Secure Name Resolution

Christian Grothoff

Berner Fachhochschule

10.11.2017

"The Domain Name System is the Achilles heel of the Web." –Tim Berners-Lee

Background: Efficient Set Union

(based on "What's the difference? Efficient Set Reconciliation without Prior Context", Eppstein et al., SIGCOMM'11)

- Alice and Bob have sets A and B
- The sets are very large
- ... but their symmetric difference $\delta := |(A - B) \cup (B - A)|$ is small
- Now Alice wants to know $B - A$ (the elements she is missing)
- ... and Bob $A - B$ (the elements he is missing)
- How can Alice and Bob do this efficiently?
  - w.r.t. communication and computation

Bad Solution

- Naive approach: Alice sends A to Bob, Bob sends $B - A$ back to Alice
- . . . or vice versa.

- Communication cost: $O(|A| + |B|)$:(
- Ideally, we want to do it in $O(\delta)$.
- First improvement: Do not send elements of A and B, but send/request hashes. Still does not improve complexity :(

- We need some more fancy data structure!

Bloom Filters

Constant size data structure that "summarizes" a set.

Operations:

$d = \mathit{NewBF}(\mathit{size})$ Create a new, empty bloom filter.

$\mathit{Insert}(d, e)$ Insert element e into the BF d.

$b = \mathit{Contains}(d, e)$ Check if BF d contains element e.

$b \in \{\text{"Definitely not in set"}, \text{"Probably in set"}\}$

BF: Insert

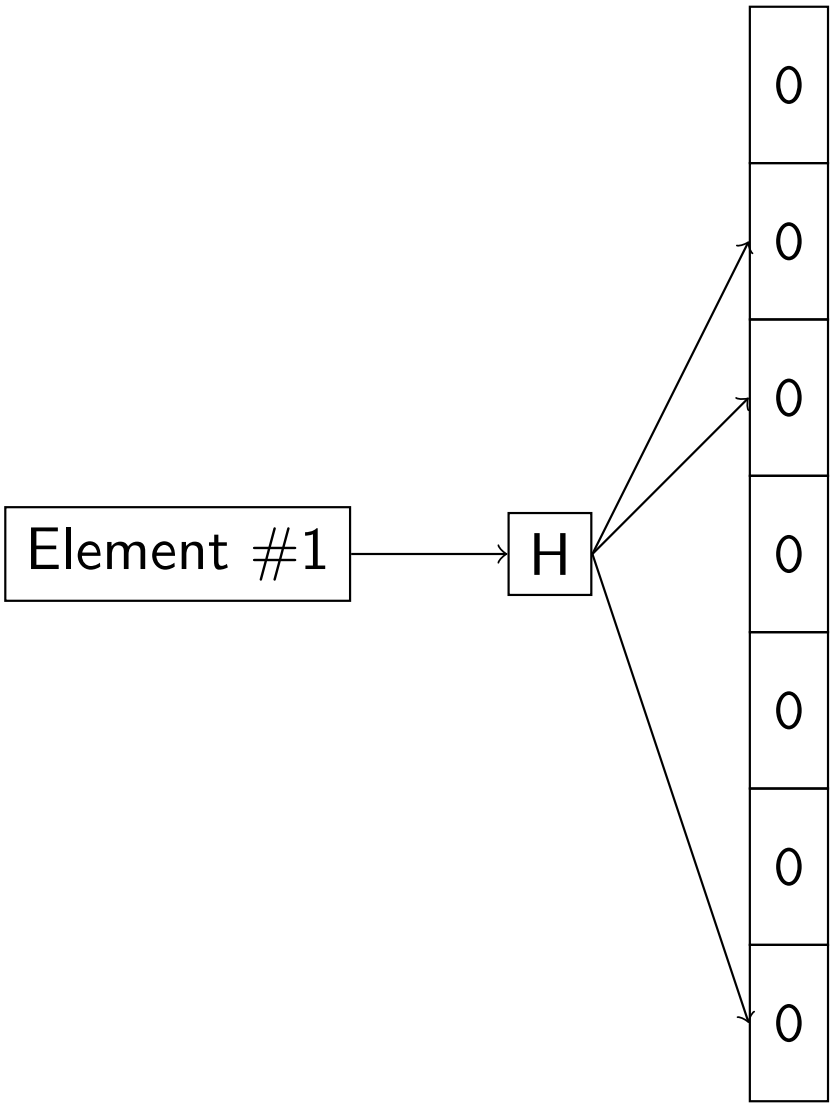

$H(\text{Element } \#1) = (2, 3, 7)$

BF: Insert

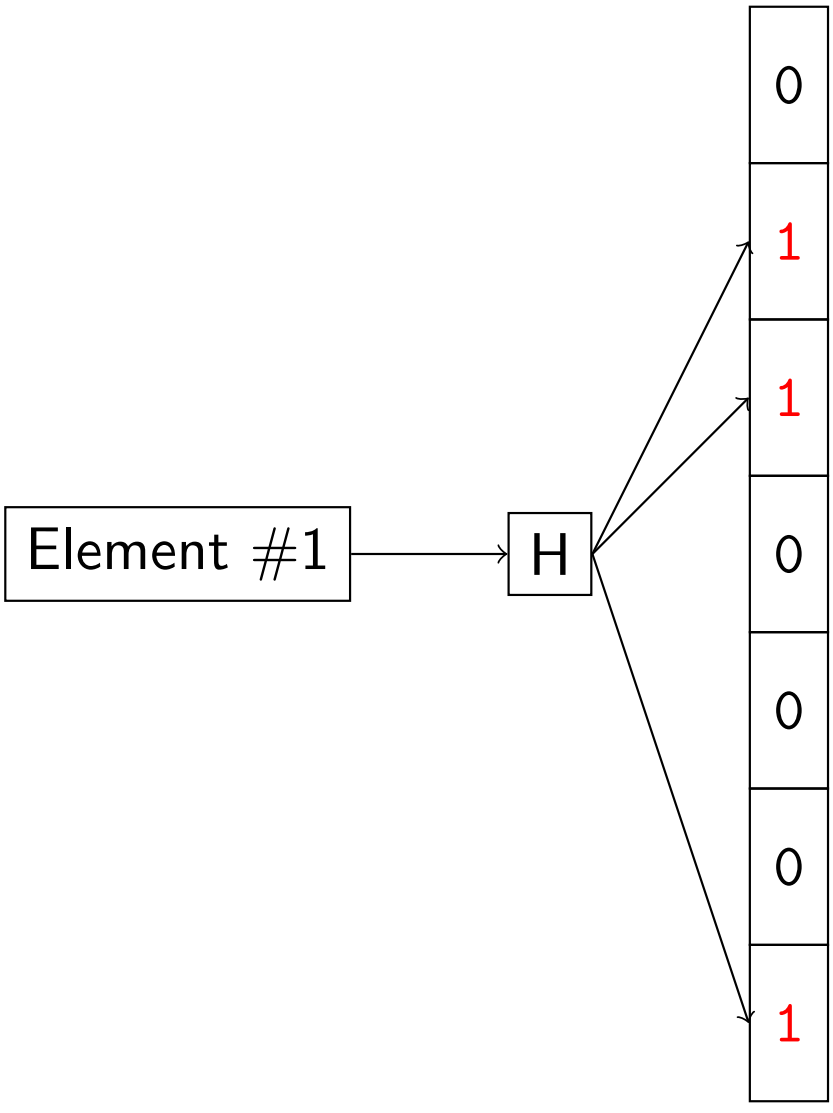

$H(\text{Element } \#1) = (2, 3, 7)$

BF: Insert

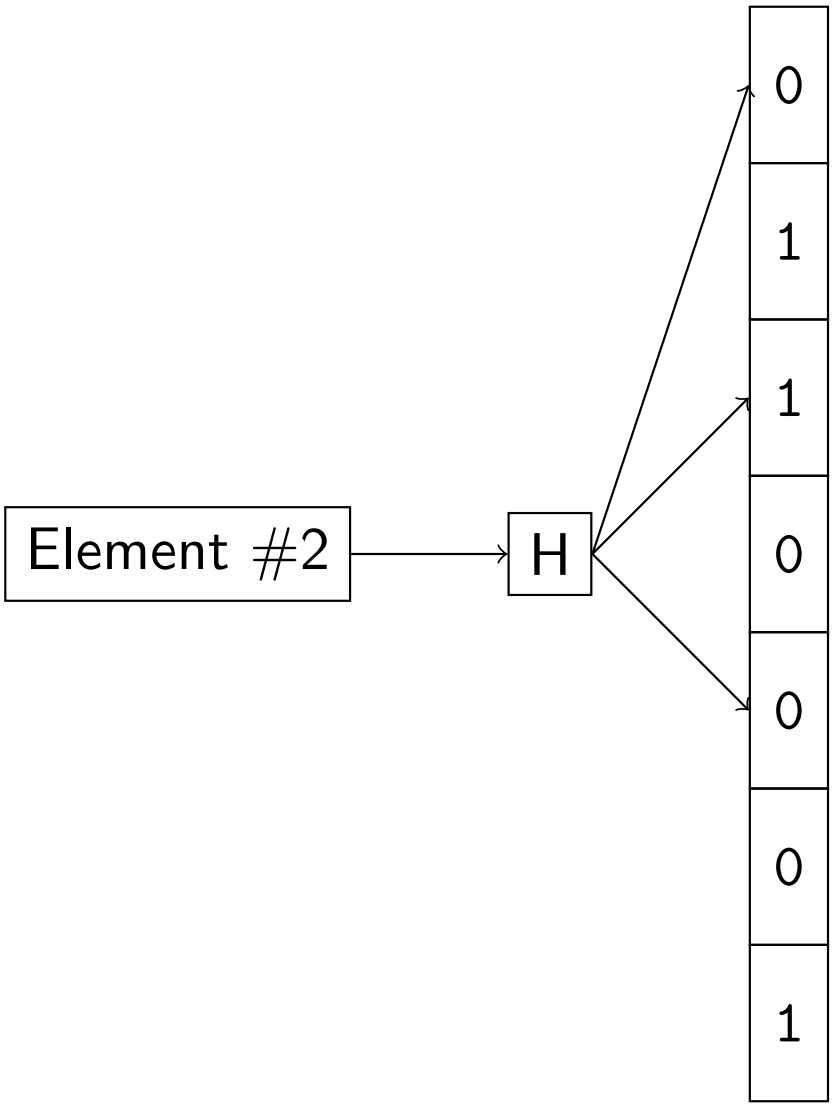

$H(\text{Element } \#1) = (2, 3, 7)$
$H(\text{Element } \#2) = (1, 3, 5)$

BF: Insert

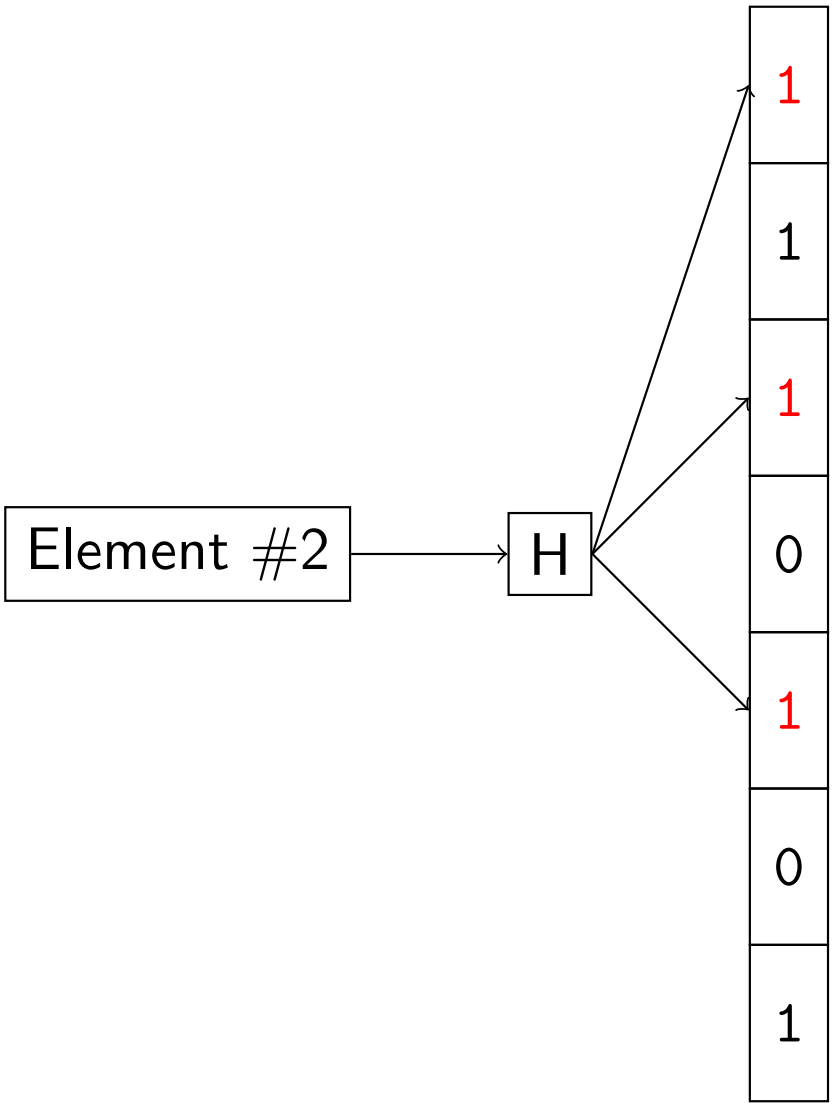

$H(\text{Element } \#1) = (2, 3, 7)$
$H(\text{Element } \#2) = (1, 3, 5)$

BF: Membership Test

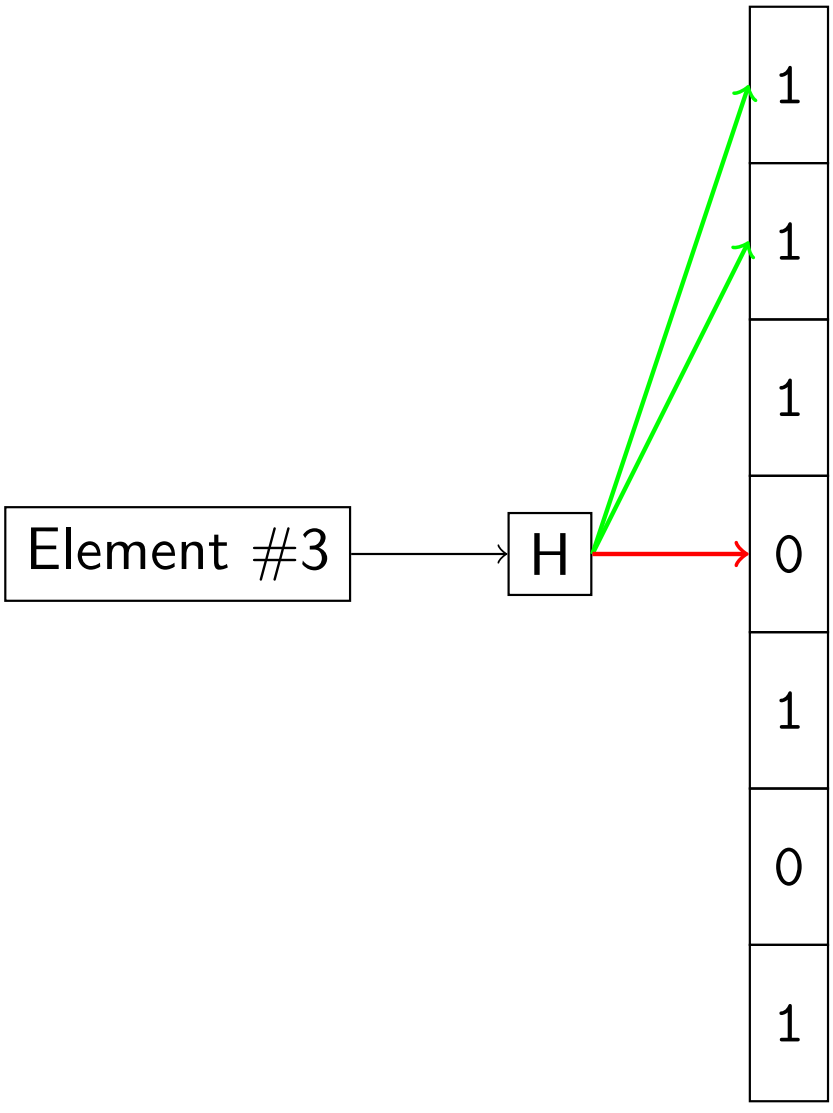

$H(\text{Element } \#1) = (2, 3, 7)$
$H(\text{Element } \#2) = (1, 3, 5)$

BF: Membership Test (false positive)

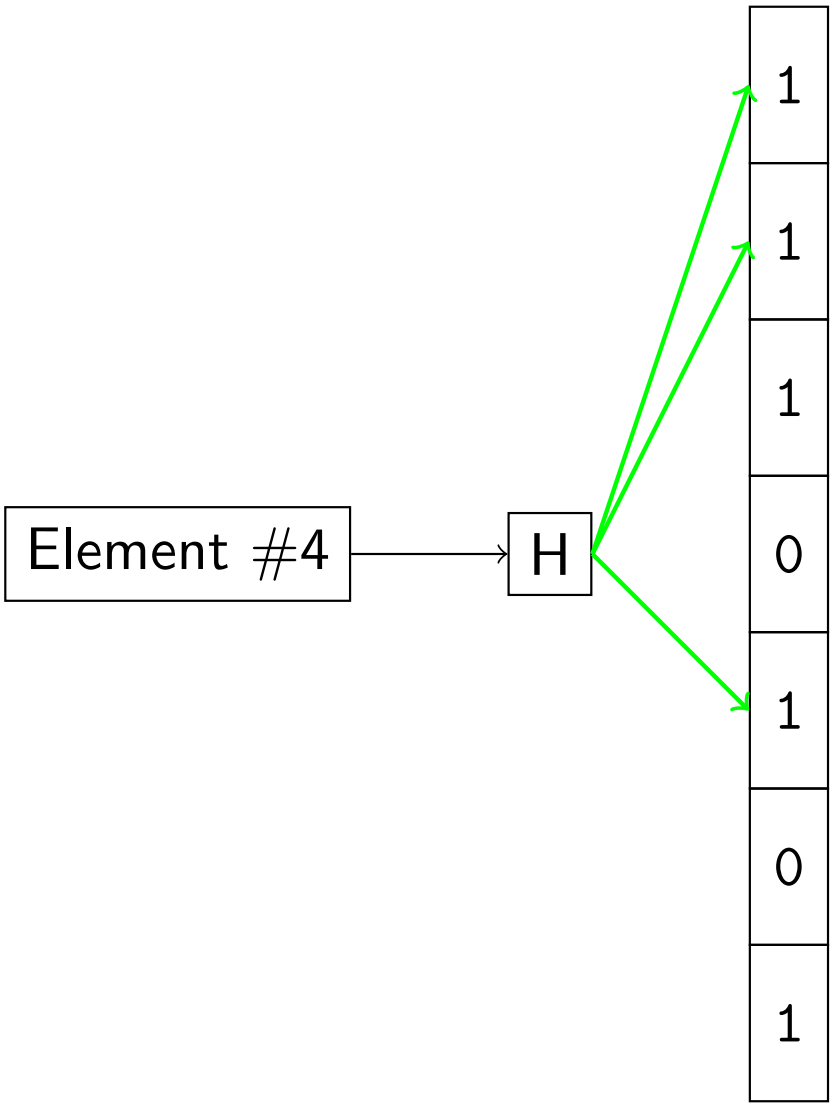

$H(\text{Element } \#1) = (2, 3, 7)$
$H(\text{Element } \#2) = (1, 3, 5)$

Counting Bloom Filters

BF where buckets hold a **positive integer**.

Additional Operation:

Remove(d, e) Remove element from the CBF d.

$\Rightarrow$ False negatives only when removing a non-existing element.

Invertible Bloom Filters

Similar to CBF, but

- Allow **negative counts**
- Additionaly store **(XOR-)sum of hashes** in buckets.

Additional Operations:

$(e, r) = \mathit{Extract}(d)$ Extract an element (e) from the IBF d, with result code $r \in \{\mathit{left}, \mathit{right}, \mathit{done}, \mathit{fail}\}$

$d' = \mathit{SymDiff}(d_1, d_2)$ Create an IBF that represents the symmetric difference of d_1 and d_2.

