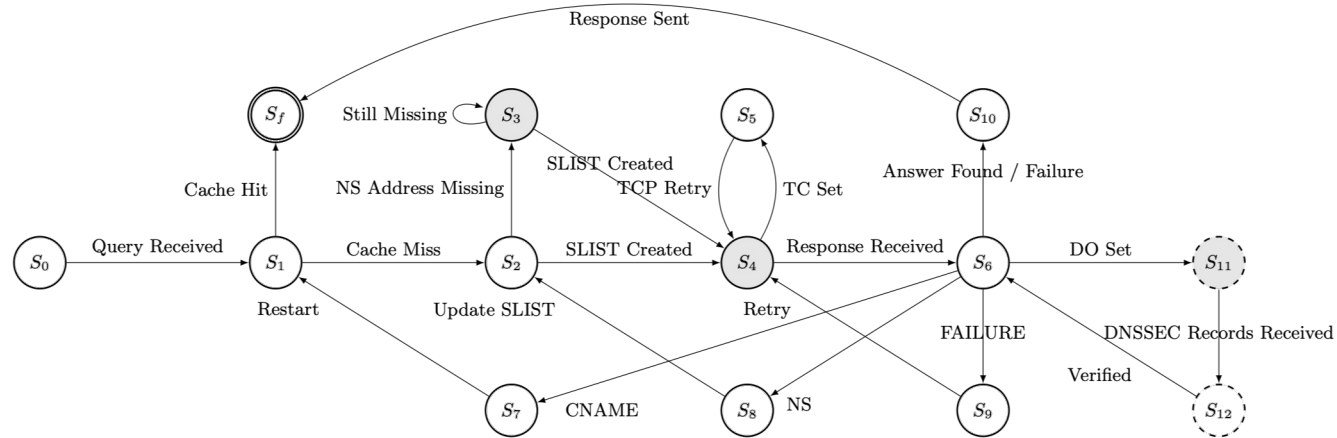
S states
 Σ inputs

stages of resolution
 server responses

δ transitions
 C cost

state transition function
 resource consumption

Extended Transition System




S states stages of resolution δ transitions state transition function
 Σ inputs server responses C cost resource consumption

V Counters
 track cycle traversal history

L Constraints
 bound the feasible state space

Extended Transition System

$$T = (S, \Sigma, V, \delta, L, C),$$



S states
 Σ inputs

stages of resolution
 server responses

δ transitions
 C cost

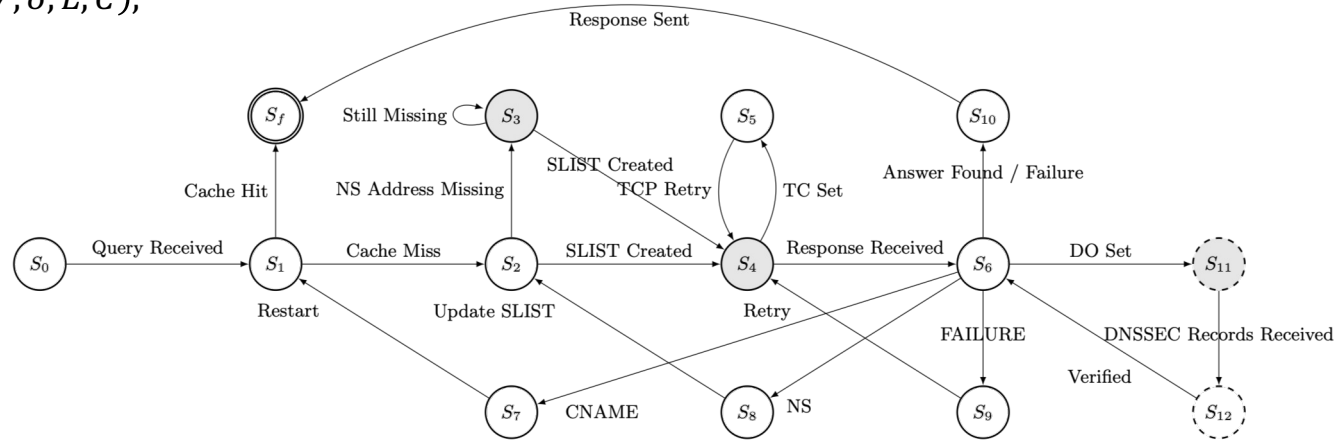
state transition function
 resource consumption

