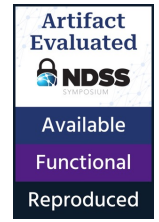

IRIS: Dynamic Privacy Preserving Search in Authenticated Chord Peer-to-Peer Networks

Angeliki Aktypi, Kasper Rasmussen

Network and Distributed System Security (NDSS) Symposium
27th February 2025

UNIVERSITY OF OXFORD
Linacre College
DEPARTMENT OF COMPUTER SCIENCE
CENTRE FOR DOCTORAL TRAINING in CYBER SECURITY

Privacy-Preserving Search in Chord

- Chord: a P2P lookup service
  - Protocols: BOOTSTRAP, LOOKUP, FETCH, PUSH, etc.
  - Algorithms: RETRIEVE, STORE, JOIN, LEAVE, etc.

- Acknowledged for it its simplicity and high performance

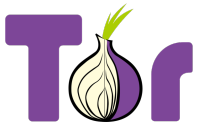

- Goal: RETRIEVE (and STORE) does not leak **what** a requester searches for

  … even if nodes have long-term identities

System & Adversary Model

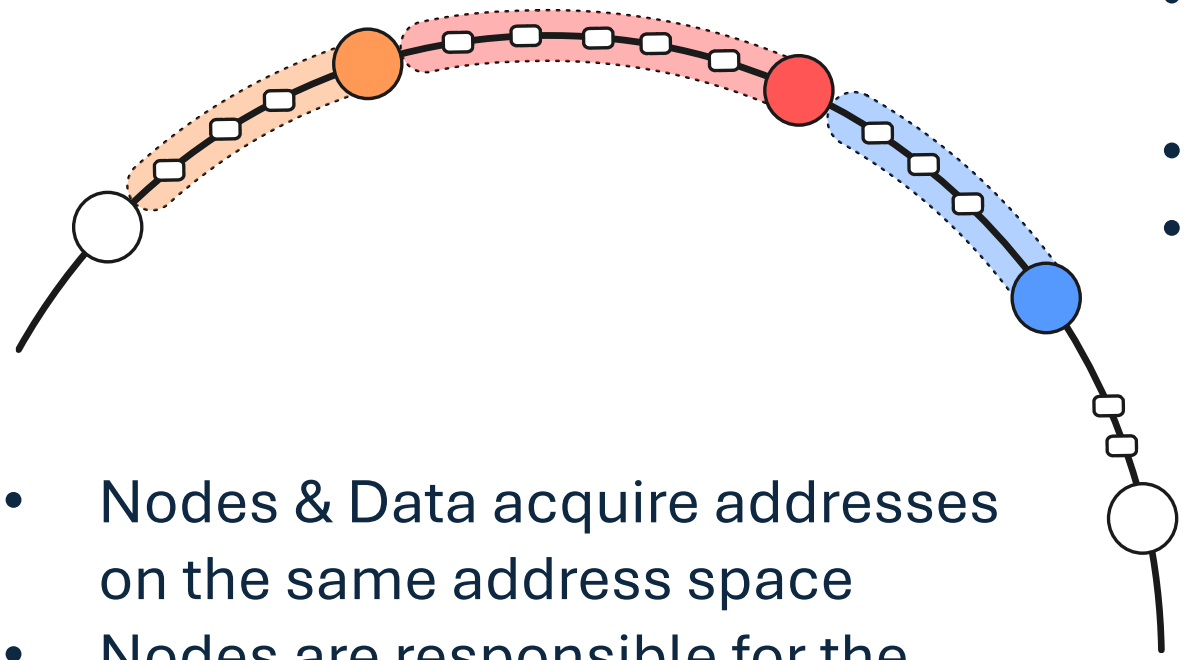

- Nodes & Data acquire addresses on the same address space
- Nodes are responsible for the Data that are ‘close’ to them

- **Internal colluding attacker** who knows all the Chord algorithms nodes use
- Nodes have **long-term** identities
- Communication is **confidential** and **integrity** protected