IBF: Insert

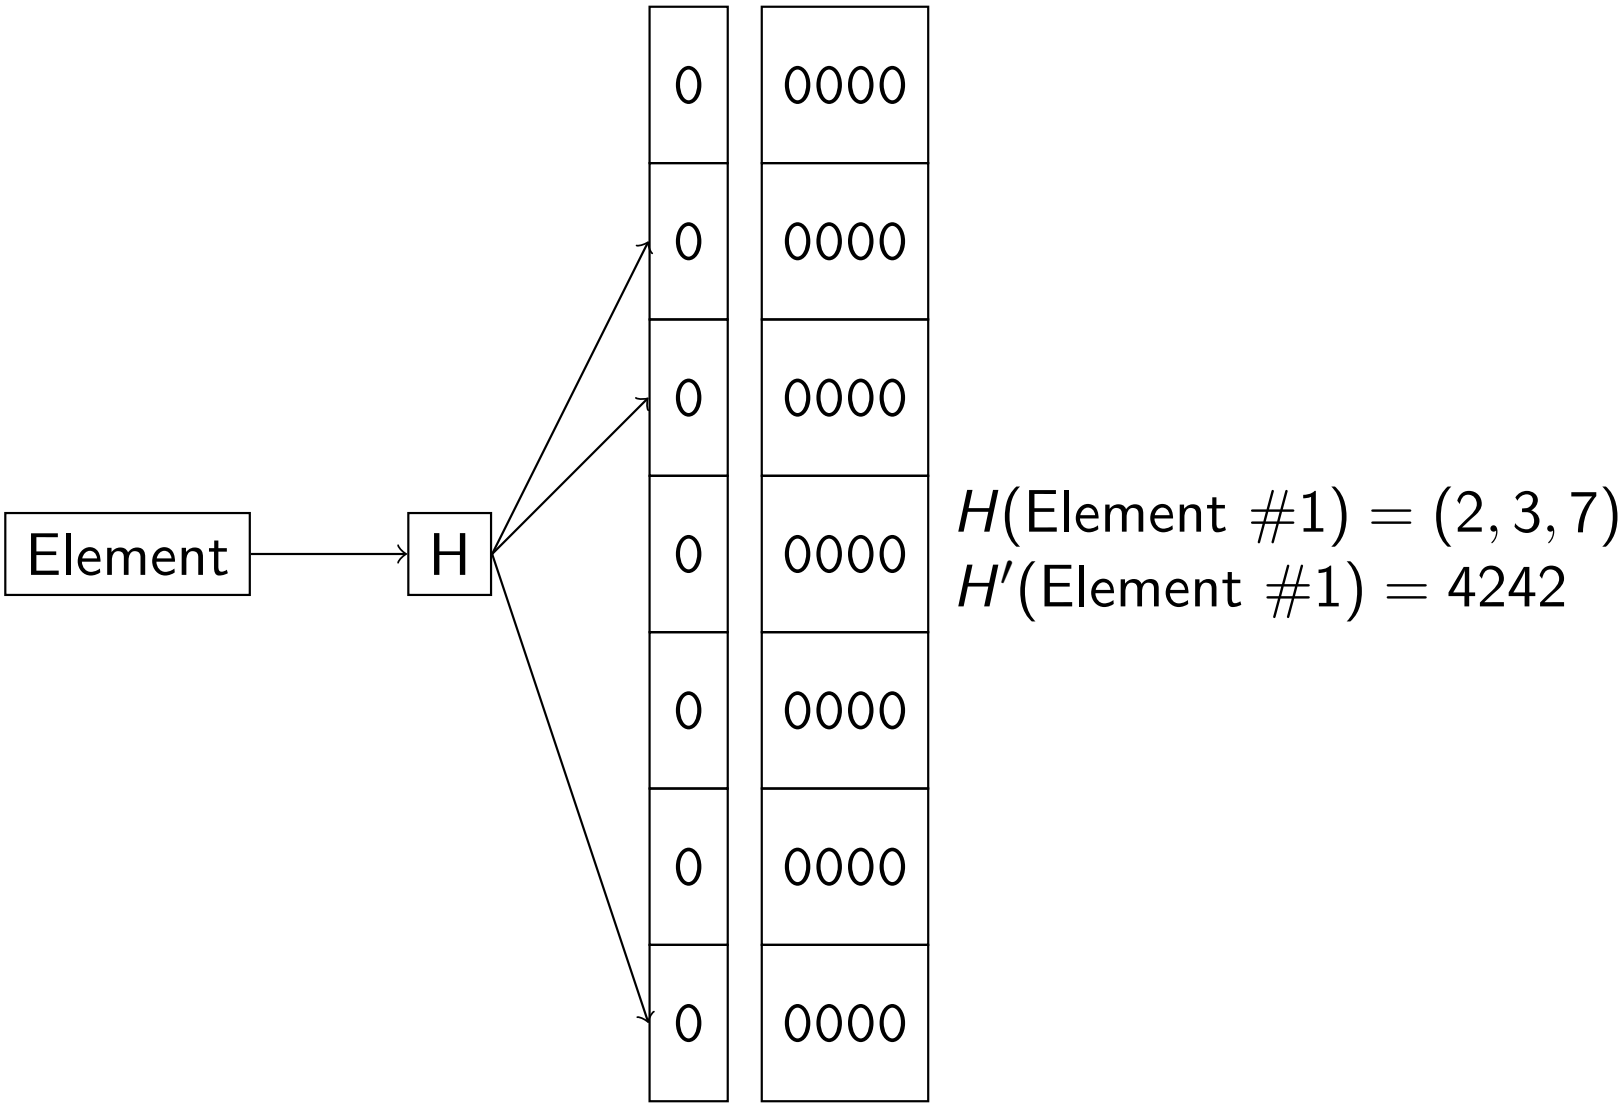
Element
H
0
0
0
0
0
0
0
0000
0000
0000
0000
0000
0000
0000
$H(\text{Element } \#1) = (2, 3, 7)$
$H'(\text{Element } \#1) = 4242$

IBF: Insert

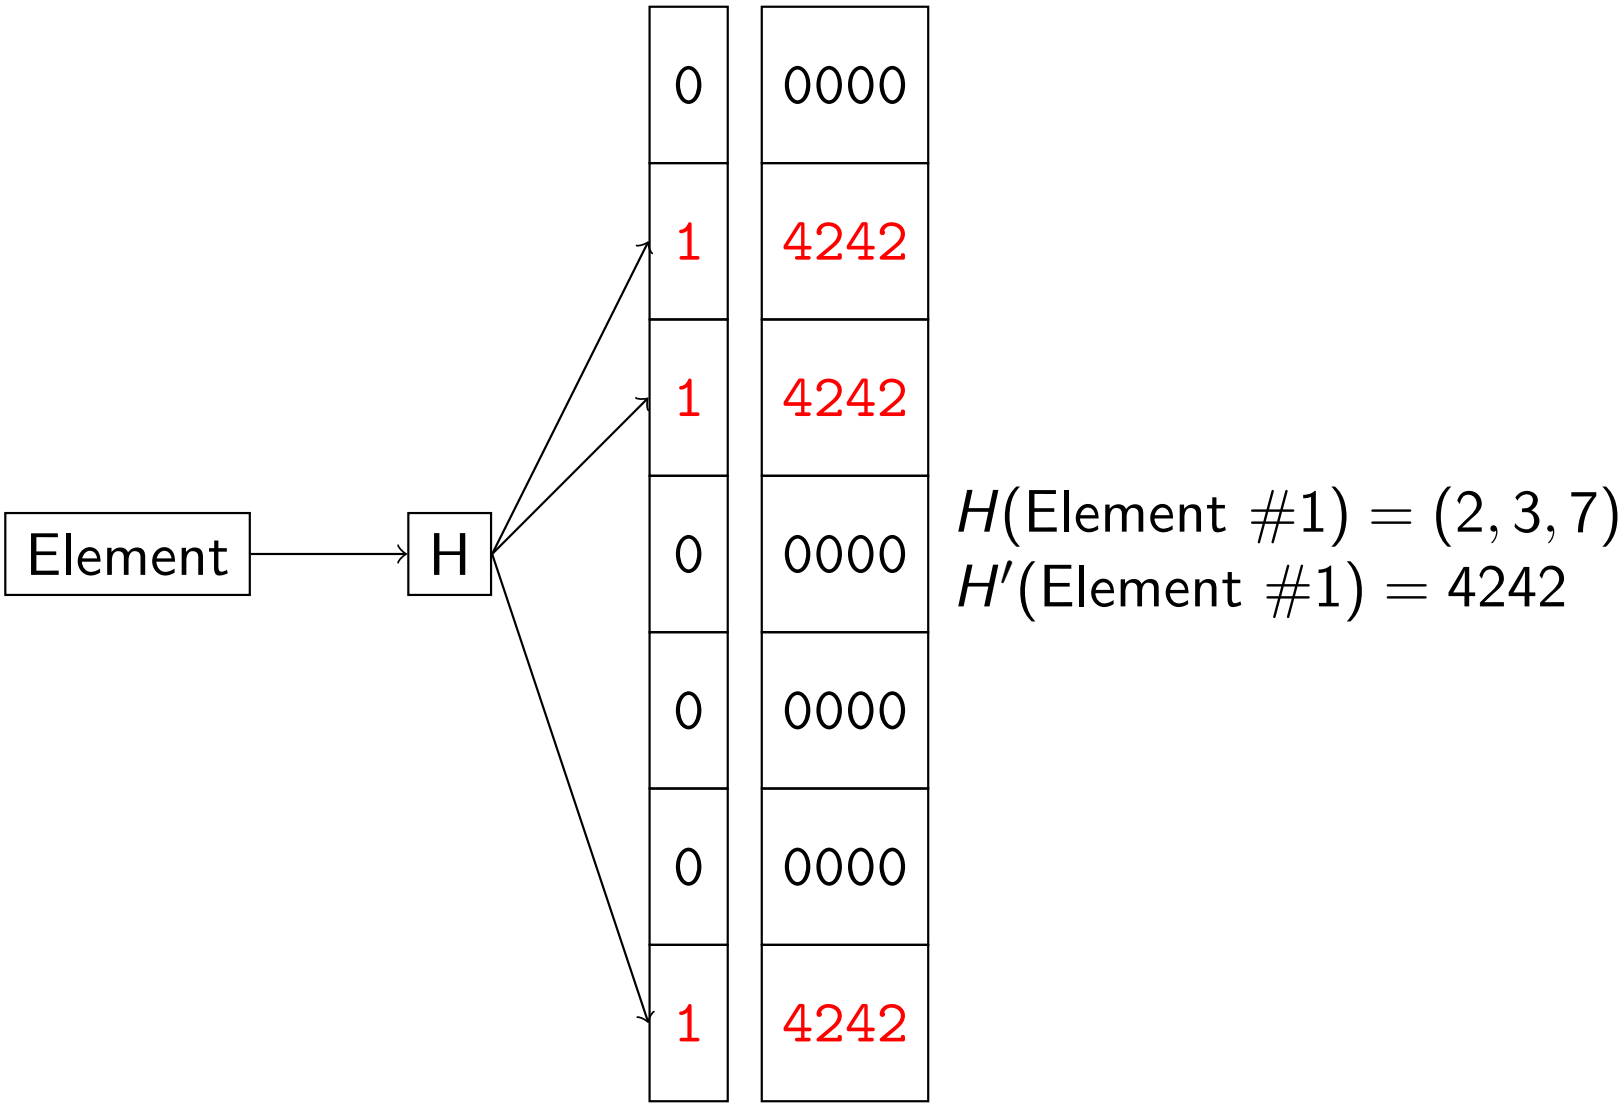
Element
H
0
0000
1
4242
1
4242
0
0000
0
0000
0
0000
1
4242
$H(\text{Element } \#1) = (2, 3, 7)$
$H'(\text{Element } \#1) = 4242$

IBF: Insert

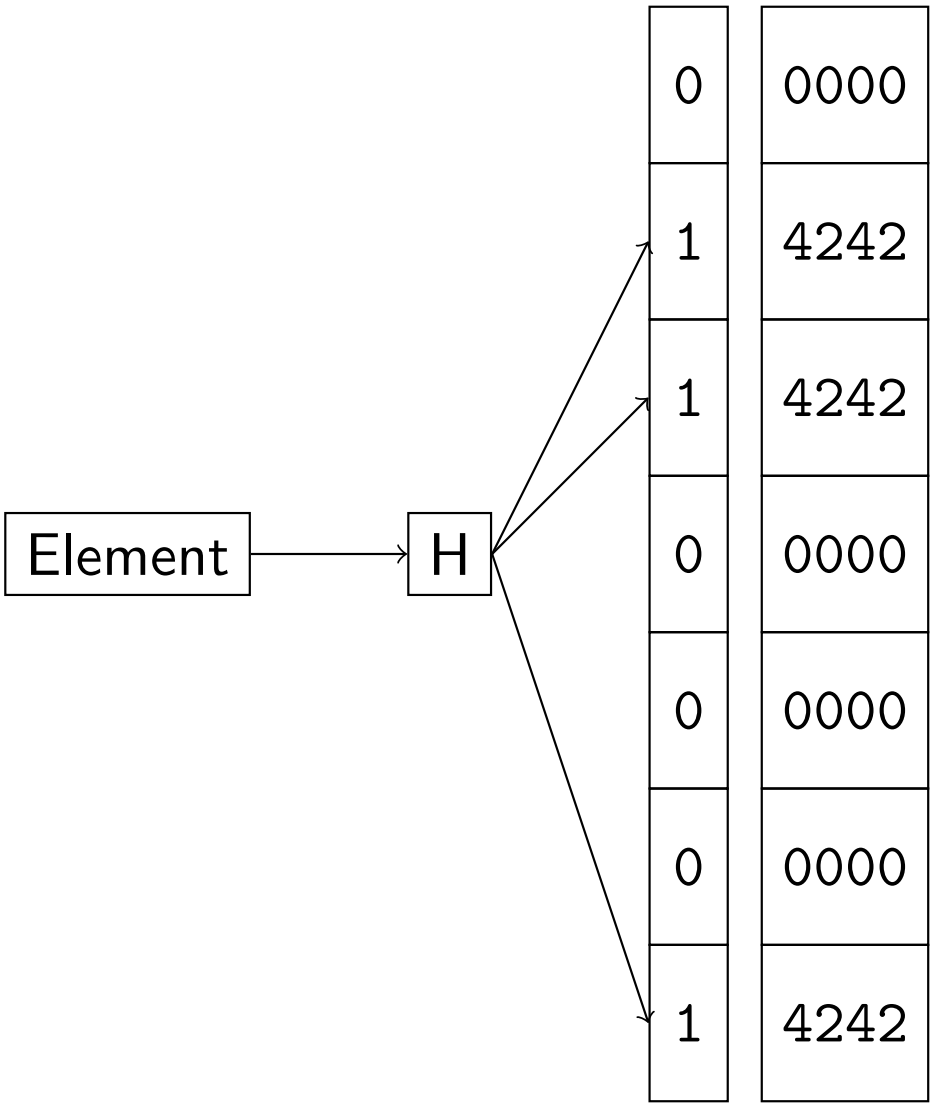

$H(\text{Element } \#1) = (2, 3, 7)$
$H'(\text{Element } \#1) = 4242$
$H(\text{Element } \#2) = (1, 3, 5)$
$H'(\text{Element } \#2) = 0101$

IBF: Insert

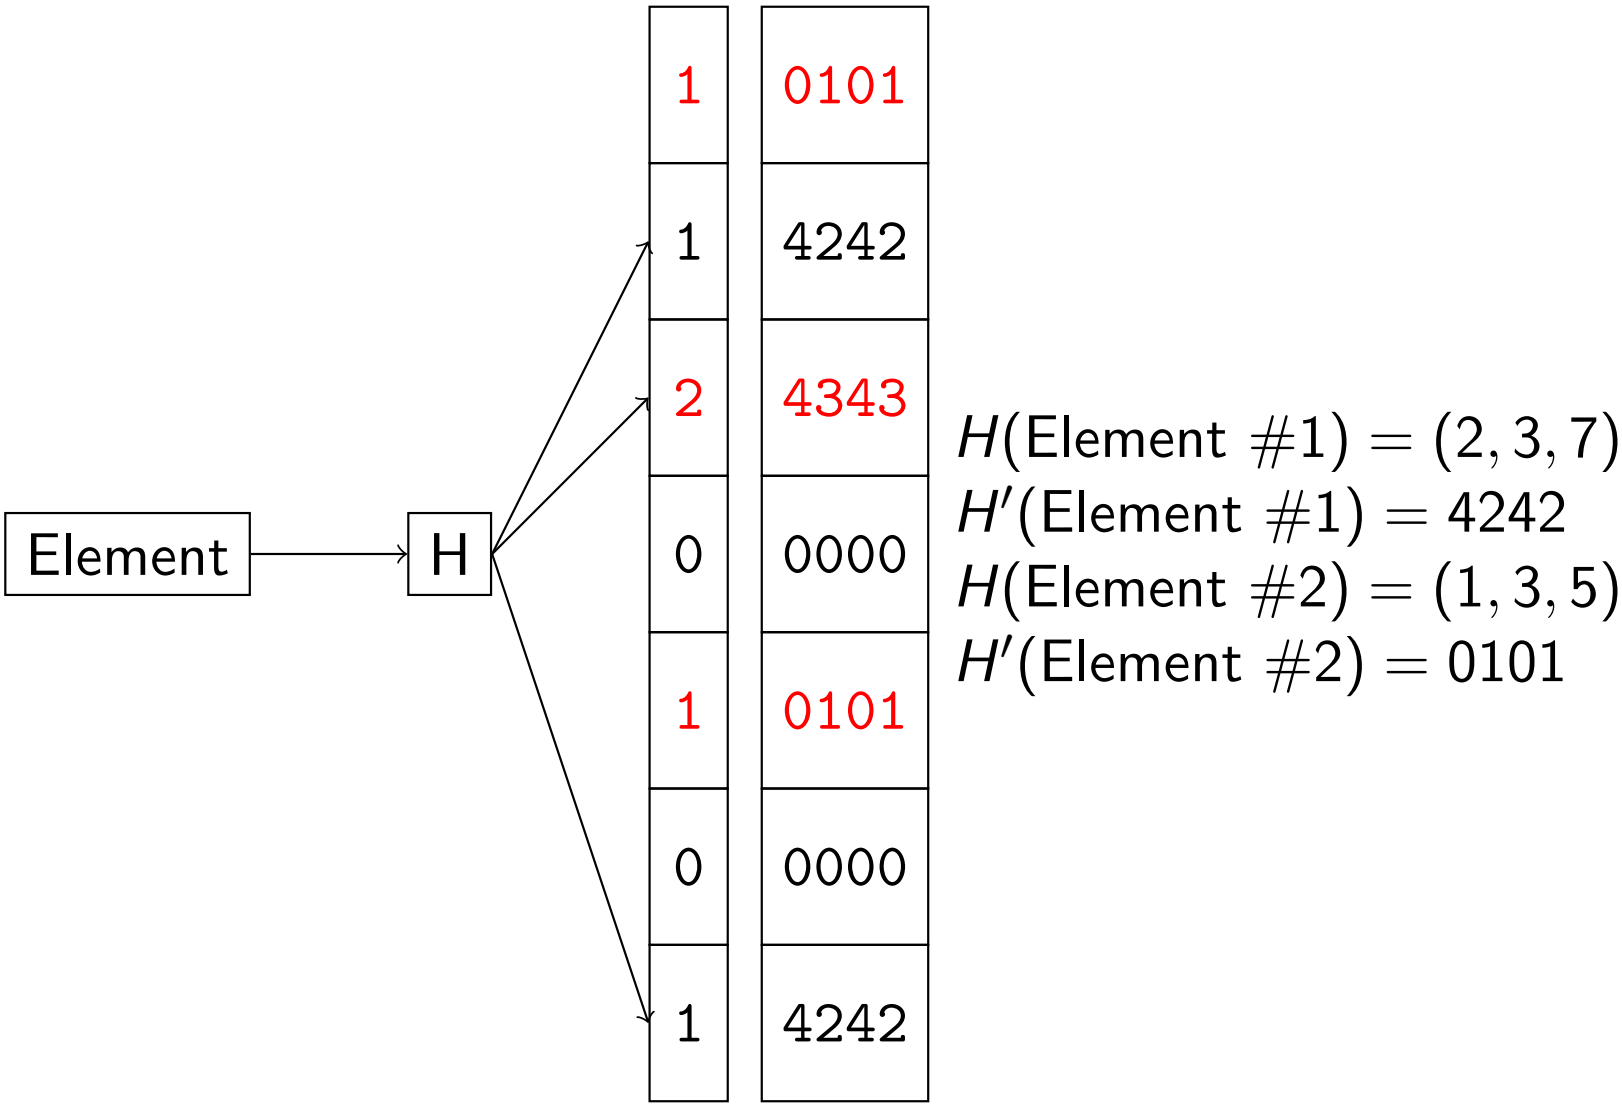

$H(\text{Element } \#1) = (2, 3, 7)$

$H'(\text{Element } \#1) = 4242$

$H(\text{Element } \#2) = (1, 3, 5)$

$H'(\text{Element } \#2) = 0101$

IBF: Extract

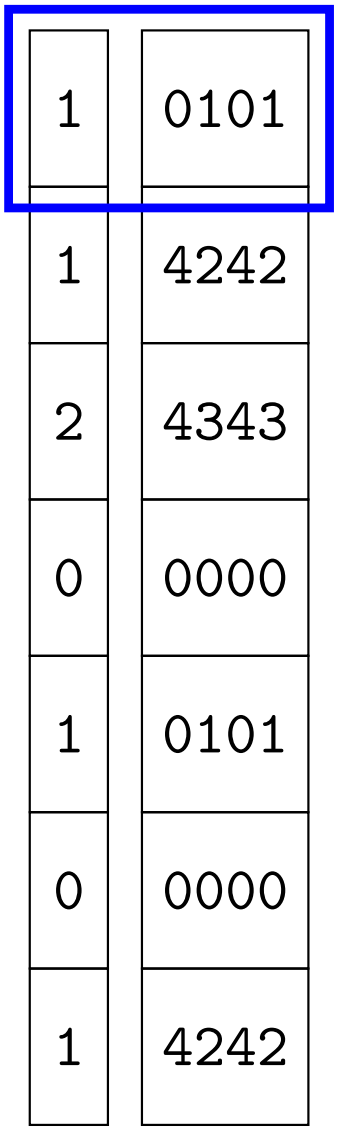

pure bucket

- Pure bucket $\Rightarrow$ extractable element hash
- Extraction $\Rightarrow$ more pure buckets (hopefully/probably)
- Less elements $\Rightarrow$ more chance for pure buckets

Symmetric Difference

We can directly compute the symmetric difference without extraction.