V Counters
 track cycle traversal history

L Constraints
 bound the feasible state space

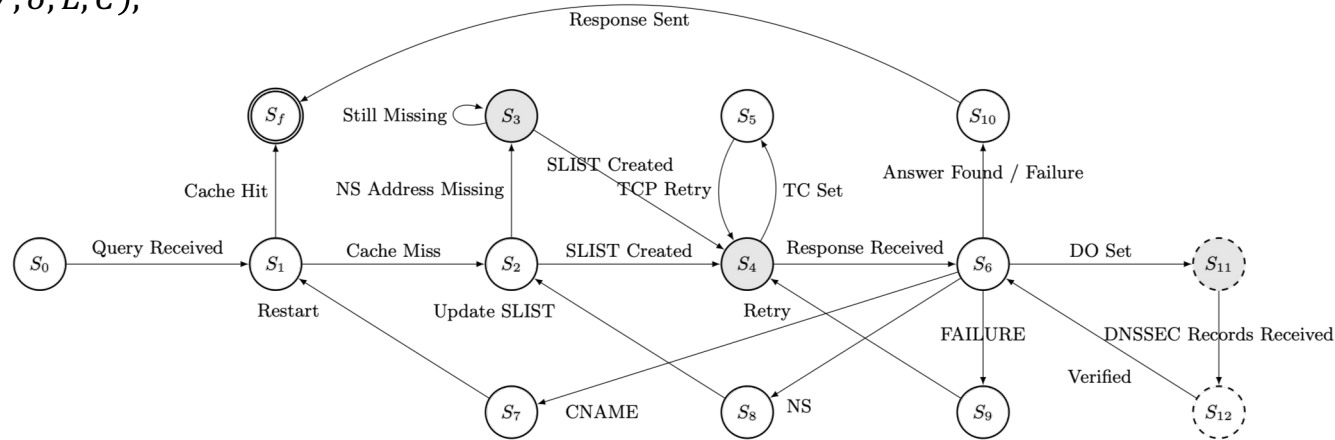
Problem Formulation

$$T = (S, \Sigma, V, \delta, L, C),$$



Problem Formulation

$$T = (S, \Sigma, V, \delta, L, C),$$



The total work of a resolution is the sum of costs along its state transition path.

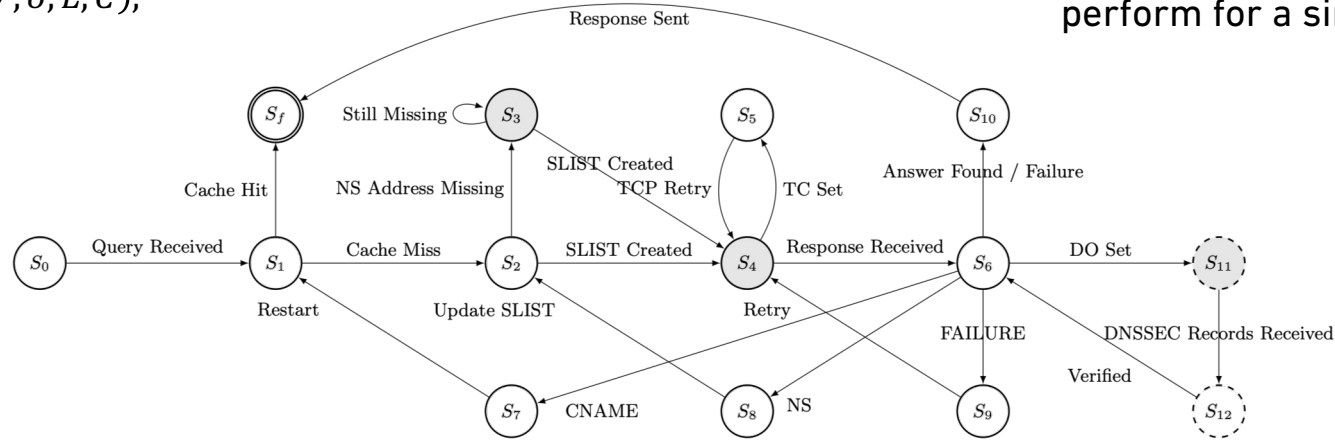
Work

$$W = \sum_{t=1}^T w_t$$

Problem Formulation

$$T = (S, \Sigma, V, \delta, L, C),$$



What is the maximum amount of work a resolver can perform for a single query?