Chord RETRIEVE Algorithm

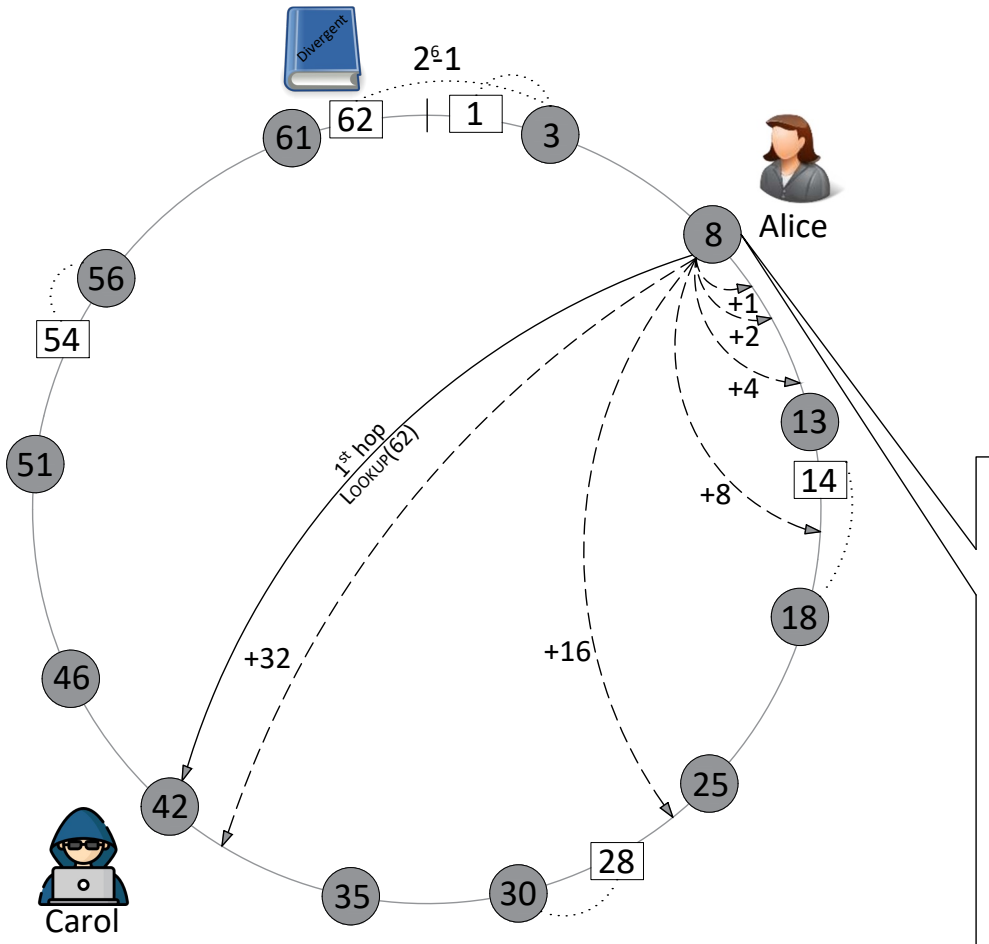

Predecessor		Comm. Info
3		ip_3, $port_3$
Routing Table		
Address	**Node**	**Comm. Info**
8+1	13	ip_{13} , $port_{13}$
8+2	13	ip_{13} , $port_{13}$
8+4	13	ip_{13} , $port_{13}$
8+8	18	ip_{18} , $port_{18}$
8+16	30	ip_{30} , $port_{30}$
8+32	42	ip_{42} , $port_{42}$

Chord RETRIEVE Algorithm

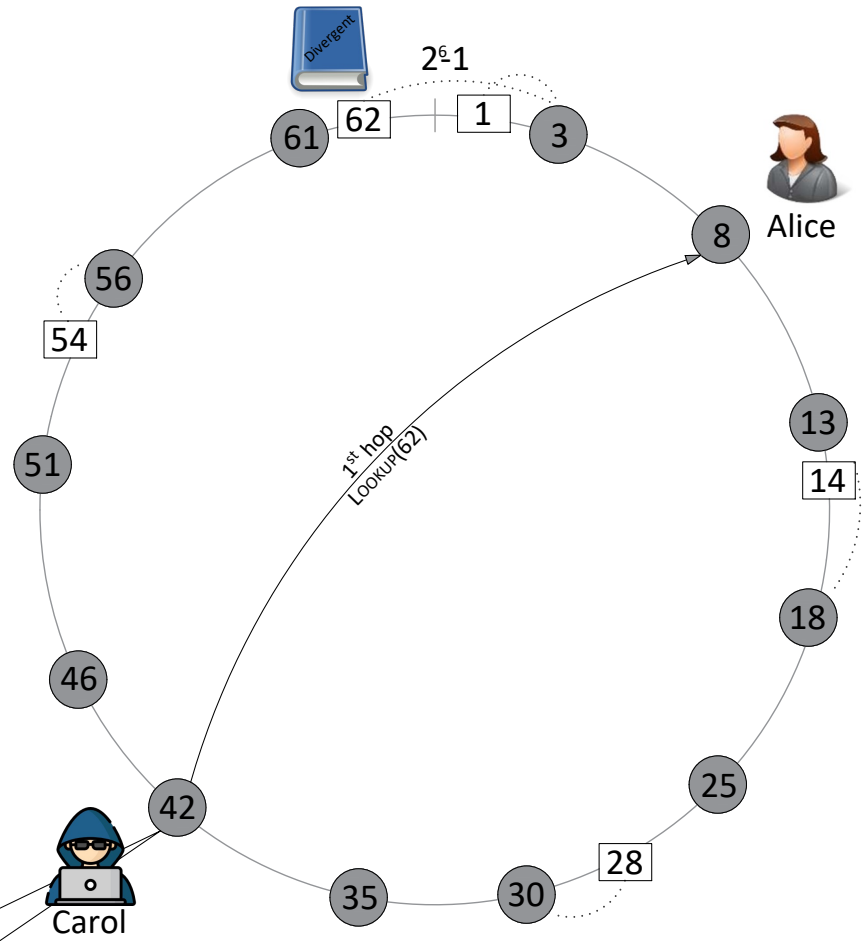

Predecessor		Comm. Info
35		ip_{35}, $port_{35}$
Routing Table		
Address	**Node**	**Comm. Info**
42+1	46	ip_{46} , $port_{46}$
42+2	46	ip_{46} , $port_{46}$
42+4	46	ip_{46} , $port_{46}$
42+8	51	ip_{51} , $port_{51}$
42+16	61	ip_{61} , $port_{61}$
42+32	13	ip_{13} , $port_{13}$