- Subtract counts
- XOR hashes

The Set Union Protocol

1. Create IBFs
2. Compute SymDiff
3. Extract element hashes

- Amount of communication and computation only depends on δ, not $|A| + |B|$:)
- How do we choose the initial size of the IBF?
- $\Rightarrow$ Do difference estimation first!

Difference Estimation

- We need an estimator that is accurate for small differences
- Idea: re-use IBFs for difference estimation:

1. Alice and Bob create fixed number of constant-size IBFs by sampling their set. The collection of IBFs is called a Strata Estimator (SE).
   - Stratum 0 contains $1/2$ of all elements
   - Stratum 1 contains $1/4$ of all elements
   - Stratum n contains $1/(2^n)$ all elements
2. Alice receives Bob's strata estimator
3. Alice computes $SE_{diff} = SymDiff(SE_{\text{Alice}}, SE_{\text{Bob}})$
   - by pair-wise $SymDiff$ of all IBFs in the SE
4. Alice estimates the size of SE_{diff}.

Strata Estimator

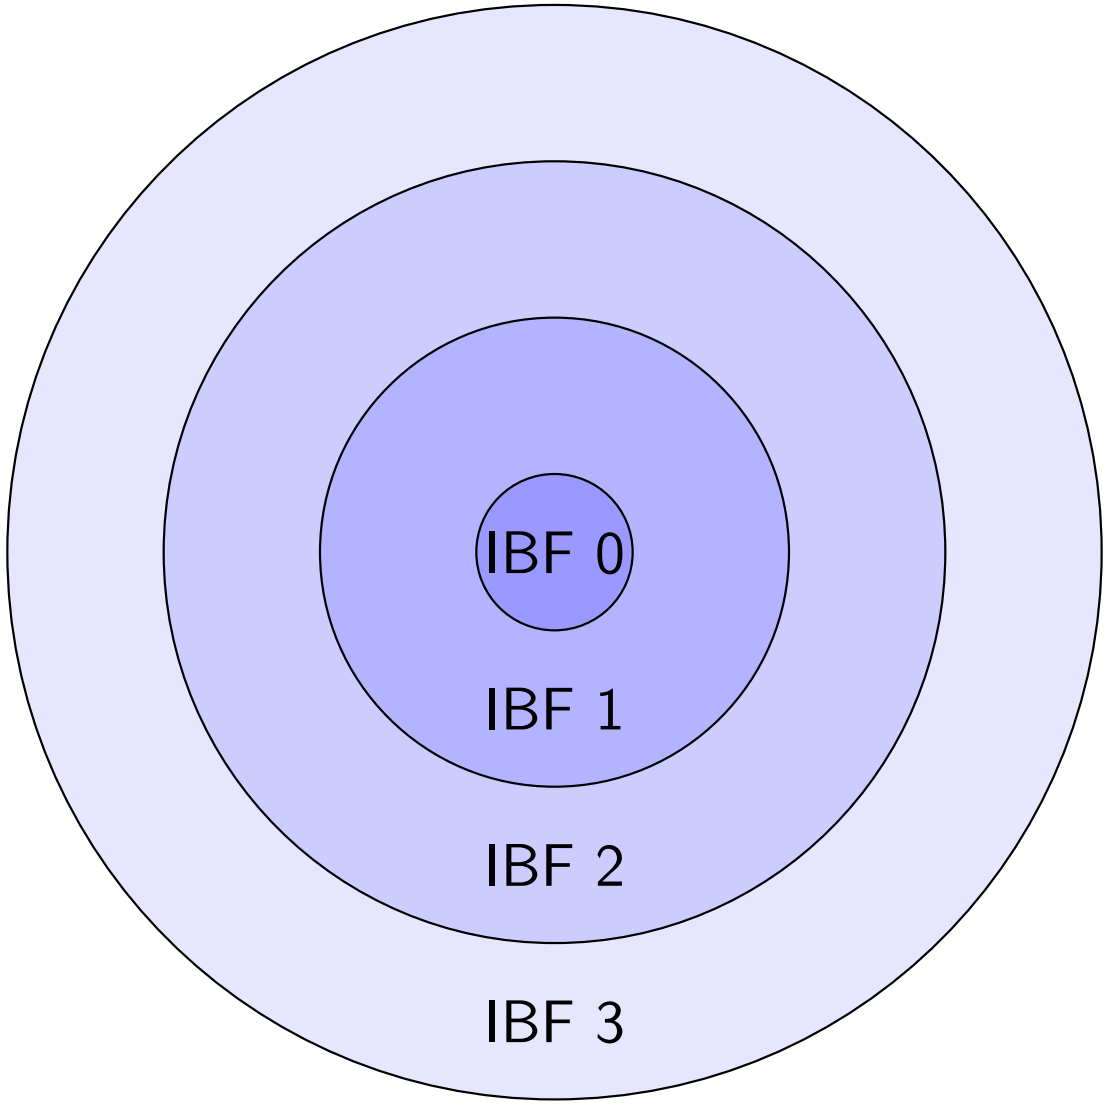

Strata Estimator

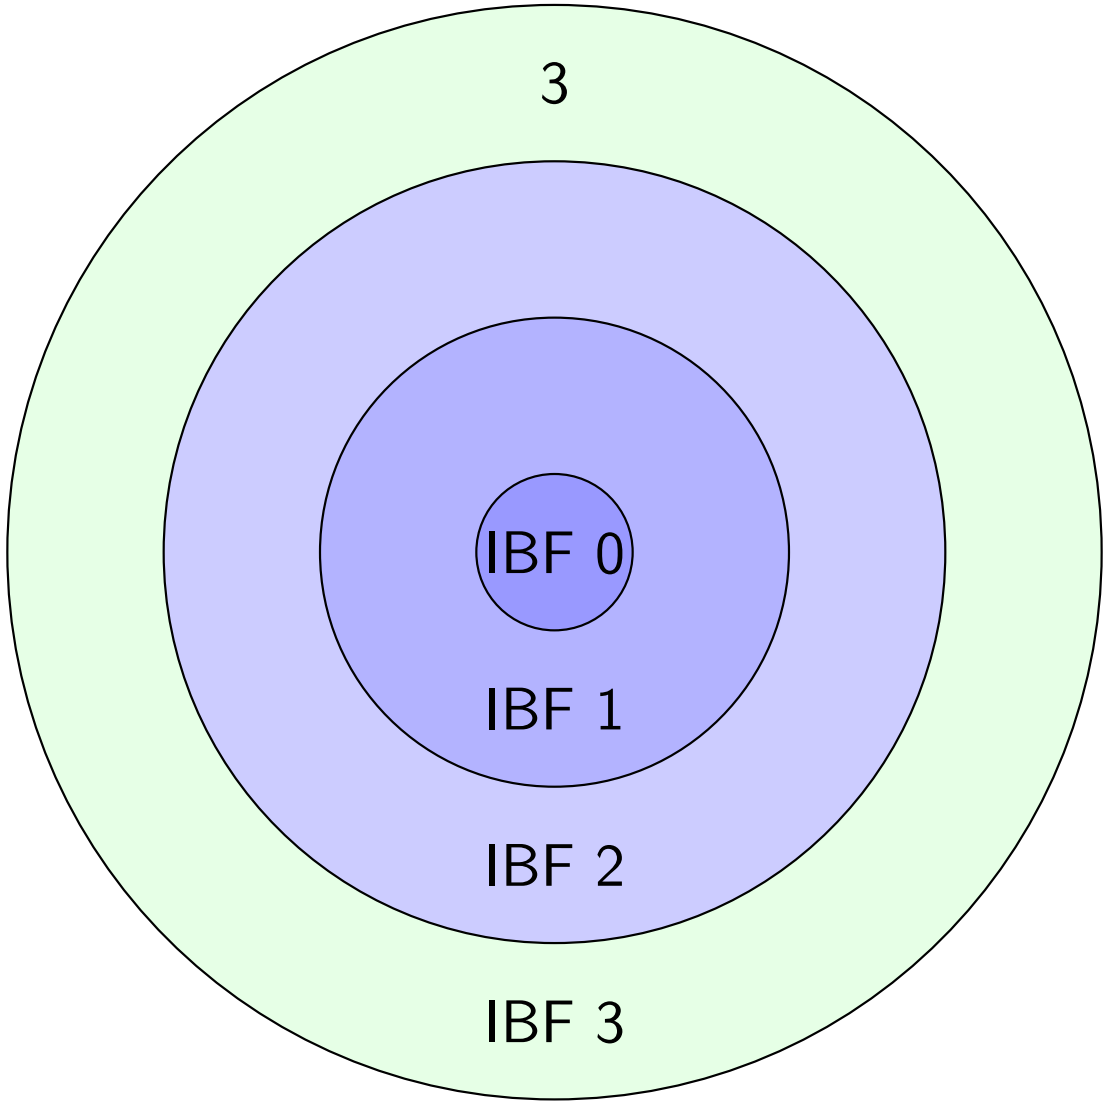

Strata Estimator

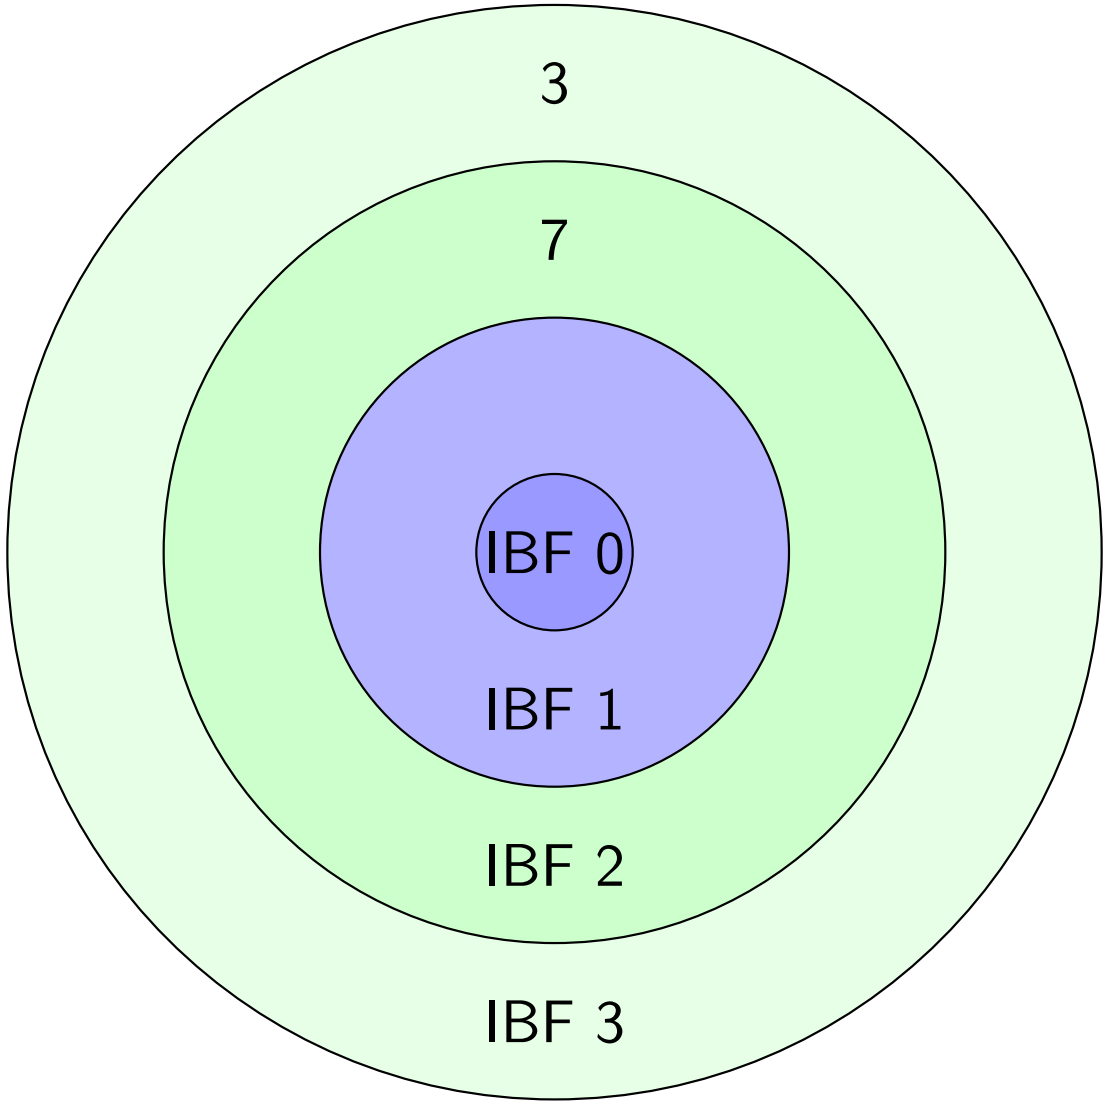

Strata Estimator

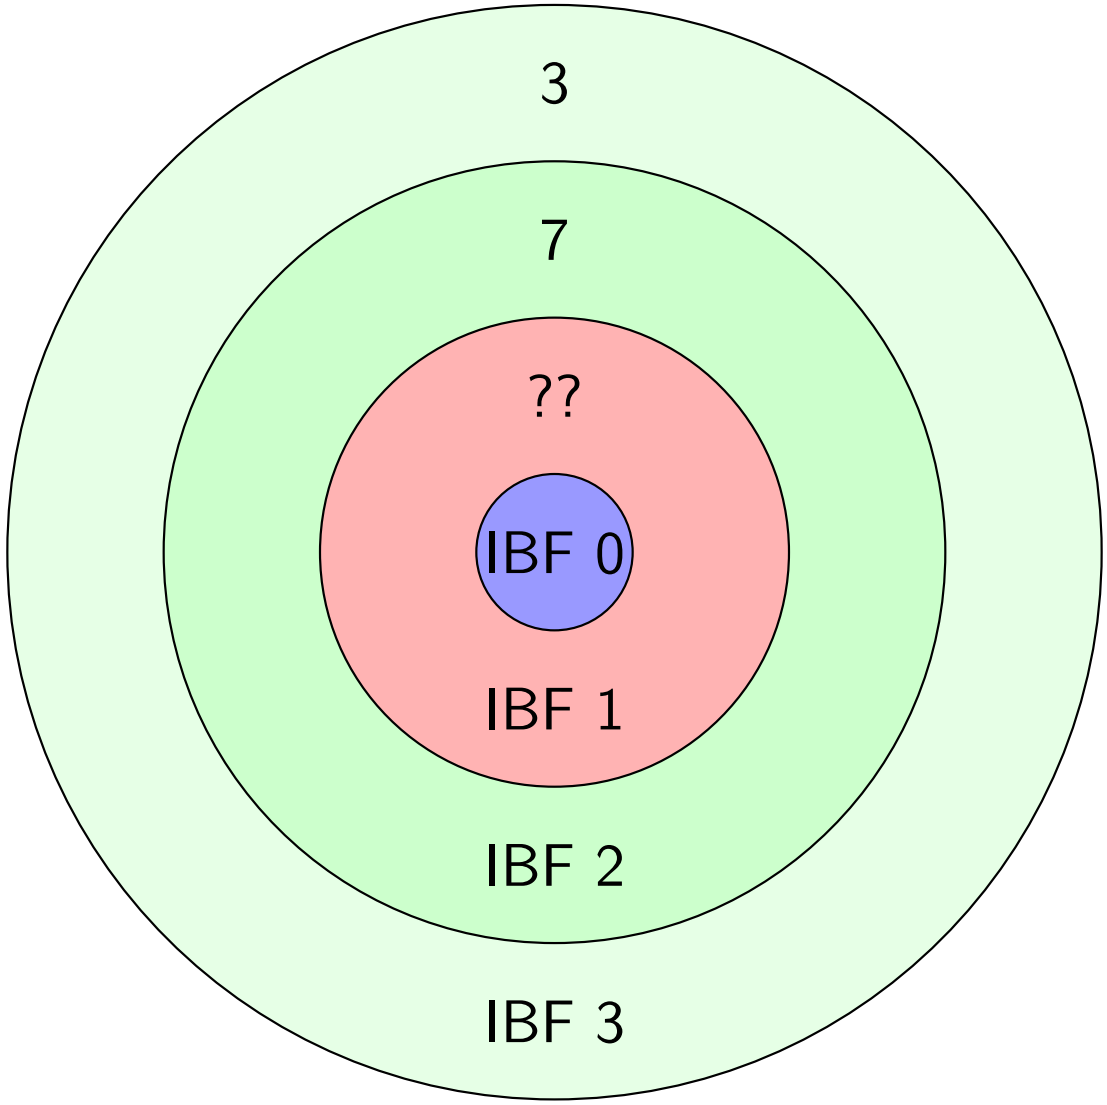

Estimation

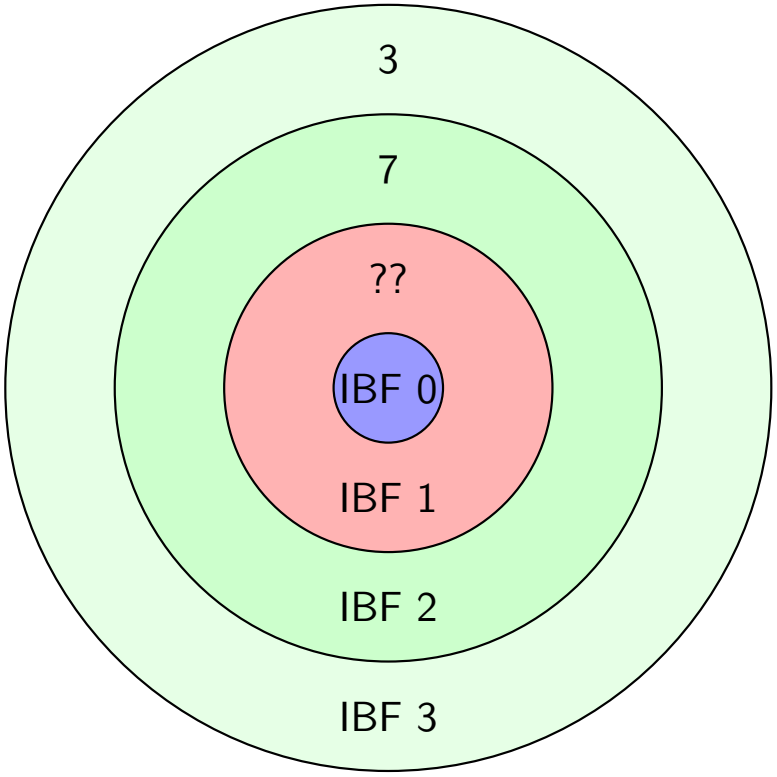

Estimate as $(3 + 7) \cdot 2^1$.
(Number of extracted hashes scaled by 2^{r-1} for r failed rounds of strata decoding.)

The Complete Protocol

1. Alice sends SE_{Alice} to Bob
2. Bob estimates the set difference δ
3. Bob computes IBF_{Bob} with size δ and sends it to Alice
4. Alice computes IBF_{Alice}
5. Alice computes $IBF_{diff} = SymDiff(IBF_{\text{Alice}}, IBF_{\text{Bob}})$
6. Alice extracts element hashes from IBF_{diff}.
   - $b = left \Rightarrow$ Send element to to Bob
   - $b = right \Rightarrow$ Send element request to to Bob
   - $b = fail \Rightarrow$ Send larger IBF (double the size) to Bob, go to (3.) with switched roles
   - $b = done \Rightarrow$ We're done ...