The total work of a resolution is the sum of costs along its state transition path.

$$\begin{array}{ccc}
 \text{Work} & & \text{Upper Bound} \\
 W = \sum_{t=1}^T w_t & \longrightarrow & W^* = \max_{S_0 \rightarrow S_f} W
 \end{array}$$

Finding the Worst Case





States \rightarrow Nodes

Transitions \rightarrow Directed Edges

Cost \rightarrow Distance

Constraints \rightarrow Feasibility Bounds


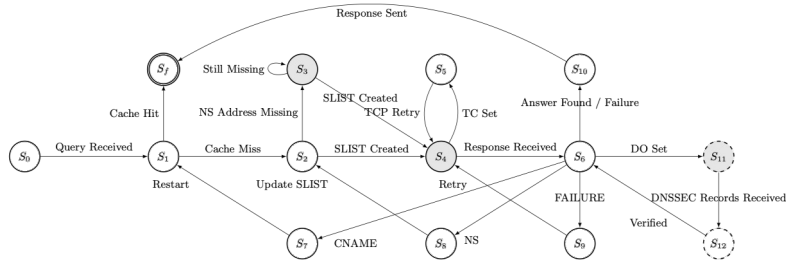
Finding the Worst Case



States \rightarrow Nodes Transitions \rightarrow Directed Edges Cost \rightarrow Distance Constraints \rightarrow Feasibility Bounds

Maximum Work \longrightarrow **Length of the Longest Path**

Finding the Worst Case





States \rightarrow Nodes Transitions \rightarrow Directed Edges Cost \rightarrow Distance Constraints \rightarrow Feasibility Bounds

Maximum Work \longrightarrow **Length of the Longest Path**

Constrained Longest Path Search

- Priority search — favor high-cost paths
- Constraint checking — prune infeasible branches at every step
- Termination guarantee — bounded by L

Finding the Worst Case



States \rightarrow Nodes Transitions \rightarrow Directed Edges Cost \rightarrow Distance Constraints \rightarrow Feasibility Bounds

Maximum Work \longrightarrow **Length of the Longest Path**

$W^*, \pi^*: S_0 \rightarrow S_1 \rightarrow S_2 \rightarrow \dots \rightarrow S_f$
— provably optimal traces

Constrained Longest Path Search

- Priority search — favor high-cost paths
- Constraint checking — prune infeasible branches at every step
- Termination guarantee — bounded by L

Correctness Guarantee

- State Completeness
- Additivity
- Boundedness/Termination

Instantiating Resolution Traces

```
> Init:  
  S0 → S1 → S2  
  
> for r = 1 to 11:  
  > for d = 1 to 32:  
    S2 → S4 → S5 → S4 → S6  
    S6 → S11 → S12 → S6  
    S6 → S8 → S2  
    S2 → S4 → S5 → S4 → S6  
    S6 → S11 → S12 → S6  
    S6 → S7 → S1 → S2  
  
> for d = 1 to 31:  
  S2 → S4 → S5 → S4 → S6  
  S6 → S11 → S12 → S6  
  S6 → S8 → S2  
  
> Final Answer:  
  S2 → S4 → S5 → S4 → S6  
  S6 → S11 → S12 → S6  
  S6 → S10 → Sf
```

Instantiating Resolution Traces

```
> Init:
  S0 → S1 → S2

> for r = 1 to 11:
  > for d = 1 to 32:
    S2 → S4 → S5 → S4 → S6
    S6 → S11 → S12 → S6
    S6 → S8 → S2
    S2 → S4 → S5 → S4 → S6
    S6 → S11 → S12 → S6
    S6 → S7 → S1 → S2

> for d = 1 to 31:
  S2 → S4 → S5 → S4 → S6
  S6 → S11 → S12 → S6
  S6 → S8 → S2

> Final Answer:
  S2 → S4 → S5 → S4 → S6
  S6 → S11 → S12 → S6
  S6 → S10 → Sf
```

A formal model and algorithm are necessary but insufficient
— standard nameservers cannot serve the traces they produce.