Chord RETRIEVE Algorithm

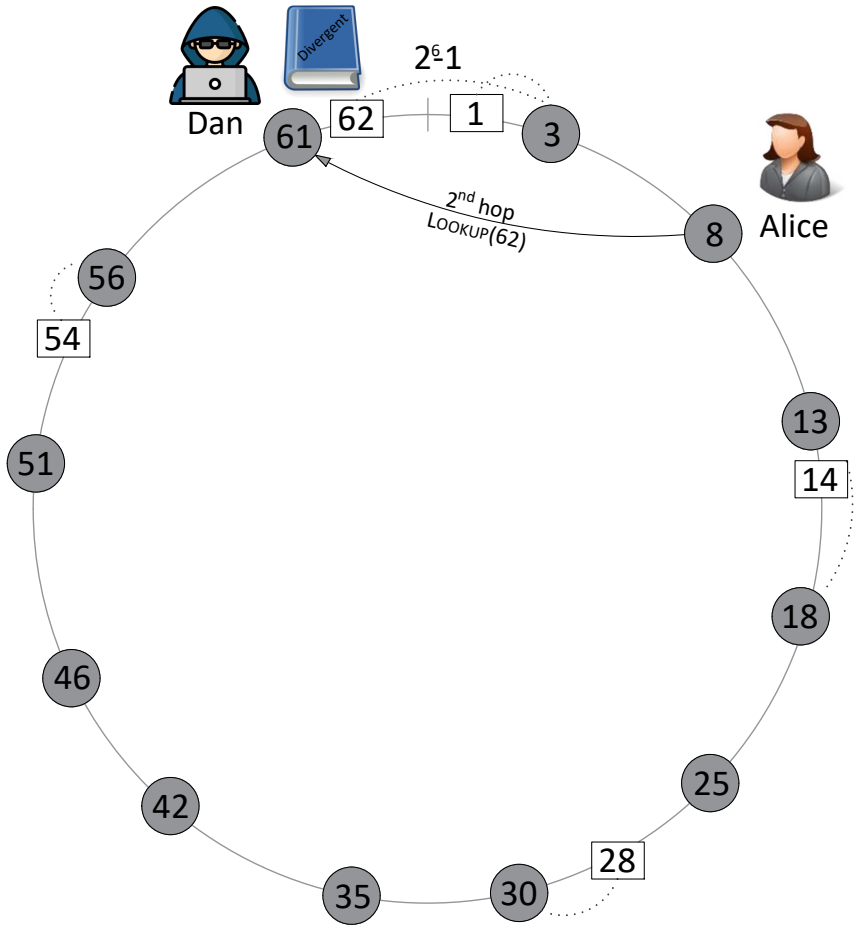

Chord RETRIEVE Algorithm

Predecessor		Comm. Info
56		ip_{56}, $port_{56}$
Routing Table		
Address	**Node**	**Comm. Info**
61+1	3	ip_3 , $port_3$
61+2	3	ip_3 , $port_3$
61+4	3	ip_3 , $port_3$
61+8	8	ip_8 , $port_8$
61+16	13	ip_{13} , $port_{13}$
61+32	25	ip_{25} , $port_{25}$