Break

Security Goals for Name Systems

- Query origin anonymity
- Data origin authentication and integrity protection
- Zone confidentiality
- Query and response privacy
- Censorship resistance
- Traffic amplification resistance
- Availability

Reminder: DNSSEC

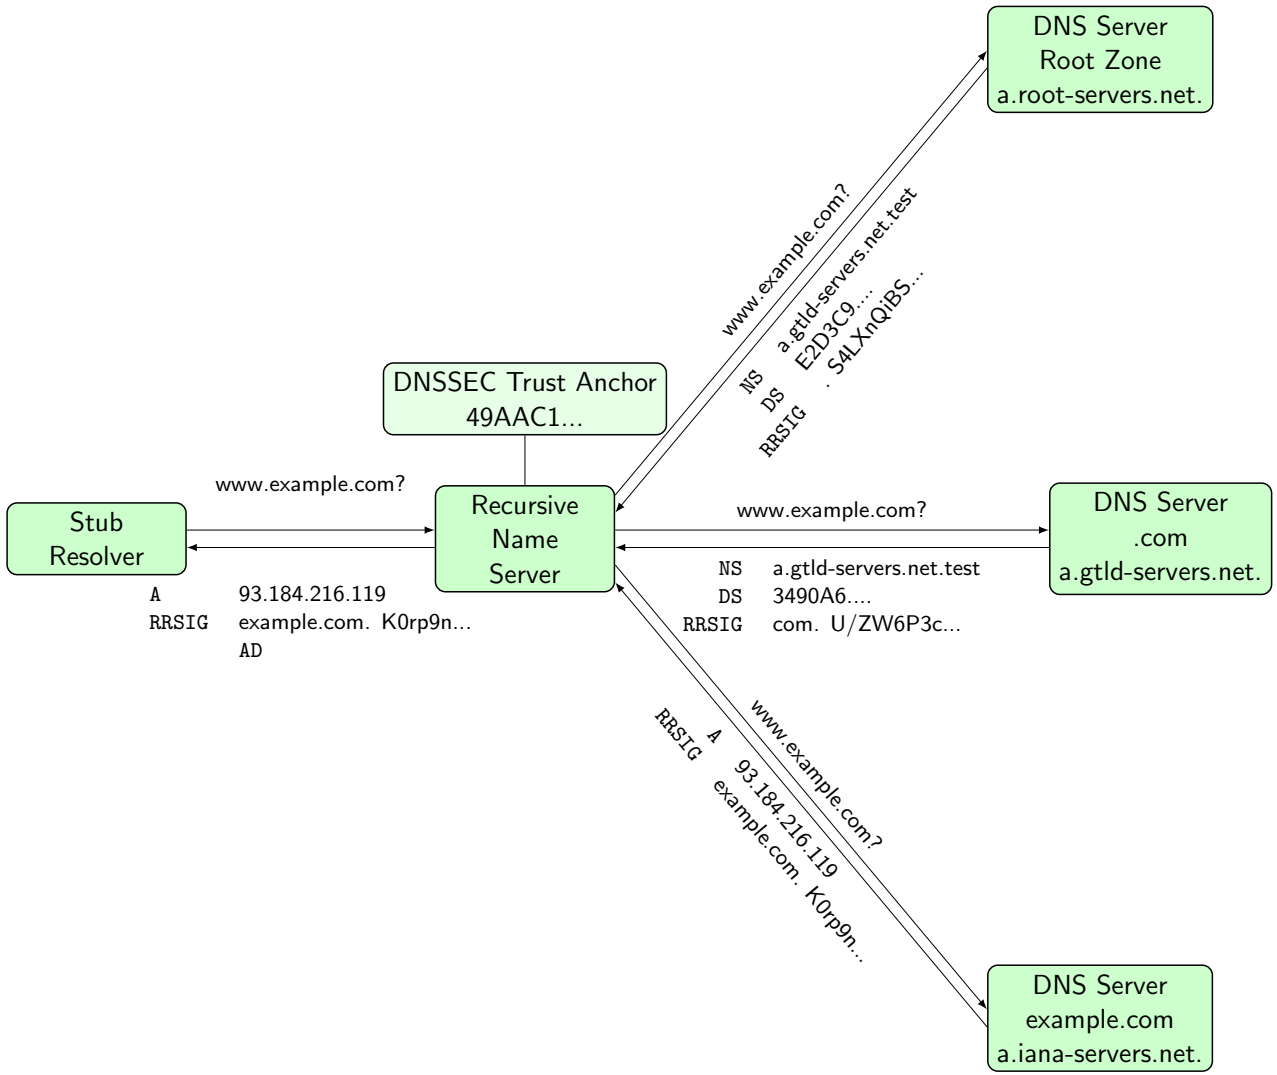

Exemplary Attacks: MORECOWBELL

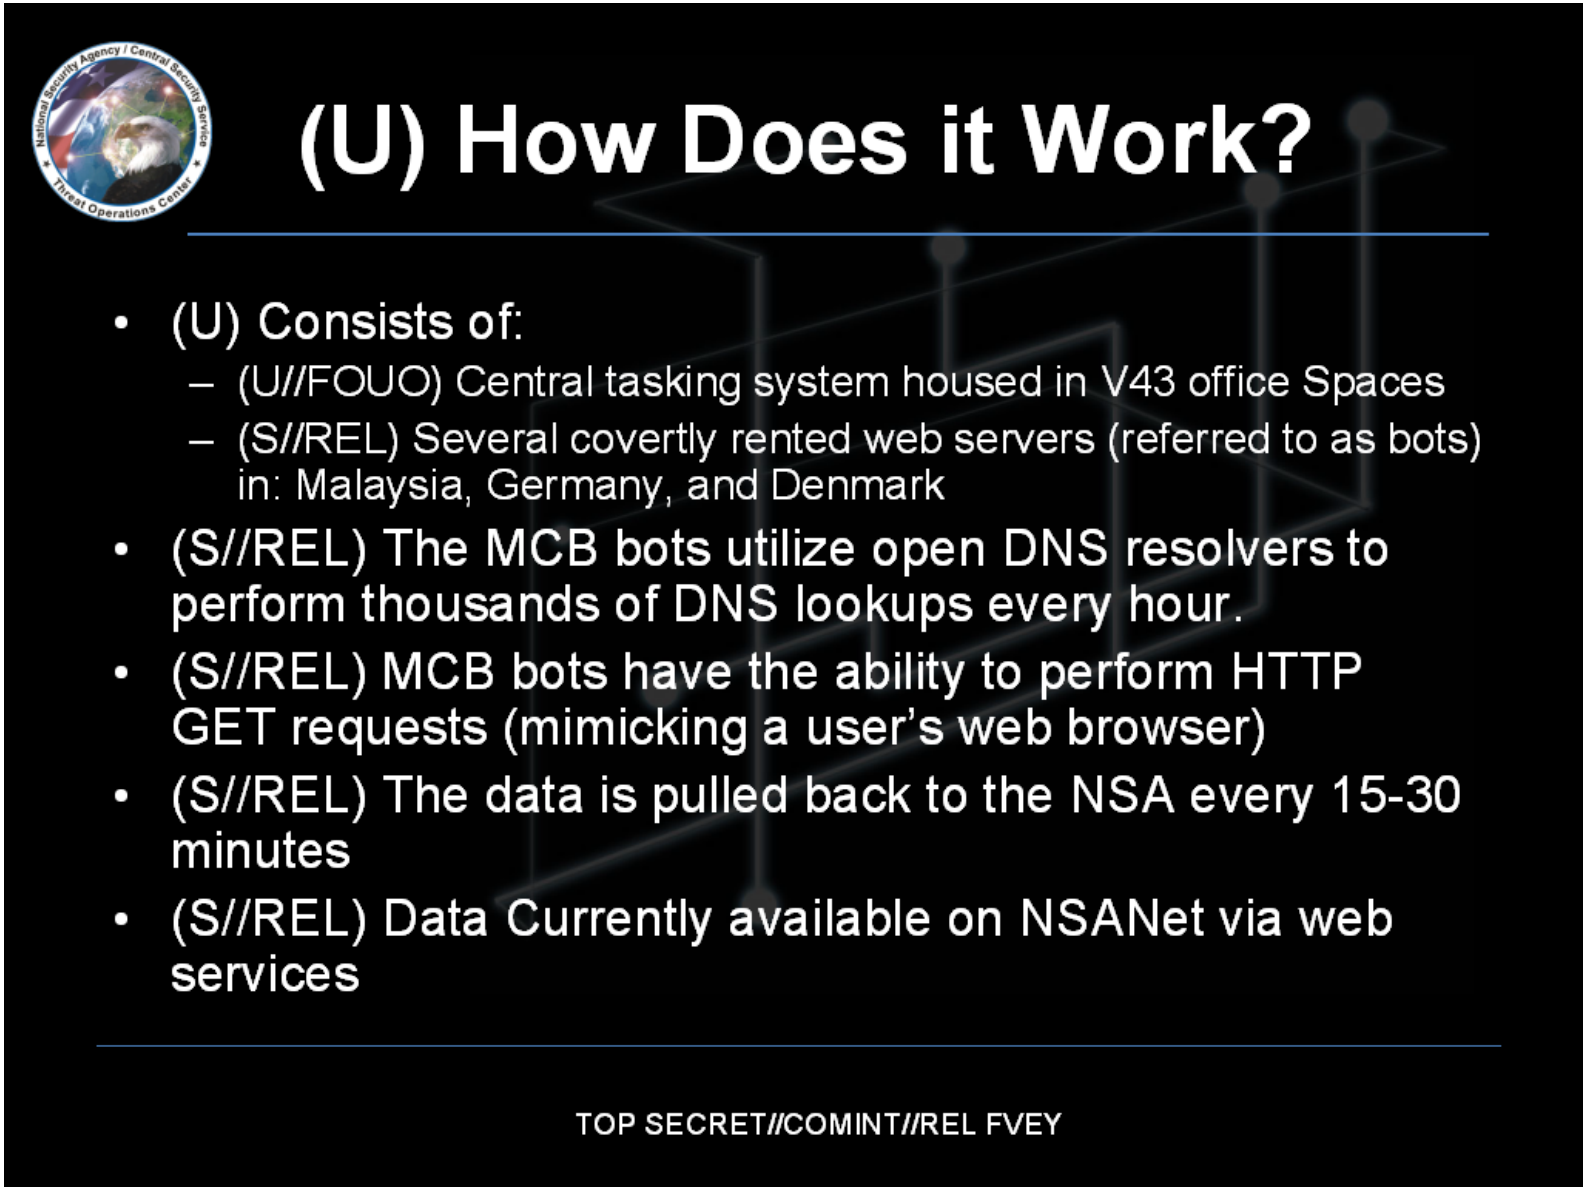

Exemplary Attacks: QUANTUMDNS

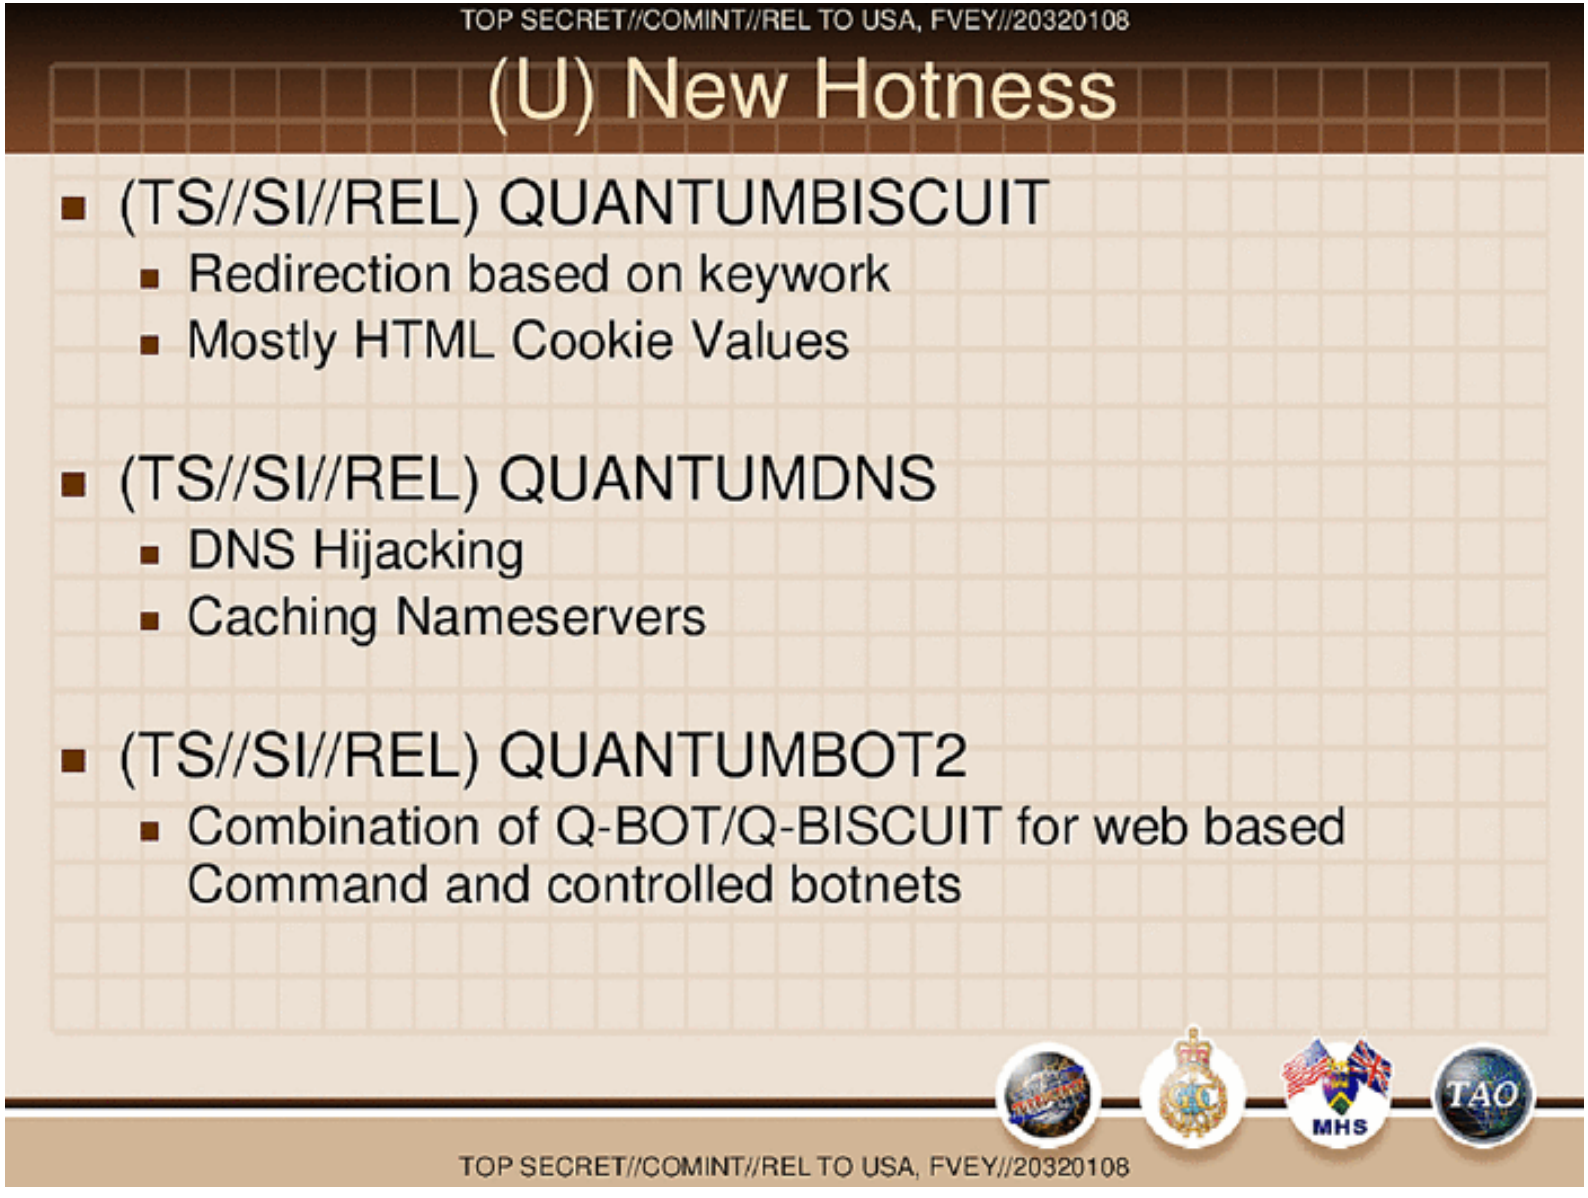

Zooko's Triangle

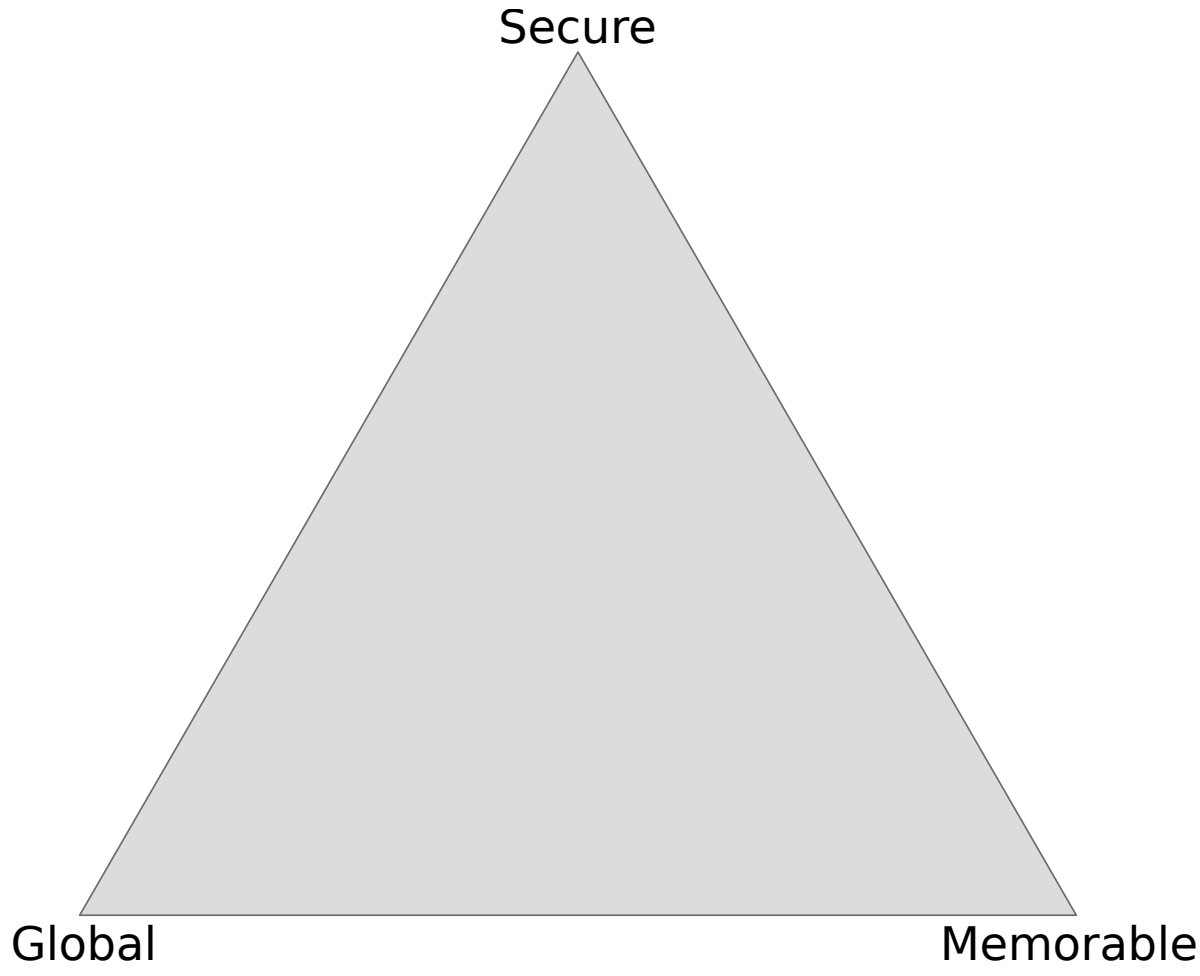

A name system can only fulfill **two**!