We need a system that can...

- Generates responses dynamically
- Session-aware across queries
- Full control over content, timing, ordering

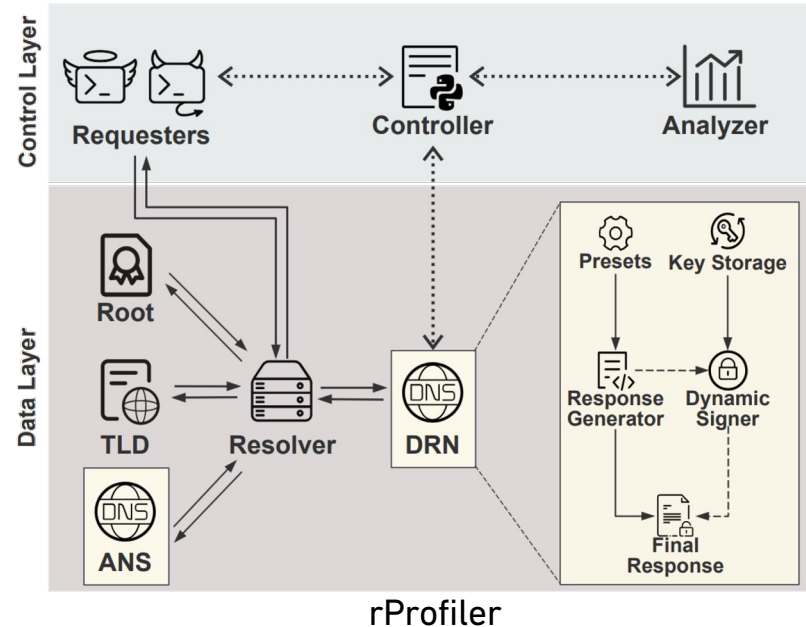
Instantiating Resolution Traces

```
> Init:  
  S0 → S1 → S2  
  
> for r = 1 to 11:  
  > for d = 1 to 32:  
    S2 → S4 → S5 → S4 → S6  
    S6 → S11 → S12 → S6  
    S6 → S8 → S2  
    S2 → S4 → S5 → S4 → S6  
    S6 → S11 → S12 → S6  
    S6 → S7 → S1 → S2  
  
> for d = 1 to 31:  
  S2 → S4 → S5 → S4 → S6  
  S6 → S11 → S12 → S6  
  S6 → S8 → S2  
  
> Final Answer:  
  S2 → S4 → S5 → S4 → S6  
  S6 → S11 → S12 → S6  
  S6 → S10 → Sf
```

A formal model and algorithm are necessary but insufficient
— standard nameservers cannot serve the traces they produce.

We need a system that can...

- Generates responses dynamically
- Session-aware across queries
- Full control over content, timing, ordering



From Constraints to Measured Resolver Work

On State Transitions

Parameters	BIND 9.21.1	Unbound 1.21.1	PowerDNS 5.1.3
max_restart	11	11	10
max_depth	32	128	50
max_fetch	5	6	10
max_referral	7	4	16
cname_reset	True	True	False
max_outgoing_query	384	128	50
tcp_retry_included	False	False	True
dnssec_included	False	False	True

From Constraints to Measured Resolver Work

On State Transitions

Parameters	BIND 9.21.1	Unbound 1.21.1	PowerDNS 5.1.3
max_restart	11	11	10
max_depth	32	128	50
max_fetch	5	6	10
max_referral	7	4	16
cname_reset	True	True	False
max_outgoing_query	384	128	50
tcp_retry_included	False	False	True
dnssec_included	False	False	True

On Transition Cost

Response	Type	Bind9 9.21.1	Unbound 1.21.1	PowerDNS 5.1.3
Answer	QTYPE	✓	✓	✓, ↔, ↓
	RRSIG	○	○	○
Authoritative	NS	✓	✓	✓, ↔, ↓
	DS	○	○	○
	RRSIG	○	○	○
Additional	A	✓	✓	○
	AAAA	✓	✓	✓, ↔, ↓
	RRSIG	○	○	○