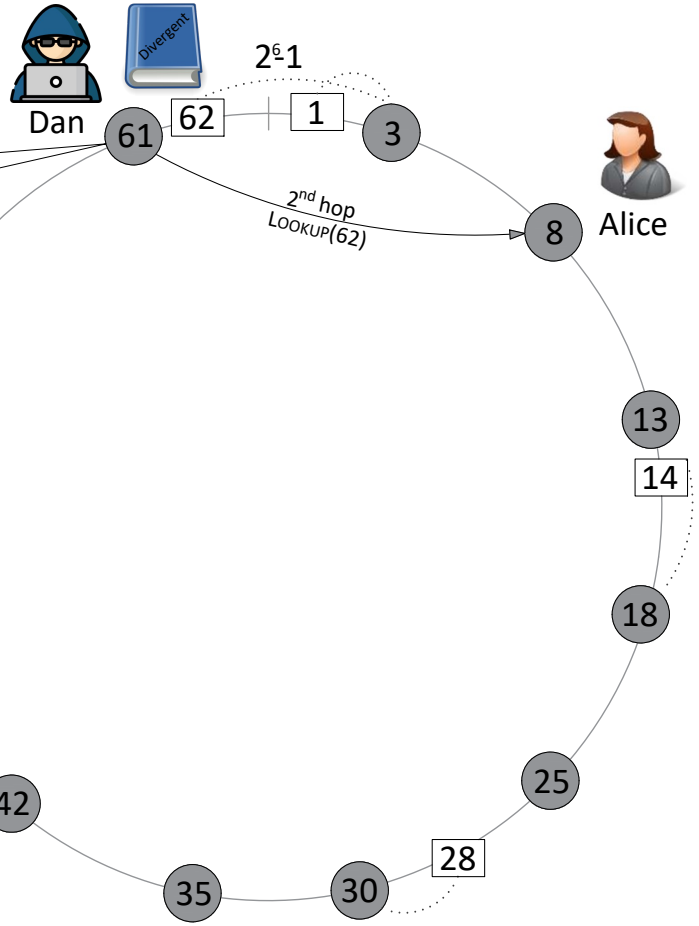

Chord RETRIEVE Algorithm

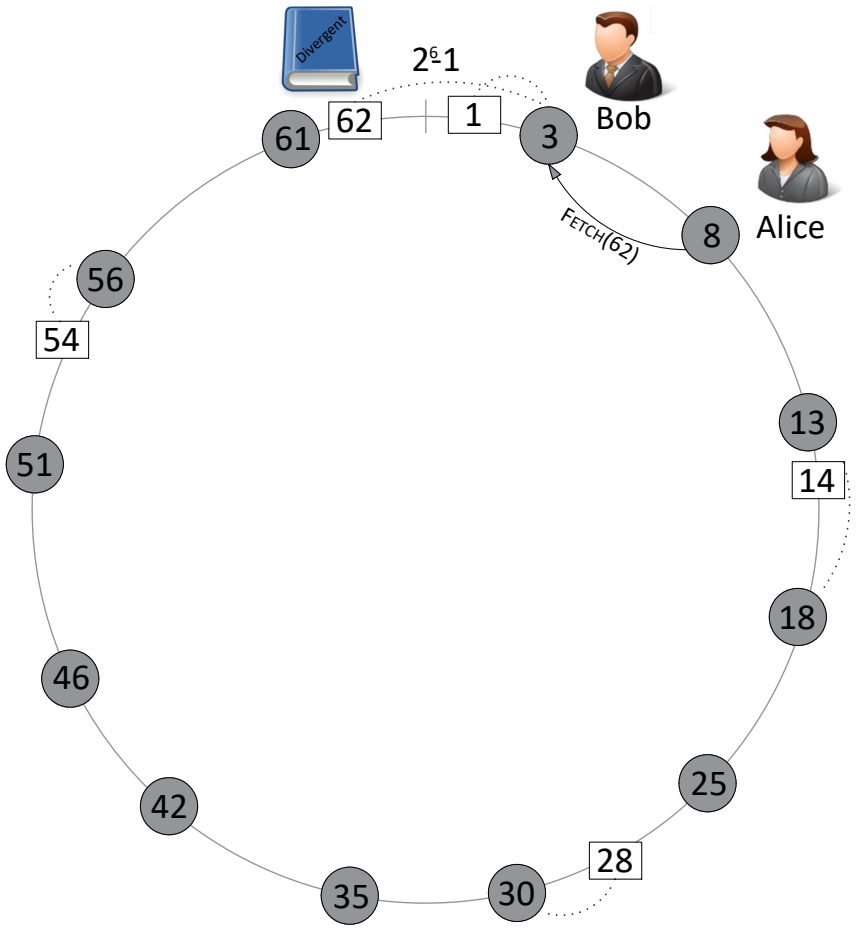

Chord RETRIEVE Algorithm

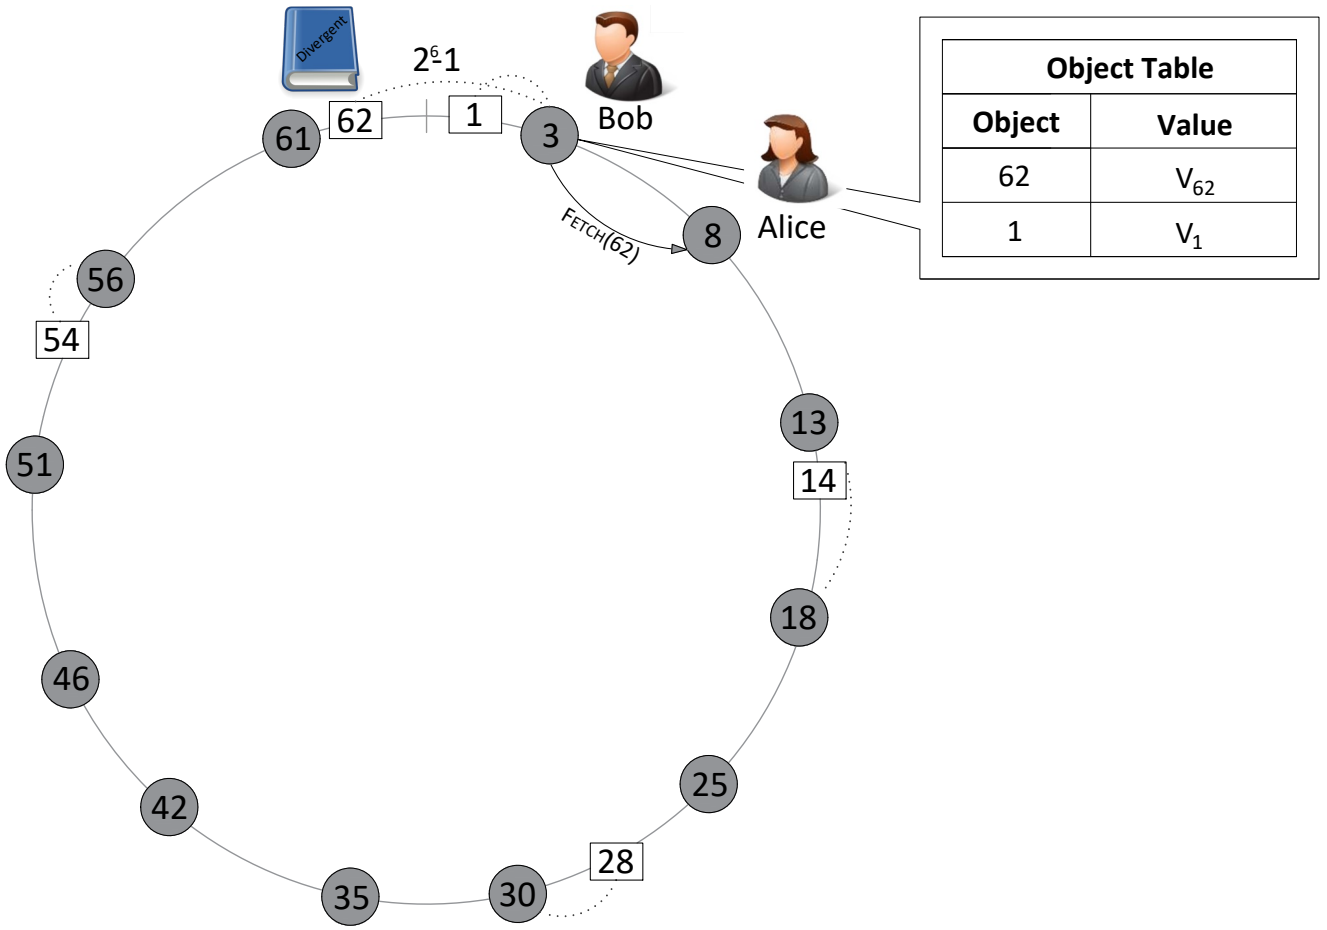

Privacy Problem

- All the nodes that take part in the RETRIEVE algorithm know **what** the requester is searching for
  - The nodes can request the same hash and see what comes back

- A malicious node can:
  - Log requests
  - Identify popular/new content
  - Track usage patterns

IRIS

1: **function** IRIS_RETRIEVE (RT_r, O_p, α, δ)
2: $N_i \leftarrow$ SELECTSTARTNODE(RT_r, O_p, δ)
3: **repeat**
4: $R_i \leftarrow$ RANDOMADDRESSBETWEEN(N_i, O_p)
5: $I_i = (1-\alpha)R_i + \alpha N_i$
6: $N_i \leftarrow$ LOOKUP(N_i, I_i)
7: **until** N_i owns O_p
8: **return** FETCH(N_i, O_p)
9: **end function**

RT_r : the requester's Routing Table, O_p : the target object, α & δ : the privacy parameters

IRIS

1: **function** IRIS_RETRIEVE (RT_r, O_p, α, δ)
2: N_i←SELECTSTARTNODE(RT_r, O_p, δ)
3: **repeat**
4: R_i←RANDOMADDRESSBETWEEN(N_i, O_p)
5: $I_i = (1-\alpha)R_i + \alpha N_i$
6: N_i←LOOKUP(N_i, I_i)
7: **until** N_i owns O_p
8: **return** FETCH(N_i, O_p)
9: **end function**