Zooko's Triangle

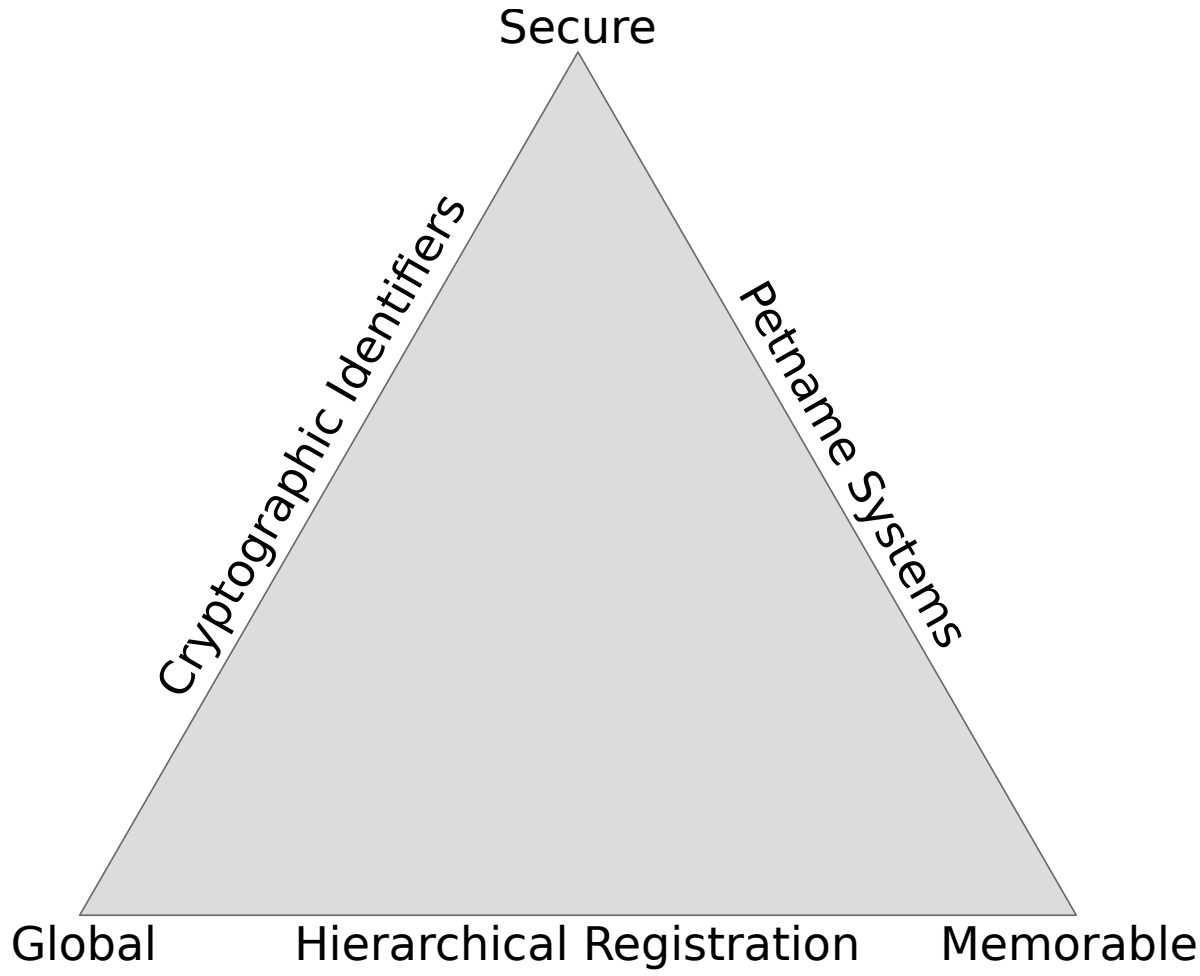

DNS, ".onion" IDs and `/etc/hosts/` are representative designs.

Zooko's Triangle

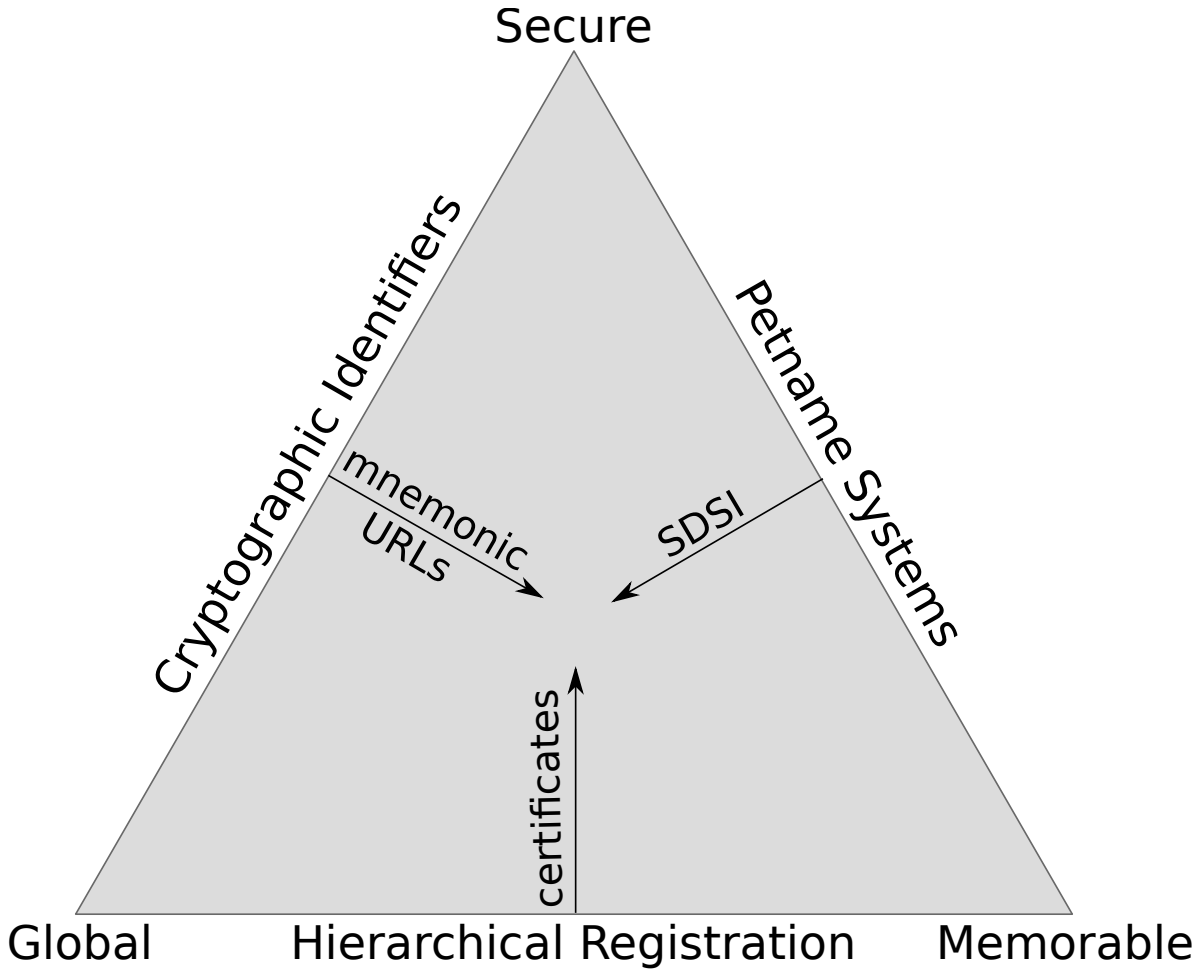

Query Name Minimization

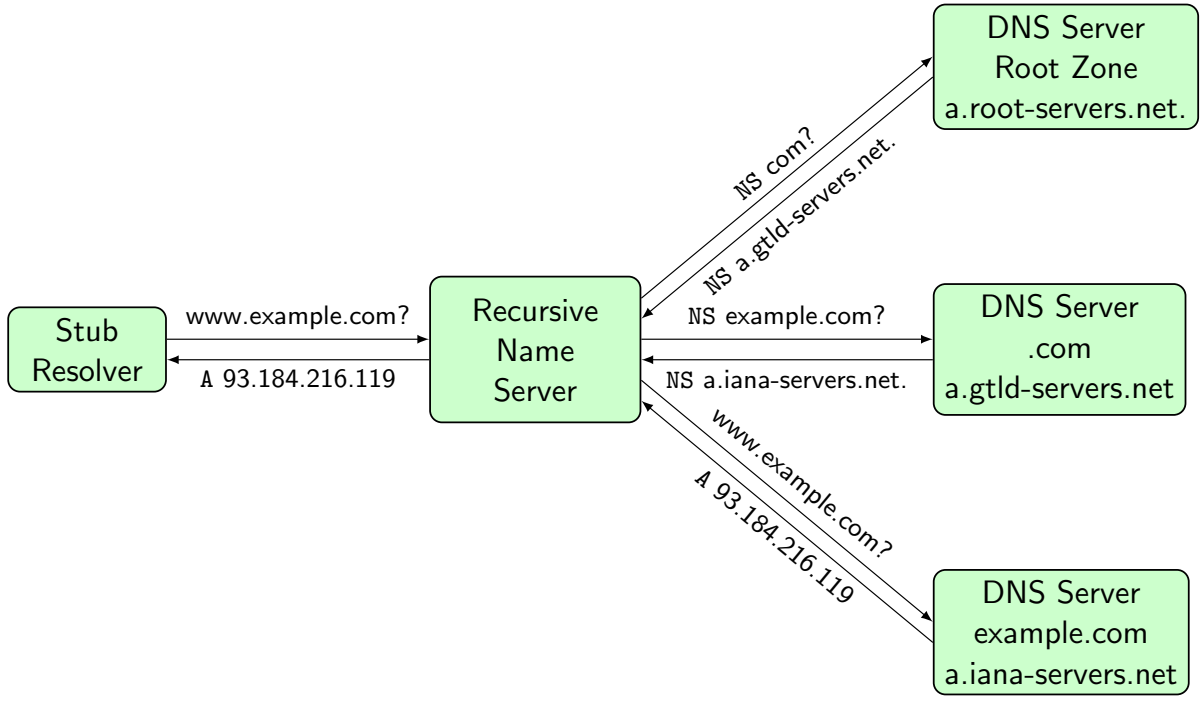

DNS over TLS

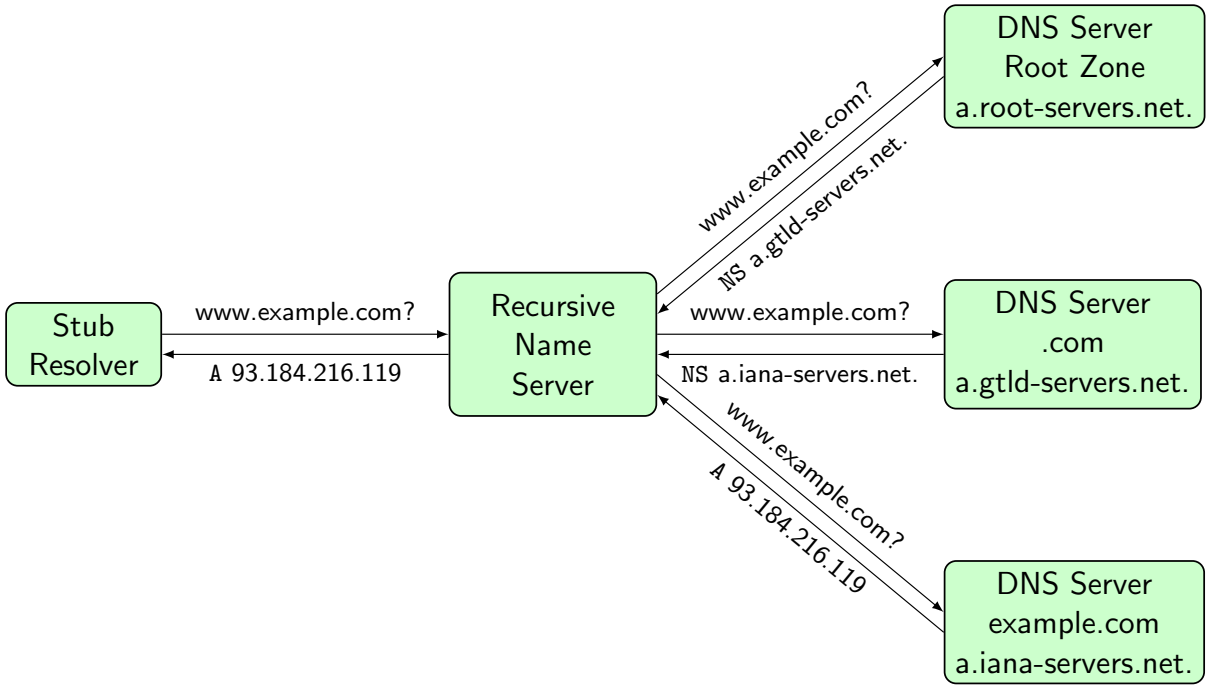
Stub Resolver
www.example.com?
A 93.184.216.119
Recursive Name Server
www.example.com?
NS a.gtld-servers.net.
DNS Server Root Zone a.root-servers.net.
www.example.com?
NS a.iana-servers.net.
DNS Server .com a.gtld-servers.net.
www.example.com?
A 93.184.216.119
DNS Server example.com a.iana-servers.net.

The Textbook Version of the Internet

Layering, $\approx$ 1990

	HTTPS
DNS	TLS
UDP	TCP
IPv4	
Ethernet	
Phys. Layer	

The Textbook Version of the Internet

Layering, ≈ 1990

	HTTPS
DNS	TLS
UDP	TCP
IPv4	
Ethernet	
Phys. Layer	

"Layering", ≈ 2020

HTTPS	libmicrohttpd
TLS-with-DANE	libgnutls
DNS-over-TLS	libunbound
TLS*	libnss
TCP	Linux
IPv6	Linux
Ethernet	
Phys. Layer	

* = castrated version without RFC 6125 or RFC 6394, possibly NULL cipher, see TLS profiles draft.

DNSCurve

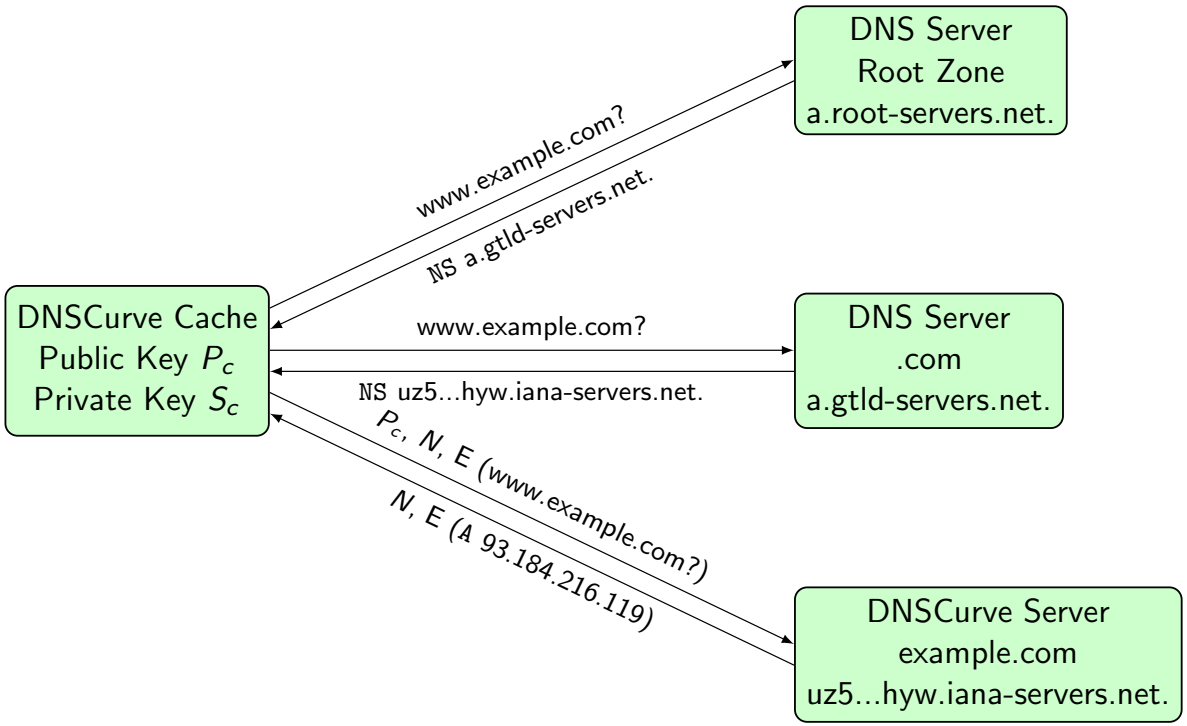
DNS Server
Root Zone
a.root-servers.net.
www.example.com?
NS a.gtld-servers.net.
DNSCurve Cache
Public Key P_c
Private Key S_c
www.example.com?
NS uz5...hyw.iana-servers.net.
DNS Server
.com
a.gtld-servers.net.
P_c, N, E (www.example.com?)
N, E (A 93.184.216.119)
DNSCurve Server
example.com
uz5...hyw.iana-servers.net.

Namecoin

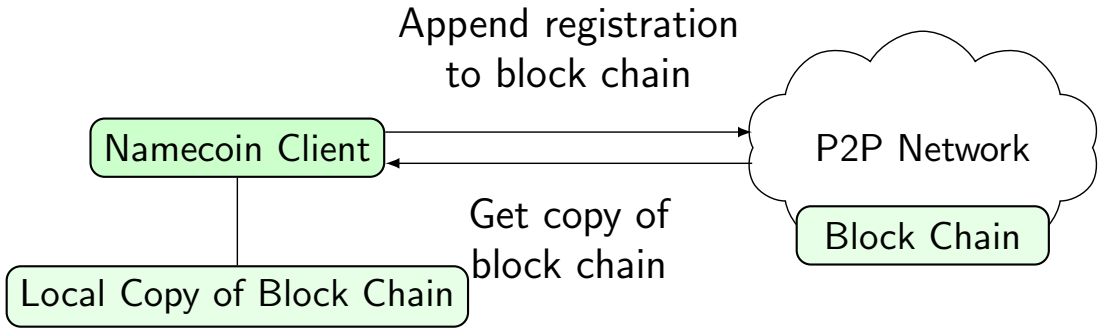
Append registration
to block chain
Namecoin Client
P2P Network
Block Chain
Get copy of
block chain
Local Copy of Block Chain

RAINS

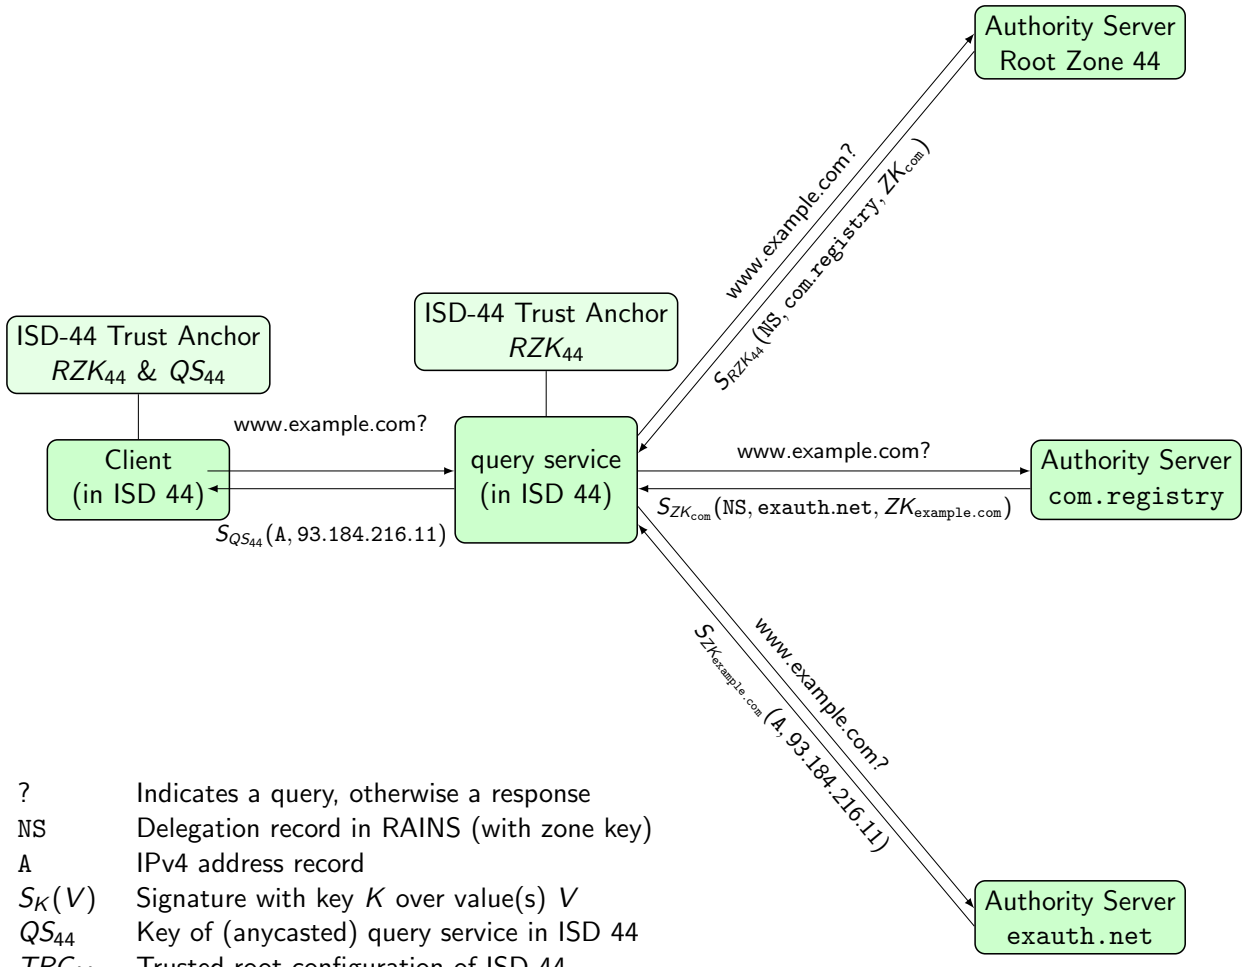

?	Indicates a query, otherwise a response
NS	Delegation record in RAINS (with zone key)
A	IPv4 address record
$S_K(V)$	Signature with key K over value(s) V
QS_{44}	Key of (anycasted) query service in ISD 44
TRC_{44}	Trusted root configuration of ISD 44
RZK_{44}	Root zone key of ISD 44
ZK_{name}	Zone key of authority for "name"