From Constraints to Measured Resolver Work

On State Transitions

Parameters	BIND 9.21.1	Unbound 1.21.1	PowerDNS 5.1.3
max_restart	11	11	10
max_depth	32	128	50
max_fetch	5	6	10
max_referral	7	4	16
cname_reset	True	True	False
max_outgoing_query	384	128	50
tcp_retry_included	False	False	True
dnssec_included	False	False	True

On Transition Cost

Response	Type	Bind9 9.21.1	Unbound 1.21.1	PowerDNS 5.1.3
Answer	QTYPE	✓	✓	✓, ↔, ↓
	RRSIG	○	○	○
Authoritative	NS	✓	✓	✓, ↔, ↓
	DS	○	○	○
	RRSIG	○	○	○
Additional	A	✓	✓	○
	AAAA	✓	✓	✓, ↔, ↓
	RRSIG	○	○	○

Some constraints are implicit —
only discoverable through probing.

From Constraints to Measured Resolver Work

On State Transitions

Parameters	BIND 9.21.1	Unbound 1.21.1	PowerDNS 5.1.3
max_restart	11	11	10
max_depth	32	128	50
max_fetch	5	6	10
max_referral	7	4	16
cname_reset	True	True	False
max_outgoing_query	384	128	50
tcp_retry_included	False	False	True
dnssec_included	False	False	True

On Transition Cost

Response	Type	Bind9 9.21.1	Unbound 1.21.1	PowerDNS 5.1.3
Answer	QTYPE	✓	✓	✓, ↔, ↓
	RRSIG	○	○	○
	NS	✓	✓	✓, ↔, ↓
Authoritative	DS	○	○	○
	RRSIG	○	○	○
	A	✓	✓	○
Additional	AAAA	✓	✓	✓, ↔, ↓
	RRSIG	○	○	○

Some constraints are implicit —
only discoverable through probing.

Sample Traces

```
# TRX = 1536
> Init:
S0 → S1 → S2
> for r = 1 to 11:
  > for d = 1 to 32:
    S2 → S4 → S5 → S4 → S6
    S6 → S11 → S12 → S6
    S6 → S8 → S2
  S2 → S4 → S5 → S4 → S6
  S6 → S11 → S12 → S6
  S6 → S7 → S1 → S2
  > for d = 1 to 31:
    S2 → S4 → S5 → S4 → S6
    S6 → S11 → S12 → S6
    S6 → S8 → S2
  > Final Answer:
    S2 → S4 → S5 → S4 → S6
    S6 → S11 → S12 → S6
    S6 → S10 → S_f
```

(a) Bind9

```
# TRX = 512
> Init:
S0 → S1 → S2
> for d = 1 to 127:
  S2 → S4 → S5 → S4 → S6
  S6 → S11 → S12 → S6
  S6 → S8 → S2
  > Final Answer:
    S2 → S4 → S5 → S4 → S6
    S6 → S11 → S12 → S6
    S6 → S10 → S_f
```

(b) Unbound

```
# TRX = 96
(with IPv6)
> Init:
S0 → S1 → S2
> for d = 1 to 15:
  S2 → S4 → S5 → S4 → S6
  S6 → S11 → S12 → S6
  S6 → S8 → S2
  > Final Answer:
    S2 → S4 → S5 → S4 → S6
    S6 → S11 → S12 → S6
    S6 → S10 → S_f
```

(c) PowerDNS

From Constraints to Measured Resolver Work

On State Transitions

Parameters	BIND 9.21.1	Unbound 1.21.1	PowerDNS 5.1.3
max_restart	11	11	10
max_depth	32	128	50
max_fetch	5	6	10
max_referral	7	4	16
cname_reset	True	True	False
max_outgoing_query	384	128	50
tcp_retry_included	False	False	True
dnssec_included	False	False	True

On Transition Cost

Response	Type	Bind9 9.21.1	Unbound 1.21.1	PowerDNS 5.1.3
Answer	QTYPE	✓	✓	✓, ↔, ↓
	RRSIG	○	○	○
Authoritative	NS	✓	✓	✓, ↔, ↓
	DS	○	○	○
	RRSIG	○	○	○
Additional	A	✓	✓	○
	AAAA	✓	✓	✓, ↔, ↓
	RRSIG	○	○	○

Some constraints are implicit —
only discoverable through probing.