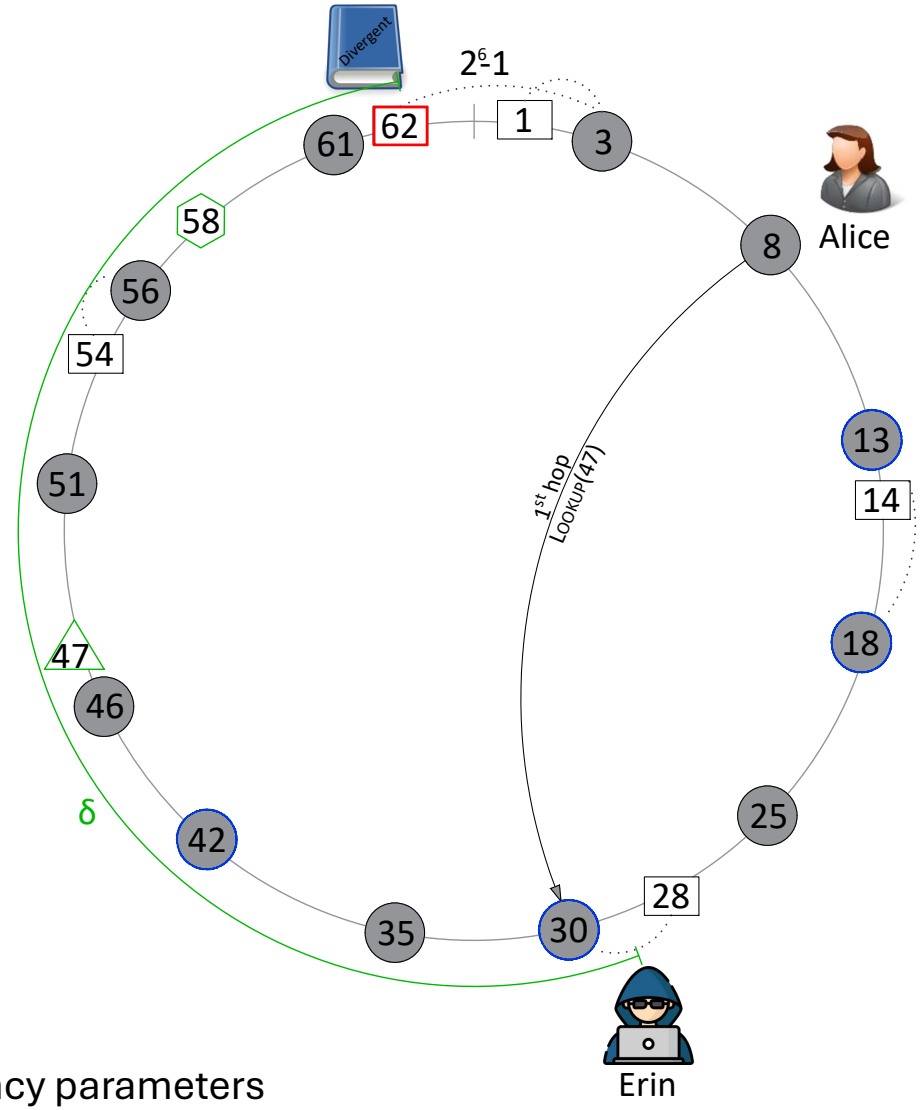

RT_r: the requester's Routing Table, O_p: the target object, α & δ : the privacy parameters

How to measure privacy?

- Iterative query process: smaller set of possible targets on every step
- A single privacy value: captures the privacy on every step
- Adopting k-anonymity: final set tiny, reflects the worst-case scenario

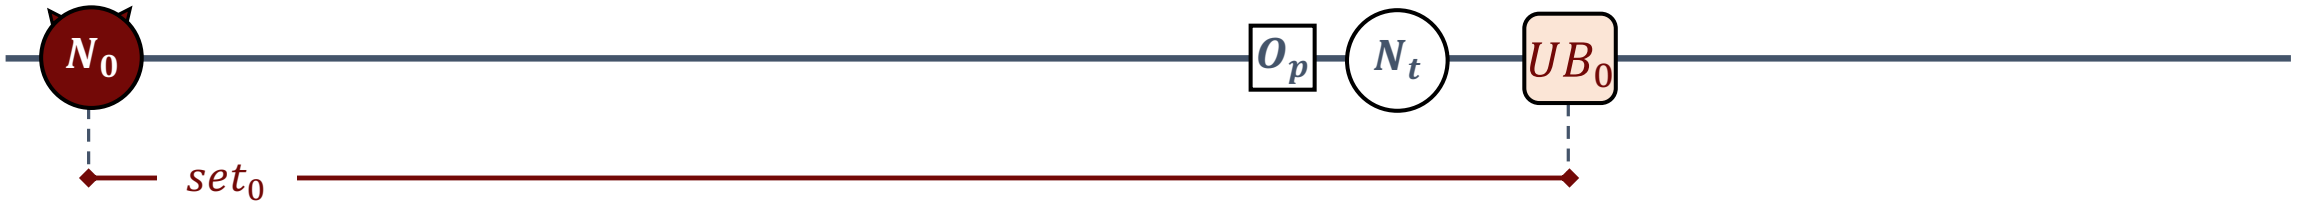

How to measure privacy?

- Iterative query process: smaller set of possible targets on every step
- A single privacy value: captures the privacy on every step
- Adopting k-anonymity: final set tiny, reflects the worst-case scenario

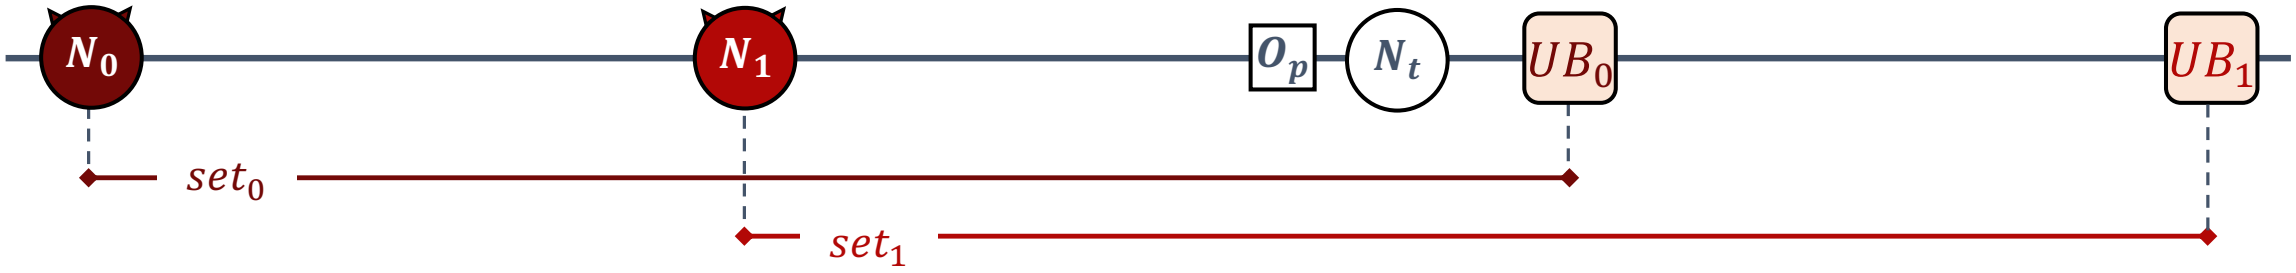

How to measure privacy?