The GNU Name System (GNS)

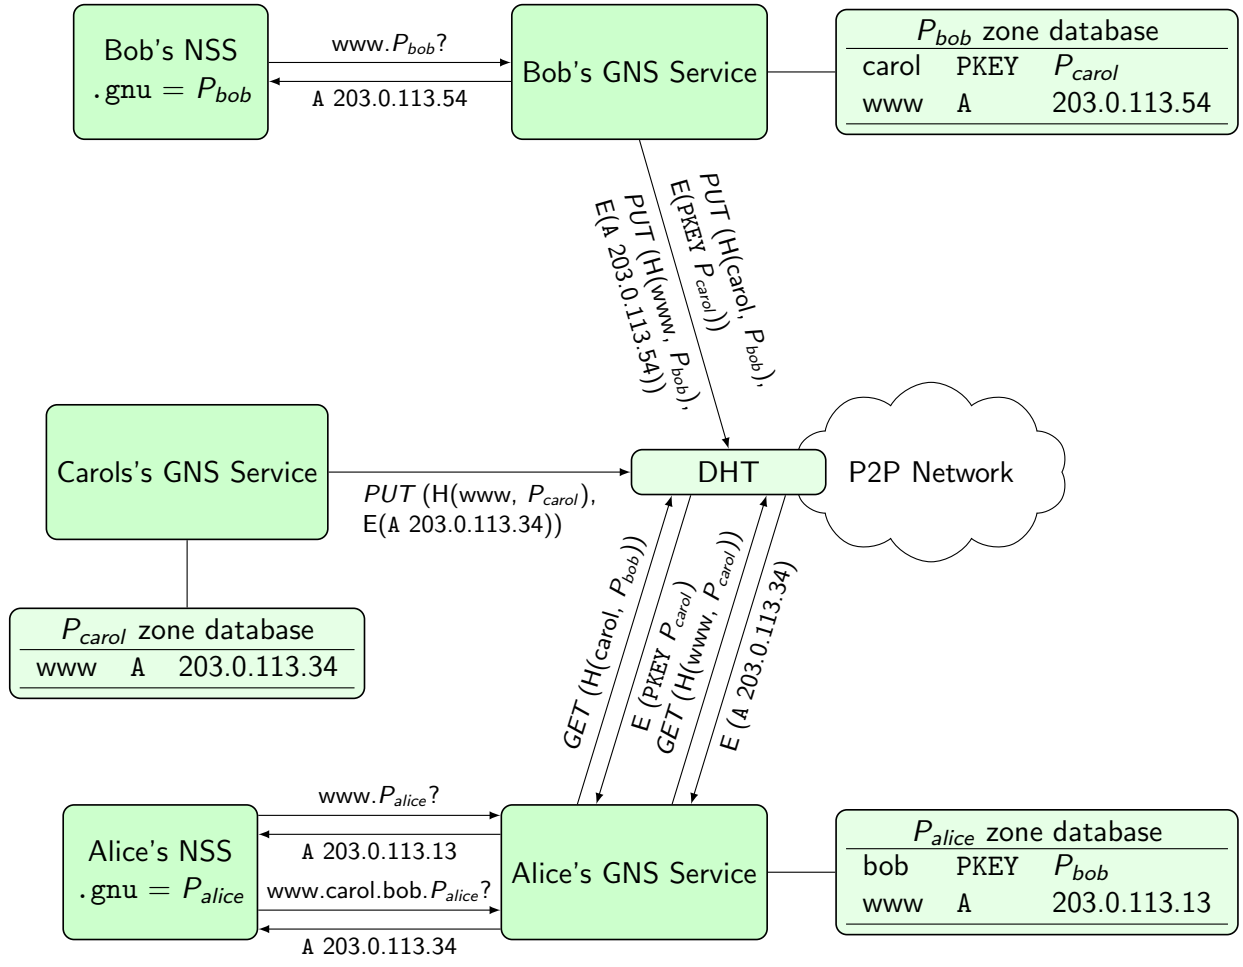

The GNU Name System[1]

Properties of GNS

- Decentralized name system with secure memorable names
- Delegation used to achieve transitivity
- Also supports globally unique, secure identifiers
- Achieves query and response privacy
- Provides alternative public key infrastructure
- Interoperable with DNS

[1] Joint work with Martin Schanzenbach and Matthias Wachs

Zone Management: like in DNS

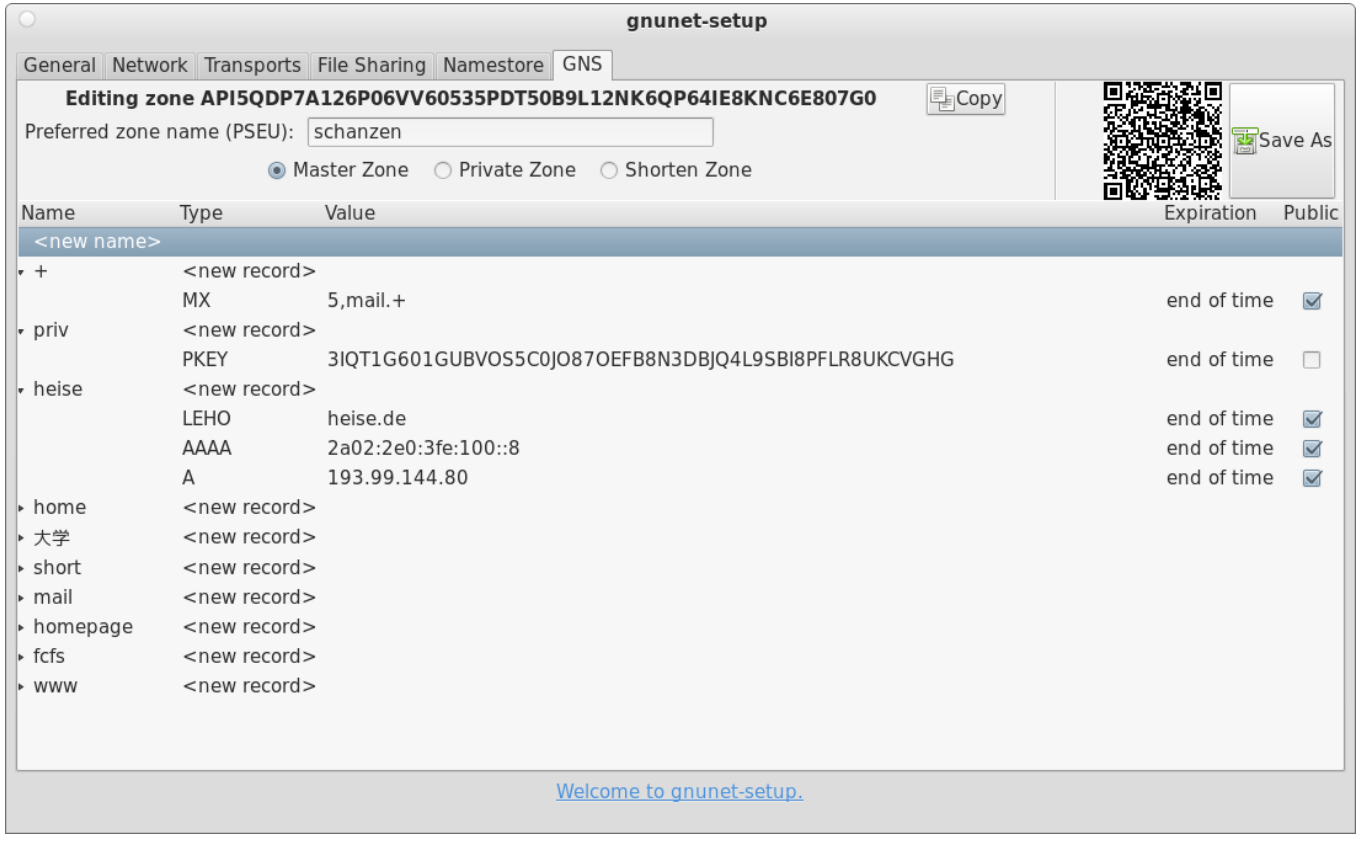

Name resolution in GNS

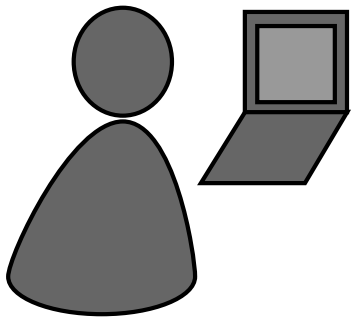

Bob

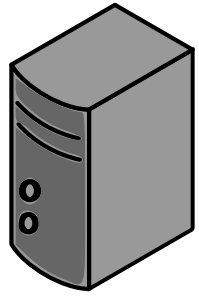

Bob's webserver

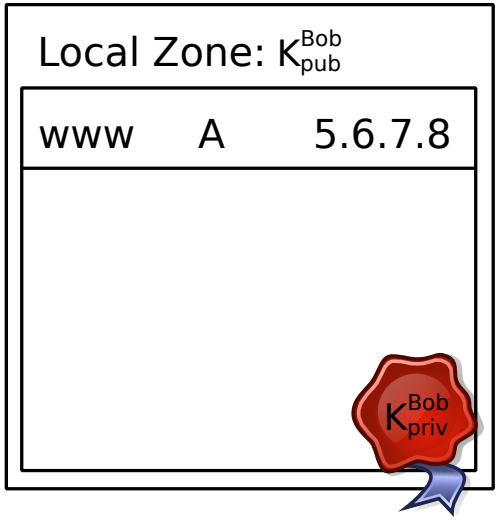

- Bob can locally reach his webserver via **www.gnu**

Secure introduction

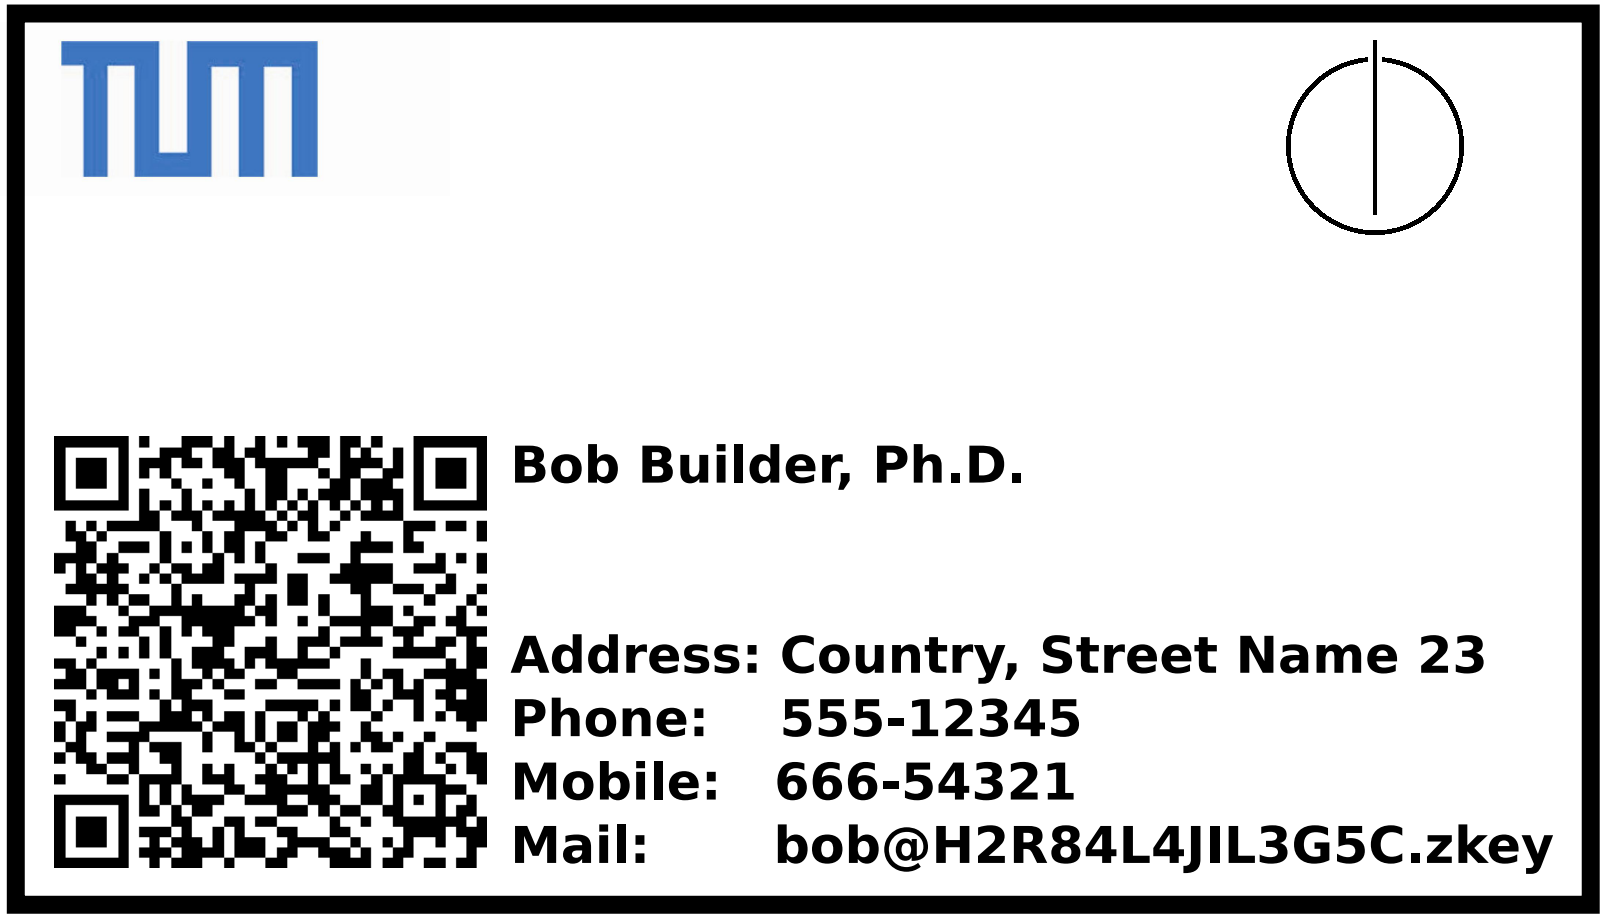

- Bob gives his public key to his **friends**, possibly via QR code

Delegation

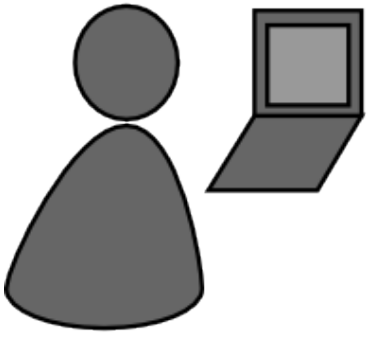

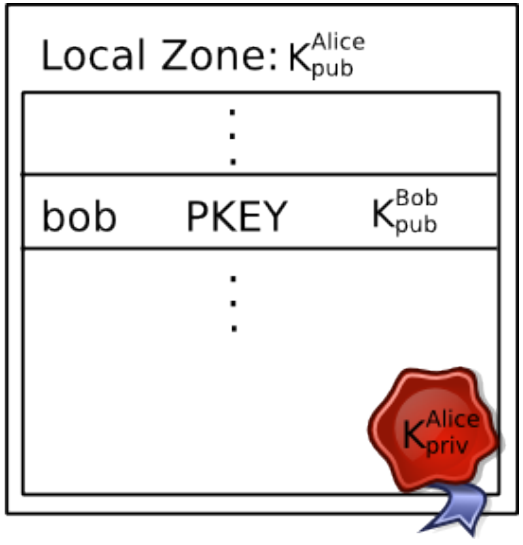

- Alice learns Bob's public key
- Alice creates delegation to zone K^{Bob}_{pub} under label **bob**
- Alice can reach Bob's webserver via **www.bob.gnu**