Sample Traces

```
# TRX = 1536
> Init:
S0 → S1 → S2
> for r = 1 to 11:
  > for d = 1 to 32:
    S2 → S4 → S5 → S4 → S6
    S6 → S11 → S12 → S6
    S6 → S8 → S2
  S2 → S4 → S5 → S4 → S6
  S6 → S11 → S12 → S6
  S6 → S7 → S1 → S2
  > for d = 1 to 31:
    S2 → S4 → S5 → S4 → S6
    S6 → S11 → S12 → S6
    S6 → S8 → S2
  > Final Answer:
    S2 → S4 → S5 → S4 → S6
    S6 → S11 → S12 → S6
    S6 → S10 → S_f
```

(a) Bind9

```
# TRX = 512
> Init:
S0 → S1 → S2
> for d = 1 to 127:
  S2 → S4 → S5 → S4 → S6
  S6 → S11 → S12 → S6
  S6 → S8 → S2
  > Final Answer:
    S2 → S4 → S5 → S4 → S6
    S6 → S11 → S12 → S6
    S6 → S10 → S_f
```

(b) Unbound

```
# TRX = 96
(with IPv6)
> Init:
S0 → S1 → S2
> for d = 1 to 15:
  S2 → S4 → S5 → S4 → S6
  S6 → S11 → S12 → S6
  S6 → S8 → S2
  > Final Answer:
    S2 → S4 → S5 → S4 → S6
    S6 → S11 → S12 → S6
    S6 → S10 → S_f
```

(c) PowerDNS


One single query can trigger...

Resolver	Transactions	Packets	Volume	Validation	Cache
BIND	1459 / 1536	12369	38.84 M	377	35.57 M
Unbound	493 / 512	5381	13.12 M	132	14.28 M
PowerDNS [†]	80 / 96	363	2.13 M	30	3.84 M

30 QPS Is More Than Enough To Exhaust Resolvers


30 QPS Is More Than Enough To Exhaust Resolvers

BIND9 9.21.1




30 QPS Is More Than Enough To Exhaust Resolvers

BIND9 9.21.1




Unbound 1.21.1




30 QPS Is More Than Enough To Exhaust Resolvers


BIND9 9.21.1



Unbound 1.21.1




PowerDNS 5.1.3




30 QPS Is More Than Enough To Exhaust Resolvers


BIND9 9.21.1



Unbound 1.21.1



PowerDNS 5.1.3



Takeaway

Resolver work can be and should be formally defined and computed

The recommended priorities for the resolver designer are:

1. Bound the amount of work (packets sent, parallel processes started)...

– RFC1034

Takeaway

Resolver work can be and should be formally defined and computed

The recommended priorities for the resolver designer are:

1. Bound the amount of work (packets sent, parallel processes started)...
- RFC1034

Resolvers remain vulnerable even with all mitigations in place

Takeaway

Resolver work can be and should be formally defined and computed

The recommended priorities for the resolver designer are:

1. Bound the amount of work (packets sent, parallel processes started)...

– RFC1034

Resolvers remain vulnerable even with all mitigations in place

Performance degradation is multi-faceted

Takeaway

Resolver work can be and should be formally defined and computed

The recommended priorities for the resolver designer are:

1. Bound the amount of work (packets sent, parallel processes started)...

– RFC1034

Resolvers remain vulnerable even with all mitigations in place

Performance degradation is multi-faceted

A principled foundation for future resolver design

Towards End-to-End Availability Guarantees

- **Understand resolver work**
- RFC Disambiguation
- Proactive resource scheduling

Takeaway

Resolver work can be and should be formally defined and computed

The recommended priorities for the resolver designer are:

1. Bound the amount of work (packets sent, parallel processes started)...

– RFC1034

Resolvers remain vulnerable even with all mitigations in place

Performance degradation is multi-faceted

A principled foundation for future resolver design

Towards End-to-End Availability Guarantees

- **Understand resolver work**
- RFC Disambiguation
- Proactive resource scheduling



Resolve the Unresolved: Systematic
Work Profiling for DNS Resolvers.

Contact: liwen.xu@inf.ethz.ch