- Iterative query process: smaller set of possible targets on every step
- A single privacy value: captures the privacy on every step
- Adopting k-anonymity: final set tiny, reflects the worst-case scenario

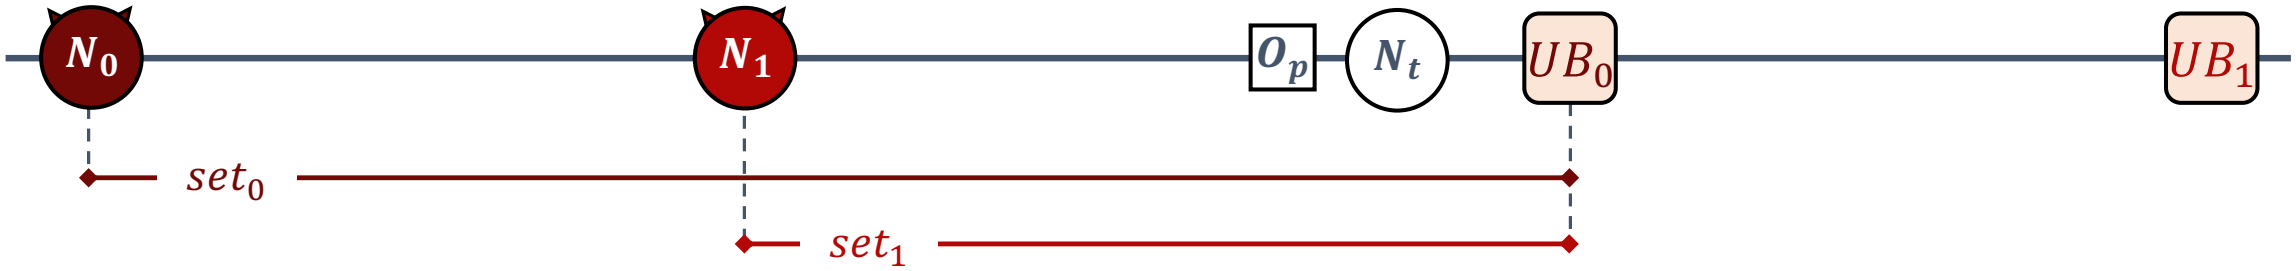

(α,δ)-privacy

- A RETRIEVE algorithm is (α,δ)-private if the following two conditions hold:
  - ✓ $prior_0 \geq \delta$ for the first queried node N_0
  - ✓ $posterior_i / prior_i \geq a$ for every iteration i>0

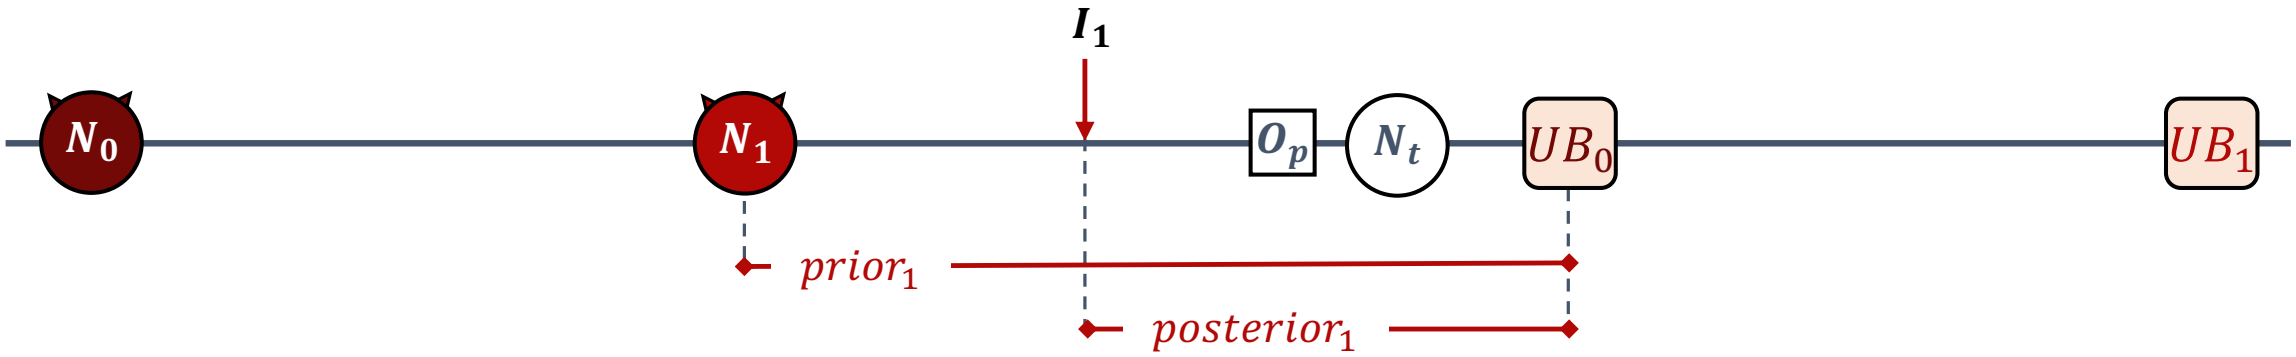

IRIS Analysis

- ✓ Correctness - IRIS converges on the target
- ✓ Privacy - IRIS is an (α,δ)-private algorithm

(Q1) How do the (α,δ)-privacy parameters influence IRIS's convergence?
(Q2) What is the queried nodes' privacy ratio distribution?

- ➢ We simulate Chord and IRIS algorithms on MATLAB:
  - Address space of 2^{23}
  - 1000 participating nodes
  - Steady state network
  - Random requester and target object
  - Open-sourced: https://github.com/angakt/iris