Name Resolution

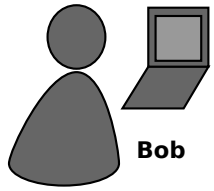

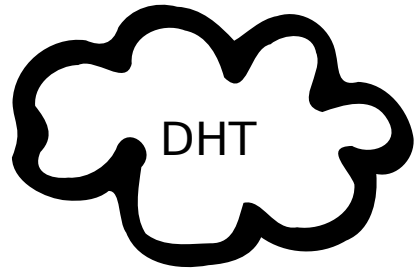

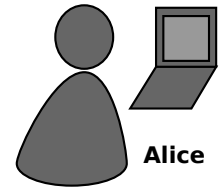

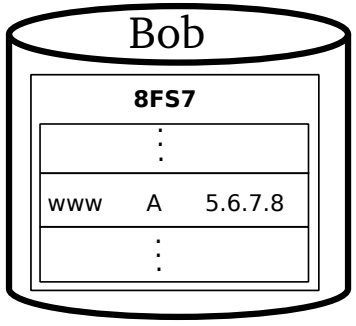

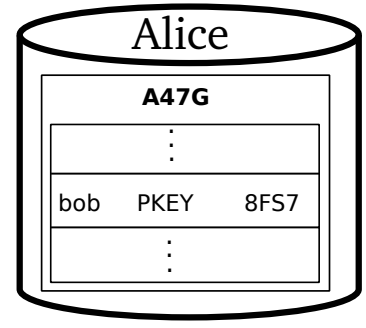

Name Resolution

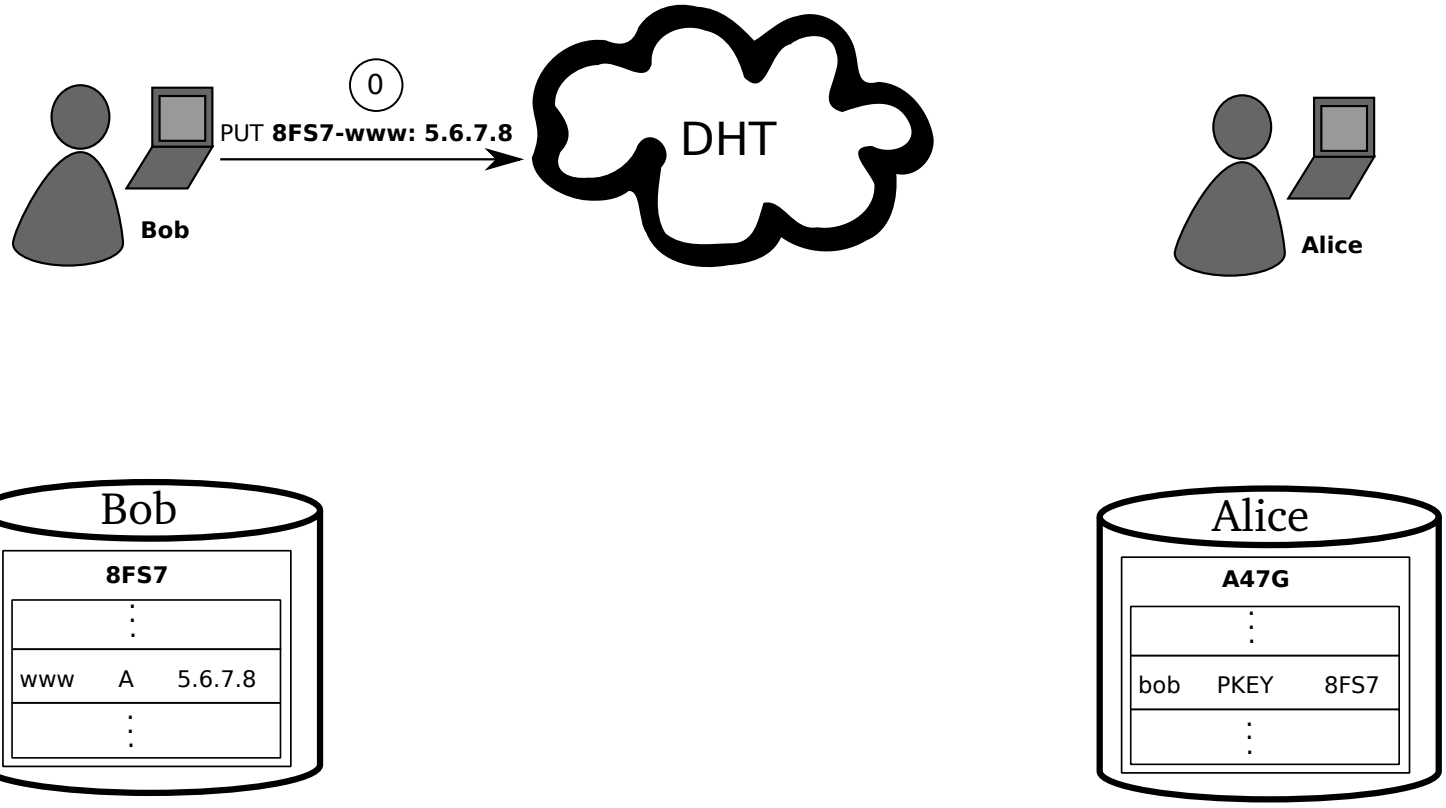

Name Resolution

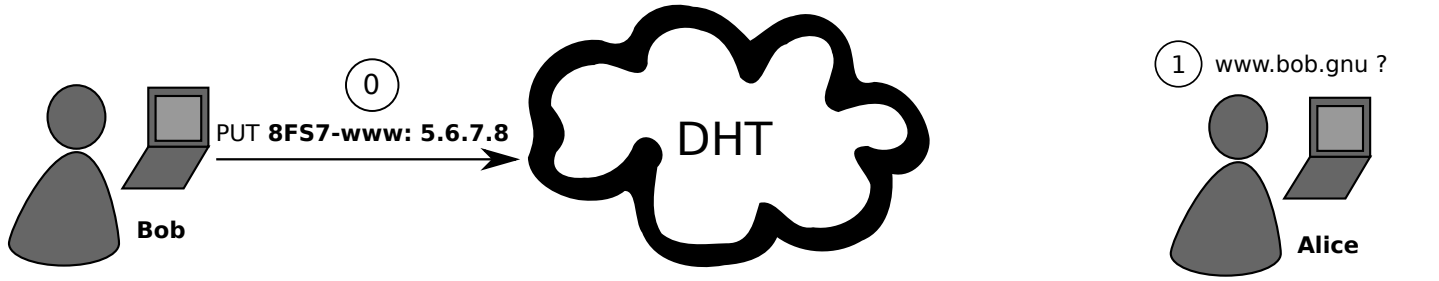

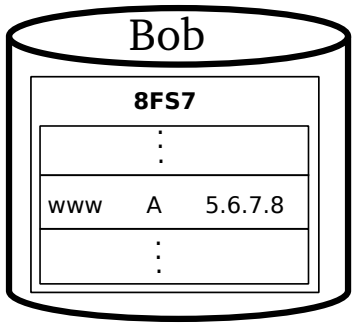

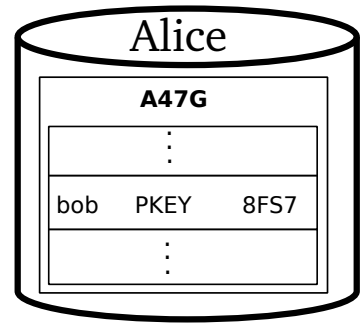

Name Resolution

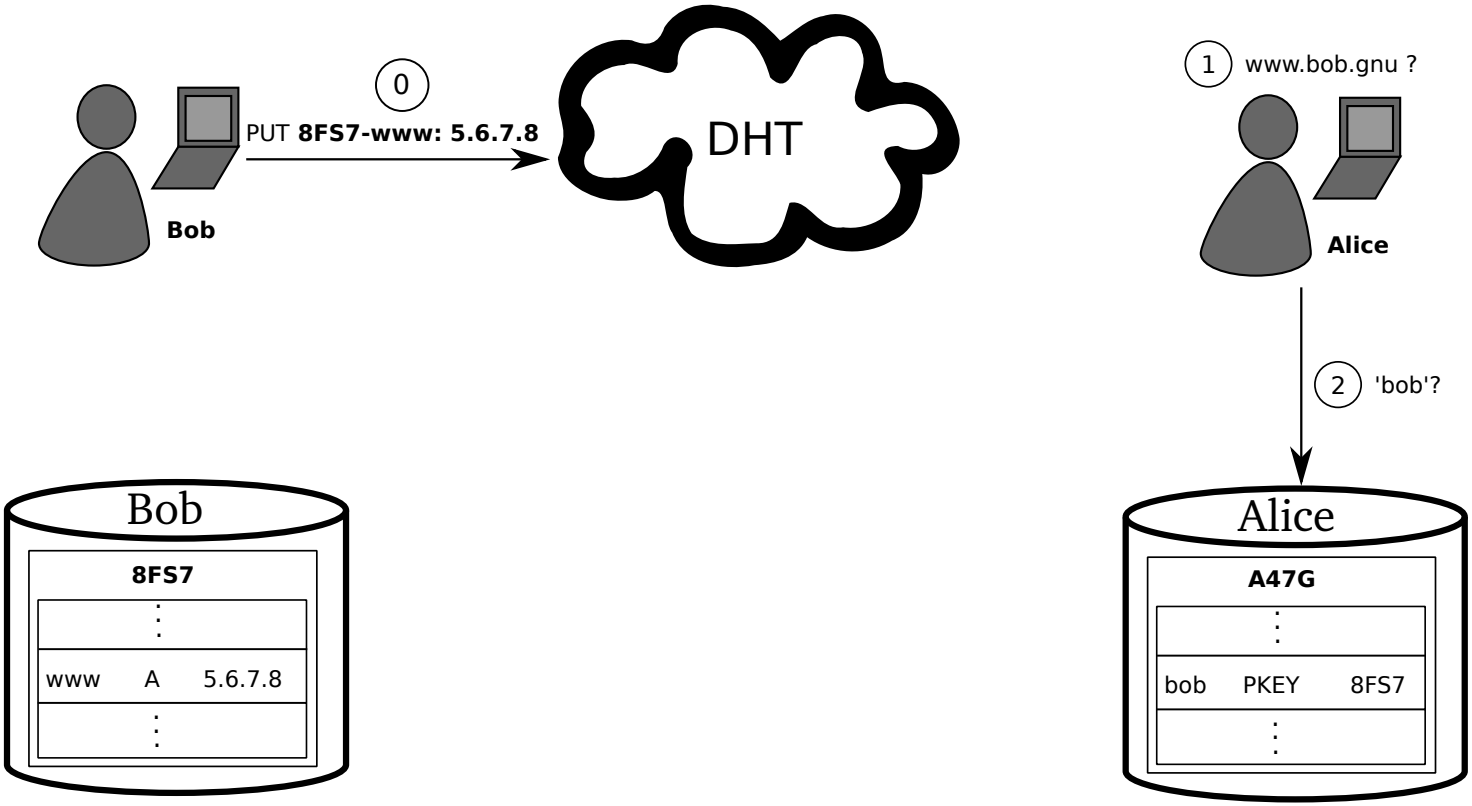

Name Resolution

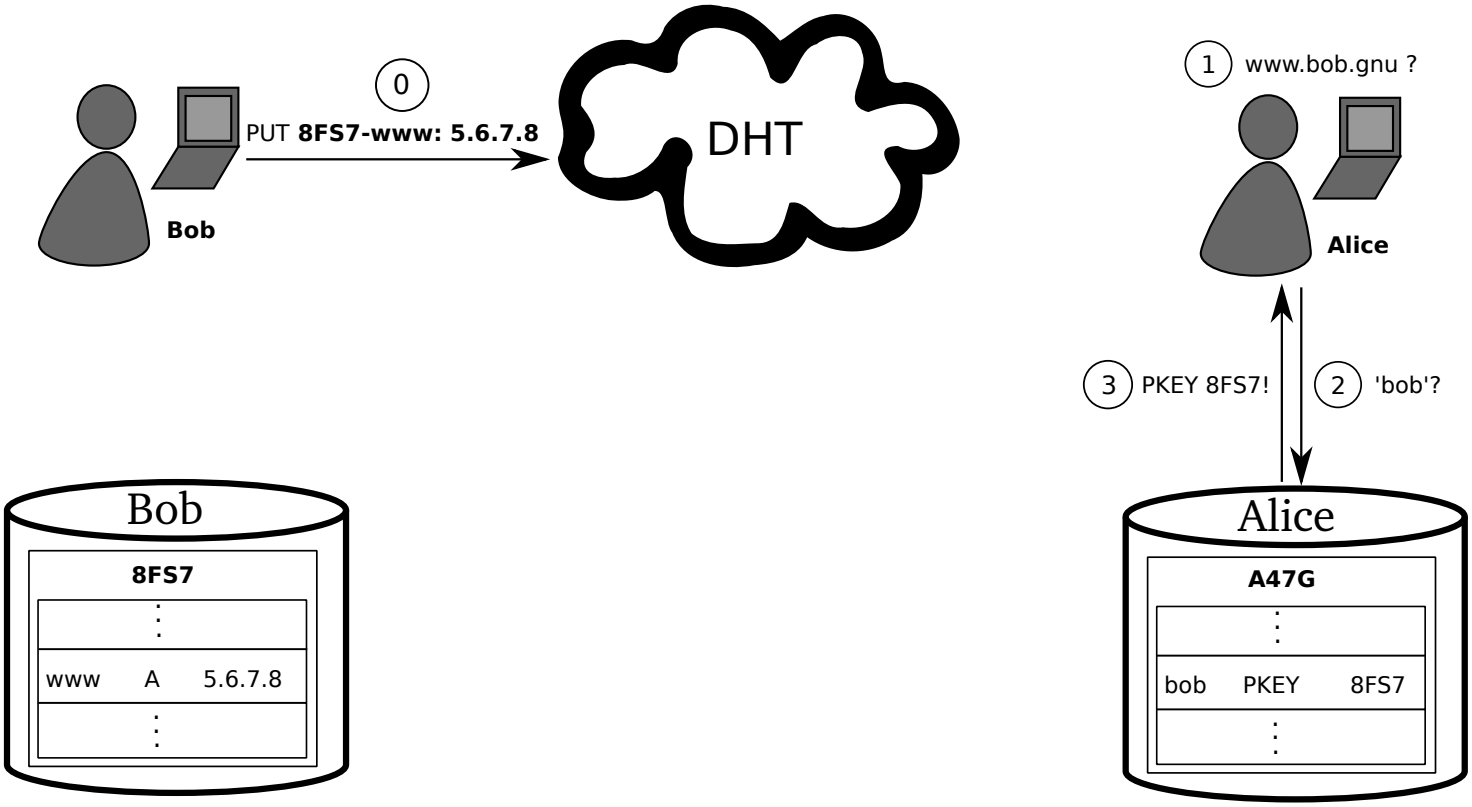

Name Resolution

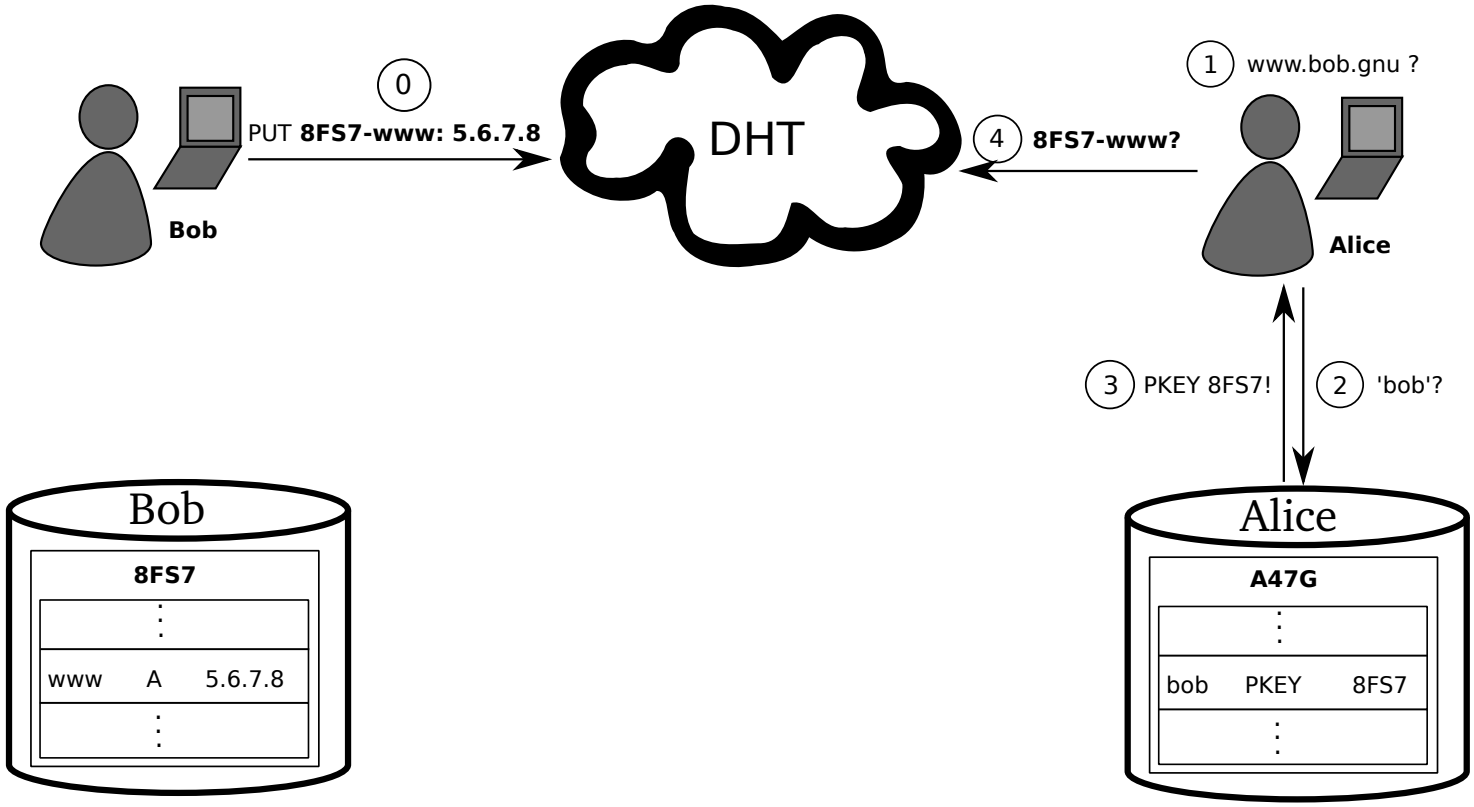

Name Resolution

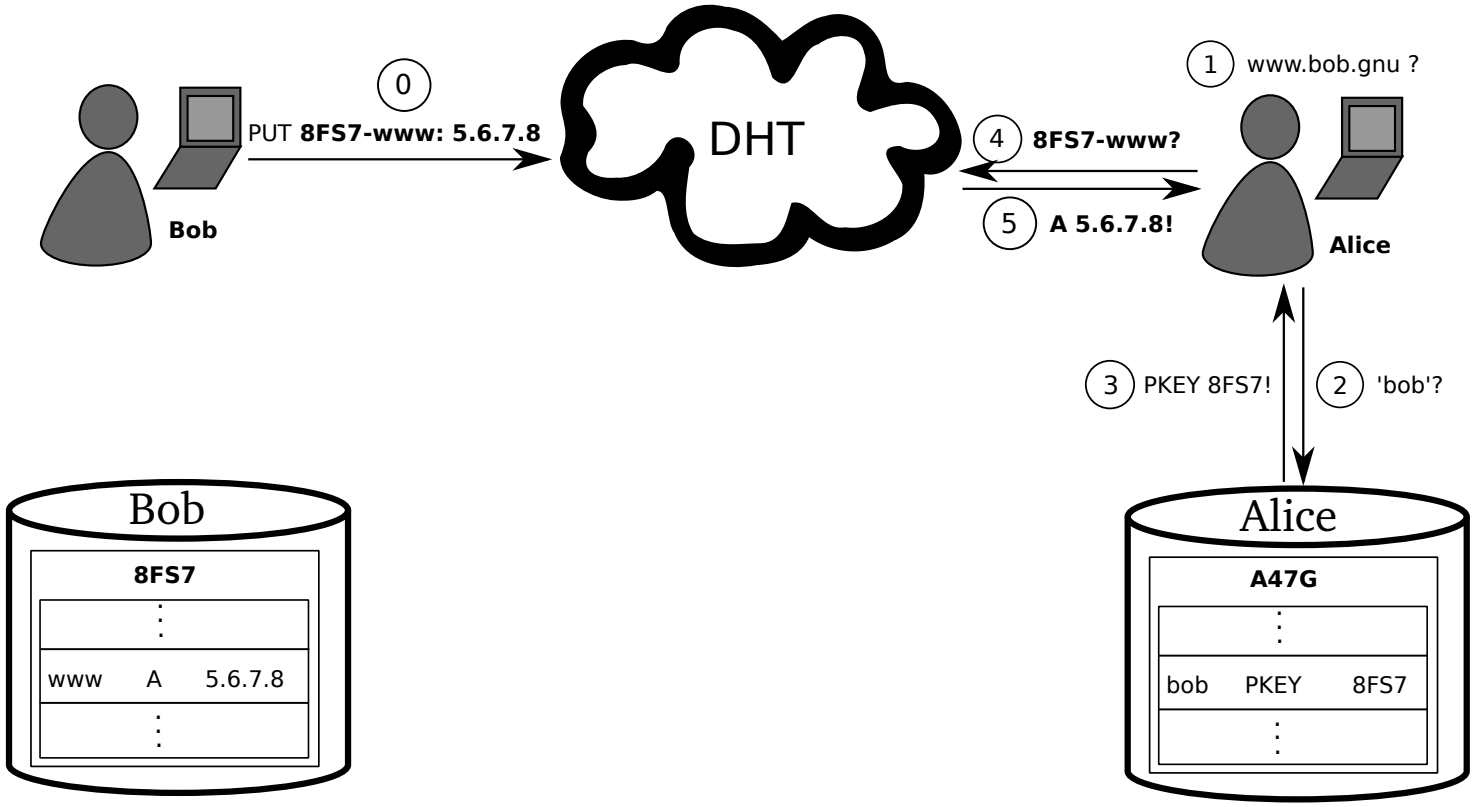

GNS as PKI (via DANE/TLSA)