Convergence Speed – Dependance on α and δ

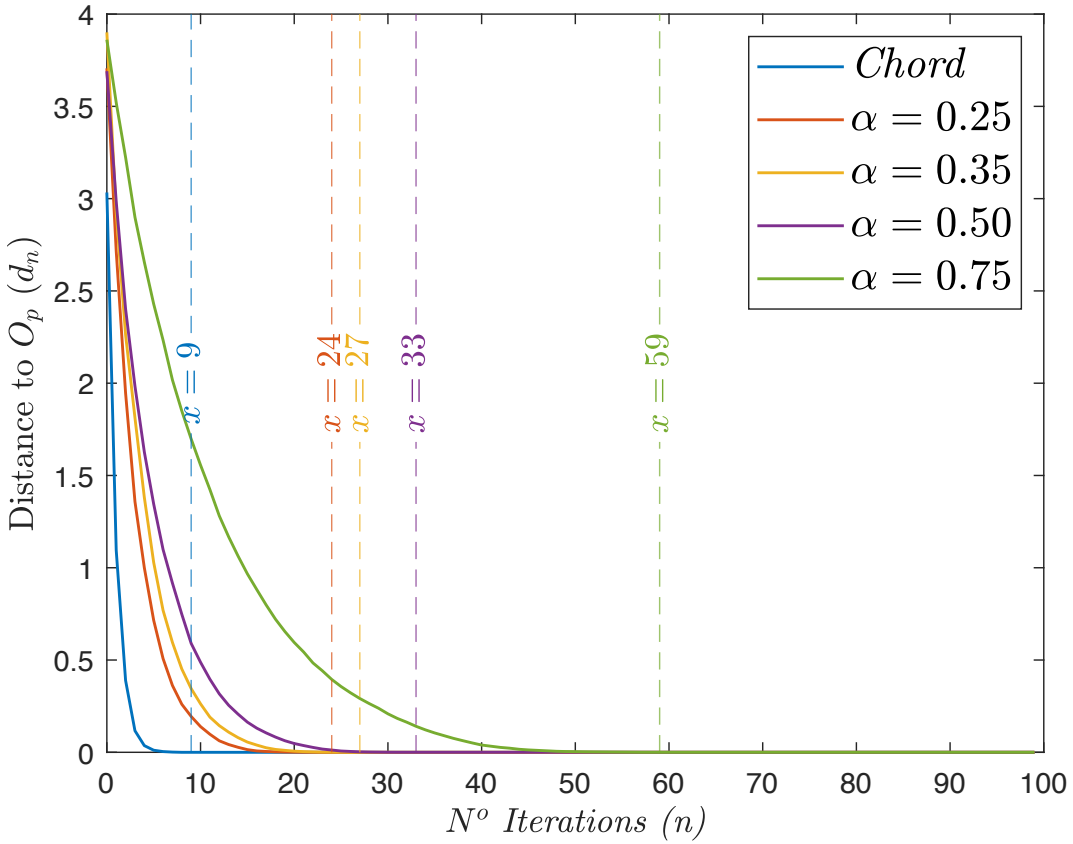

experiments done with δ= 2^{23}/16

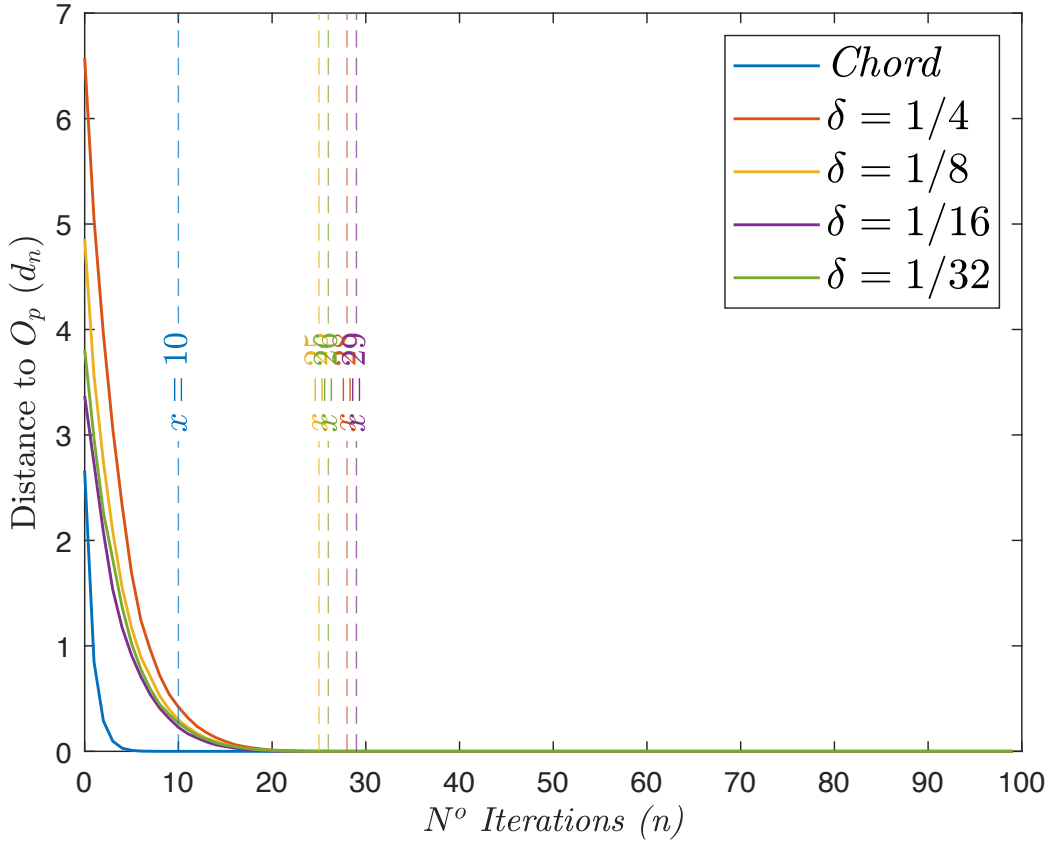

experiments done with α=0.35

Privacy – Actually we do better!

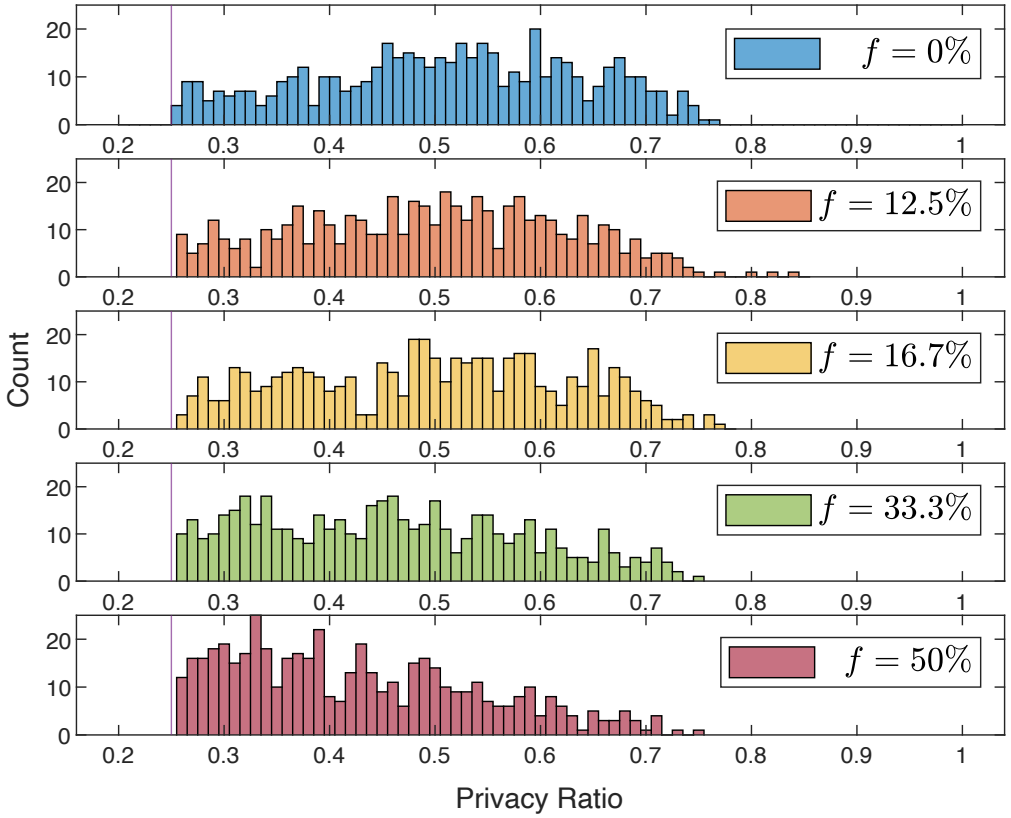

- α **limits** the information loss per iteration
- an attacker gains only a small amount of **extra** information from each request
- in most cases the information gain is **less** than α

Summary

IRIS: privacy preserving search in authenticated Chord P2P networks

Adjustable trade-off between privacy and efficiency

Works with vanilla Chord – facilitates adoption by already deployments

Summary

IRIS: privacy preserving search in authenticated Chord P2P networks

Adjustable trade-off between privacy and efficiency

Works with vanilla Chord – facilitates adoption by already deployments

Thank you!

Angeliki Aktypi
angeliki.aktypi@cs.ox.ac.uk

IRIS - Correctness

At each iteration, the distance to the final object is (on average) given by:

$$d_{i+1} = \frac{d_i}{2} + a\frac{d_i}{2} = d_i\left(\frac{1+a}{2}\right)$$

Given $d_0 = \delta$, after n iteration the distance is:

$$d_n = \delta\left(\frac{1+a}{2}\right)^n$$

IRIS converges on the target:

$$\alpha \in [0,1) \Rightarrow \lim_{n\to\infty} d_n = 0$$

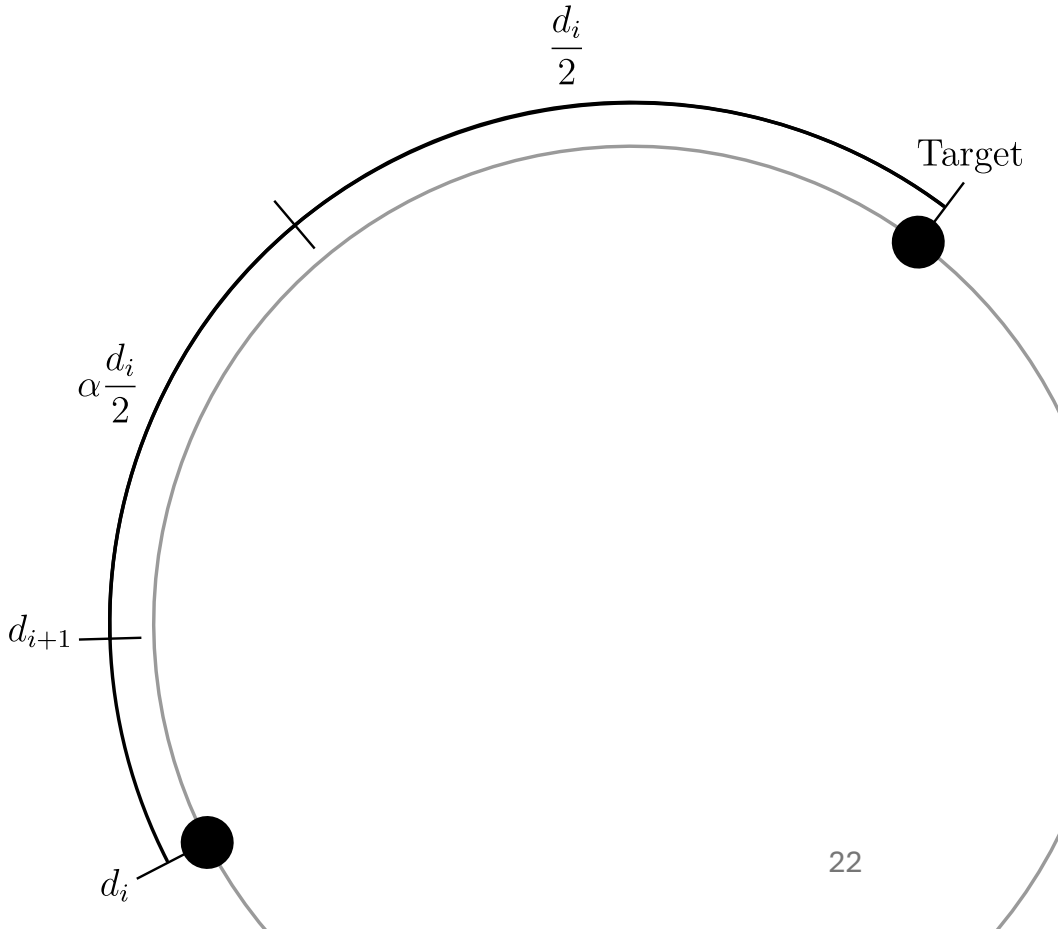

IRIS - Correctness

At each iteration, the distance to the final object is (on average) given by:

$$d_{i+1} = \frac{d_i}{2} + a\frac{d_i}{2} = d_i\left(\frac{1+a}{2}\right)$$

Given $d_0 = \delta$, after n iteration the distance is:

$$d_n = \delta\left(\frac{1+a}{2}\right)^n$$

IRIS converges on the target:

$$\alpha \in [0,1) \Rightarrow \lim_{n\to\infty} d_n = 0$$