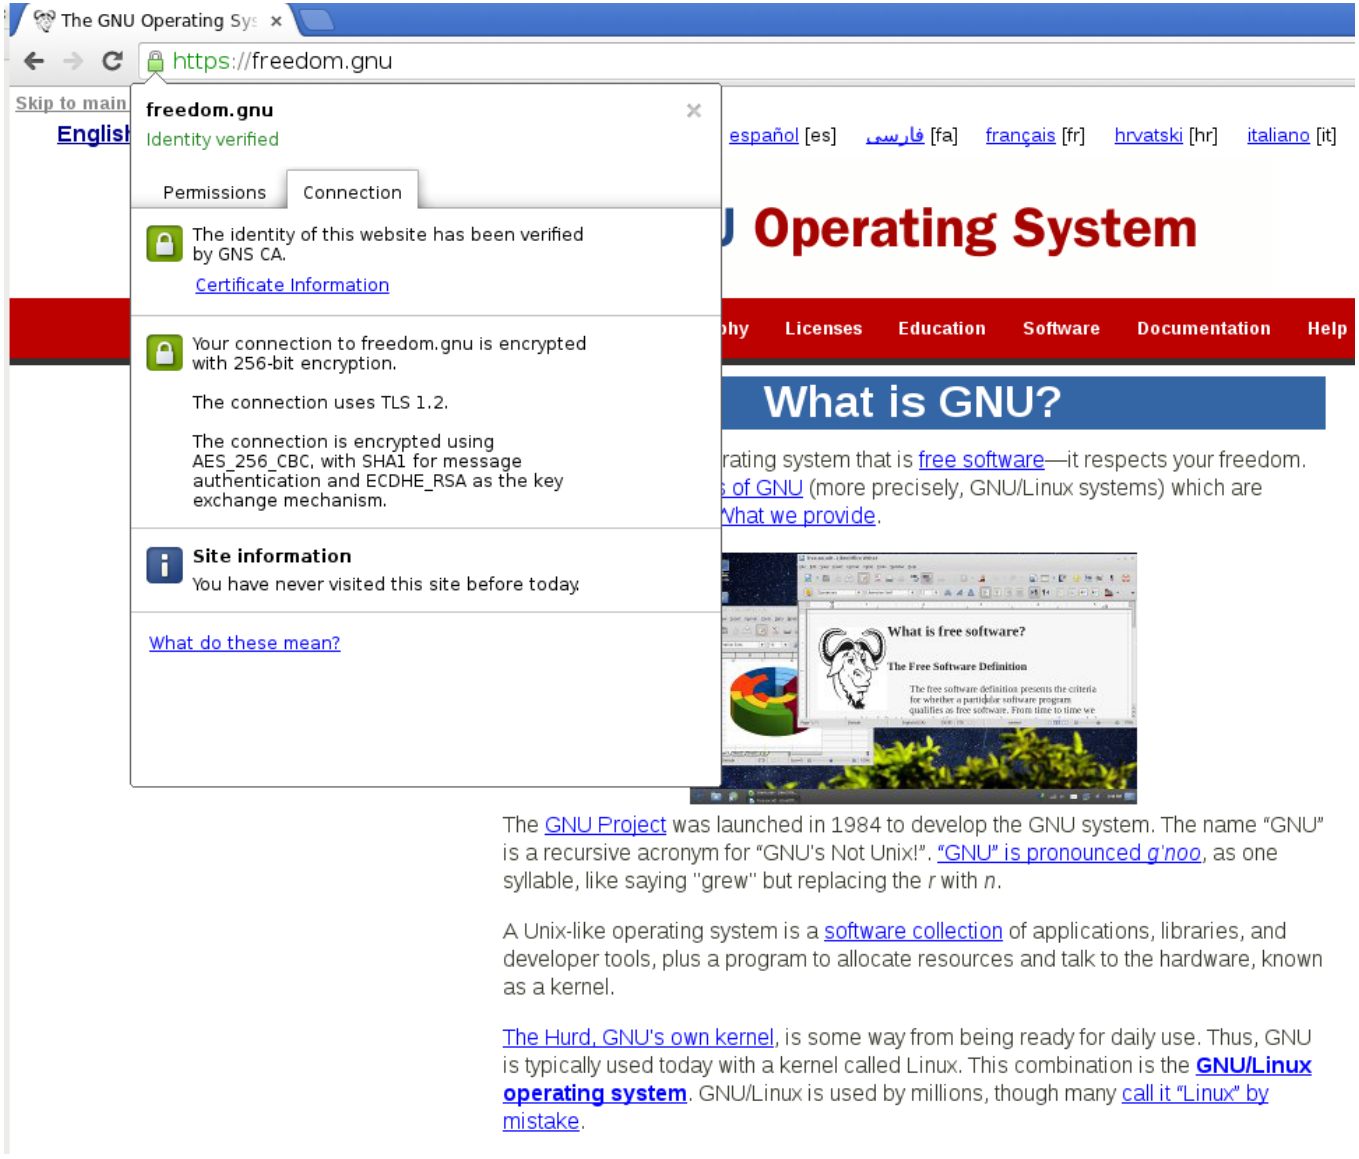

Privacy Issue: DHT

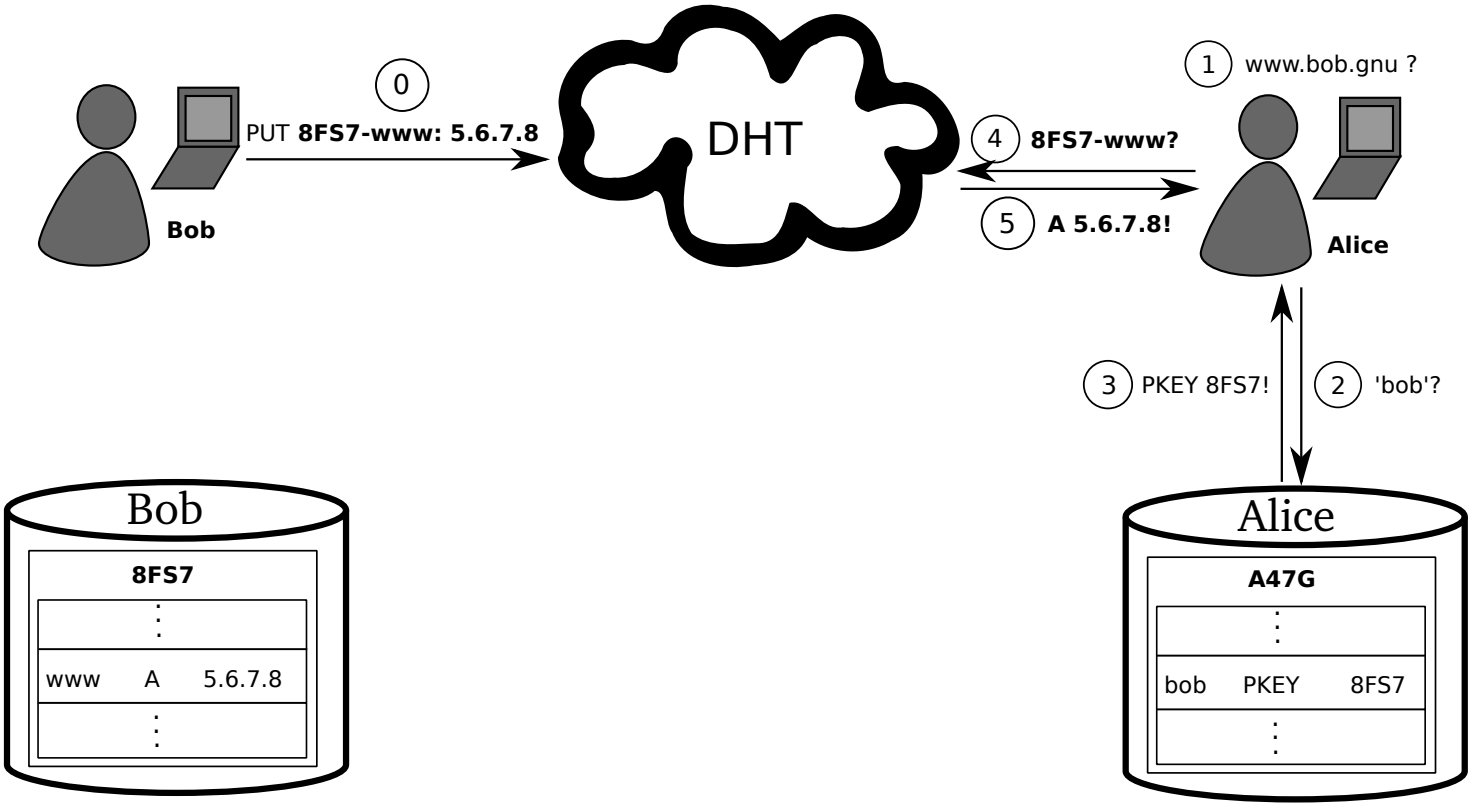

Query Privacy: Terminology

G	generator in ECC curve, a point
o	size of ECC group, $o := \|G\|$, o prime
x	private ECC key of zone ($x \in \mathbb{Z}_o$)
P	public key of zone, a point $P := xG$
l	label for record in a zone ($l \in \mathbb{Z}_o$)
$R_{P,l}$	set of records for label l in zone P
$q_{P,l}$	query hash (hash code for DHT lookup)
$B_{P,l}$	block with encrypted information for label l in zone P published in the DHT under $q_{P,l}$

Query Privacy: Cryptography

Publishing records $R_{P,I}$ as $B_{P,I}$ under key $q_{P,I}$

$$h := H(I, P) \tag{1}$$

$$d := h \cdot x \mod o \tag{2}$$

$$B_{P,I} := S_d(E_{HKDF(I,P)}(R_{P,I})), dG \tag{3}$$

$$q_{P,I} := H(dG) \tag{4}$$

Query Privacy: Cryptography

Publishing records $R_{P,I}$ as $B_{P,I}$ under key $q_{P,I}$

$$h := H(I, P) \tag{1}$$
$$d := h \cdot x \mod o \tag{2}$$
$$B_{P,I} := S_d(E_{HKDF(I,P)}(R_{P,I})), dG \tag{3}$$
$$q_{P,I} := H(dG) \tag{4}$$

Searching for records under label I in zone P

$$h := H(I, P) \tag{5}$$
$$q_{P,I} := H(hP) = H(hxG) = H(dG) \Rightarrow \texttt{obtain } B_{P,I} \tag{6}$$
$$R_{P,I} = D_{HKDF(I,P)}(B_{P,I}) \tag{7}$$

The ".zkey" Zone

- ".zkey" is another pTLD, in addition to ".gnu"
- In "LABEL.zkey", the "LABEL" is a public key of a zone
- "alice.bob.*KEY*.zkey" is perfectly legal

$\Rightarrow$ Globally unique identifiers

Key Revocation

- Revocation message signed with private key (ECDSA)
- Flooded on all links in P2P overlay, stored forever
- Efficient set reconciliation used when peers connect
- Expensive proof-of-work used to limit DoS-potential
- Proof-of-work can be calculated ahead of time
- Revocation messages can be stored off-line if desired

Shadow Records

- Records change
- Expiration time controls validity, like in DNS
- DHT propagation has higher delays, compared to DNS

Shadow Records

- Records change
- Expiration time controls validity, like in DNS
- DHT propagation has higher delays, compared to DNS
- SHADOW is a flag in a record
- Shadow records are only valid if no other, non-expired record of the same type exists

NICKnames

- “alice.bob.carol.dave.gnu” is a bit long for Edward (“.gnu”)
- Also, we need to trust Bob, Carol and Dave (for each lookup)
- Finally, Alice would have liked to be called Krista (just Bob calls her Alice)

NICKnames

- “alice.bob.carol.dave.gnu” is a bit long for Edward (“.gnu”)
- Also, we need to trust Bob, Carol and Dave (for each lookup)
- Finally, Alice would have liked to be called Krista (just Bob calls her Alice)
- “NICK” records allow Krista to specify her preferred NICKname
- GNS adds a “NICK” record to each record set automatically
- Edward learns the “NICK”, and software could automatically create “krista.gnu”

NICKnames

- "alice.bob.carol.dave.gnu" is a bit long for Edward (".gnu")
- Also, we need to trust Bob, Carol and Dave (for each lookup)
- Finally, Alice would have liked to be called Krista (just Bob calls her Alice)
- "NICK" records allow Krista to specify her preferred NICKname
- GNS adds a "NICK" record to each record set automatically
- Edward learns the "NICK", and software could automatically create "krista.gnu"
- Memorable, short trust path in the future! TOFU!
- Krista better pick a reasonably unique NICK.

Relative Names

- GNS records can contain ".+"
- CNAME: "server1.+"
- MX: "mail.+"
- ".+" stands for "relative to current zone"

Supporting this for links in browsers would be nice, too.

Legacy Hostname (LEHO) Records

LEHO records give a hint about the DNS name the server expects.

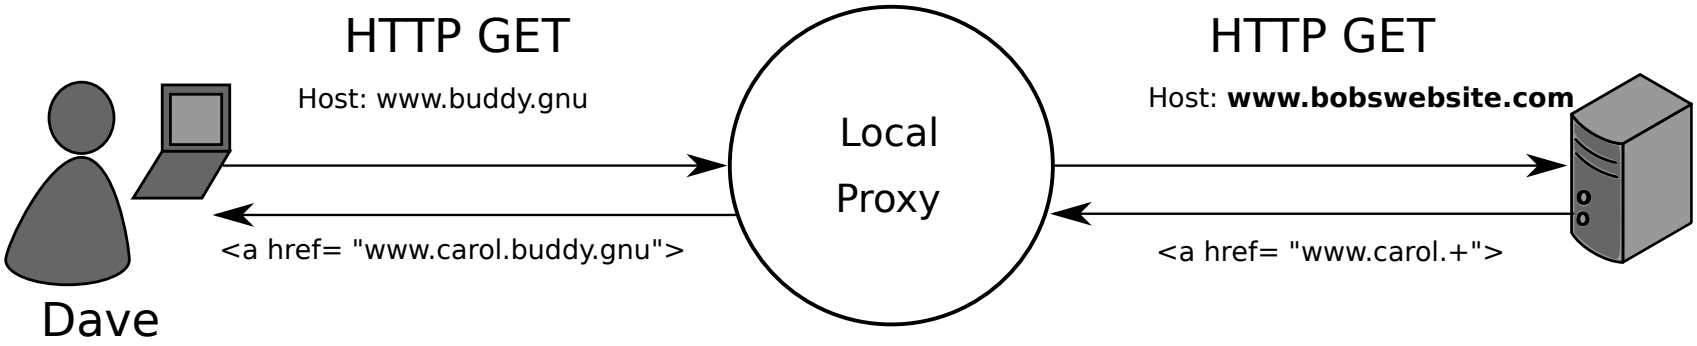

Legacy Hostname (LEHO) Records

LEHO records give a hint about the DNS name the server expects.

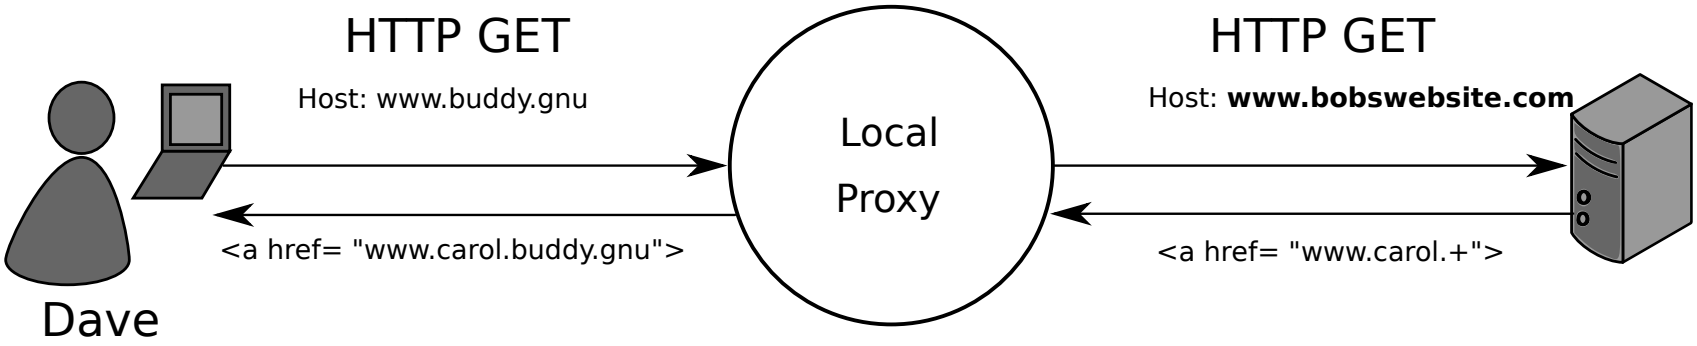

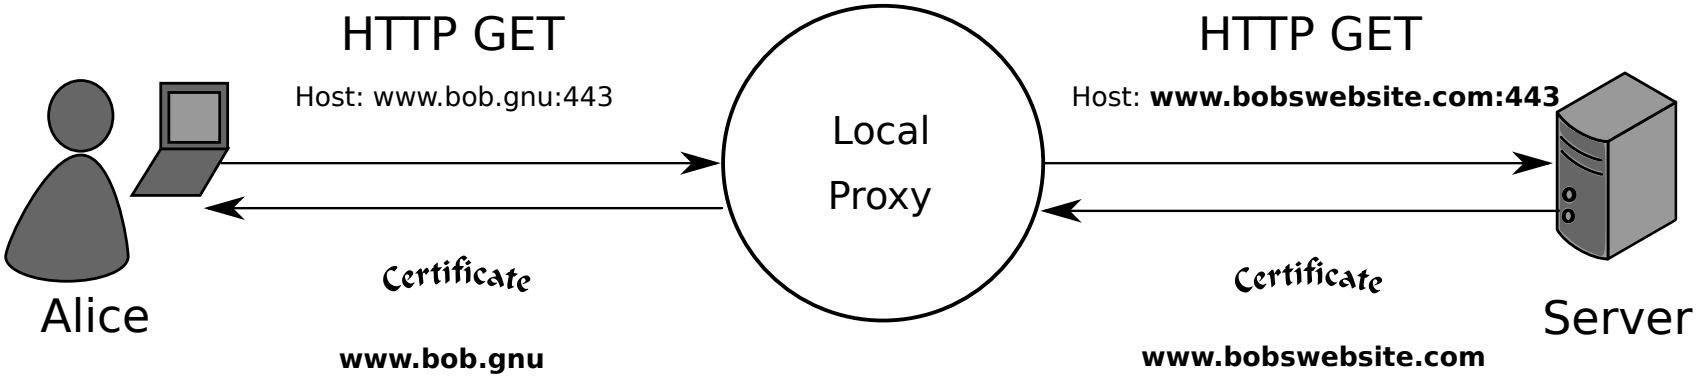

DNS Delegation

- Delegate to DNS using GNS2DNS records
- GNS2DNS record specifies:
  - Name of DNS resolver (i.e. "ns1.example.com" or "piratedns.+")
  - DNS domain to continue resolution in (i.e. "example.com" or "piratebay.org")
- GNS will first resolve DNS resolver name to A/AAAA record
- GNS will then resolve "*left.of.gns2dns.*example.com" using DNS

Fun GNS Record Types

- DNS CERT: store your GPG public key
- GNUNET VPN: TCP/IP services hosted in GNUnet
- GNUNET PHONE: have a conversation

Application Integration

- SOCKS proxy (`gnunet-gns-proxy`)
- NSS plugin
- GNS (C) API
- GNS (IPC) protocol
- GNS command-line tool

Summary

- Interoperable with DNS
- Globally unique identifiers with ".zkey"
- Delegation allows using zones of other users
- Trust paths explicit, trust agility
- Simplified key exchange compared to Web-of-Trust
- Privacy-enhanced queries, censorship-resistant
- Reliable revocation

Privacy summary

Method	Defense against MiTM	Zone privacy	Privacy vs. network	Privacy vs. operator	Traffic amplification resistance	Censorship resistance	Ease of migration
DNS	✗	✓	✗	✗	✗	✗	✓
DNSSEC	✓	✗	✗	✗	✗	✗	✗*
DNSCurve	✓	✓	✓	✗	✓	✗	✗
DNS-over-TLS	✓	n/a	✓	✗	✓	✗	✗
Namecoin	✓	✗	✓	✓	✓	✓	✗
RAINS	✓	✗	✓	✗	✓	✗	✗
GNS	✓	✓	✓	✓	✓	✓	✗

*EDNS0

Key management summary

	Suitable for personal use	Memorable	Decentralised	Modern cryptography	Understandable	Exposes metadata	Transitive
DNS	✗	✓	✗	✗	✗	✗	✓
DNSSEC	✗	✓	✗	✗	✗	✗	✓
DNSCurve	✗	✓	✗	✓	✗	✗	✓
DNS-over-TLS	✗	✓	✗	✗	✗	✗	✓
TLS-X.509	✗	✓	✗	✗	✗	✗	✓
Web of Trust	✓	✗	✓	✗	✗	✗	✓
TOFU	✓	✗	✓		✓	✓	✗
Namecoin	✗	✓	✗	✓	✓	✗	✓
RAINS	✗	✓	✗	✓	✓	✗	✓
GNS	✓	✓	✓	✓	✓	✓	✓

Ongoing and Future Work (Project 2, BS theses)

- Optimze GNUnet DHT
- Import ".fr" TLD into GNS and hijack it!
- Implement & evaluate bounded Eppstein set reconciliation
- Integrate GNS with Tor

Conclusion

- Query name minimization is low-cost, low-benefit approach, but should clearly be done
- Simple encryption schemes offer medium-cost, medium-benefit approach
- GNU Name System performance depends on the DHT
  $\Rightarrow$ need to invest more in DHT design & implementation

DNS	globalist
DNSSEC	authoritarian
Namecoin	libertarian (US)
RAINS	nationalist
GNS	anarchist

In which world do you want to live?

Do you have any questions?

References:

- Nathan Evans and Christian Grothoff. *R5N. Randomized Recursive Routing for Restricted-Route Networks*. **5th International Conference on Network and System Security**, 2011.
- Matthias Wachs, Martin Schanzenbach and Christian Grothoff. *On the Feasibility of a Censorship Resistant Decentralized Name System*. **6th International Symposium on Foundations & Practice of Security**, 2013.
- M. Schanzenbach *Design and Implementation of a Censorship Resistant and Fully Decentralized Name System*. **Master's Thesis (TUM)**, 2012.