Ting: Measuring and Exploiting Latencies Between All Tor Nodes

Frank Cangialosi Dave Levin Neil Spring

University of Maryland

Measuring latencies

Ping

Measurement Host

External Host

Limited to the nodes
we **control**

Measuring latencies

Ping

Measurement Host

External Host

Limited to the nodes we **control**

To gain **broader** insight, we can:

1. Control more nodes?

2. Estimate latencies?

Measuring latencies

Ping

Goal

Measurement Host

External Host

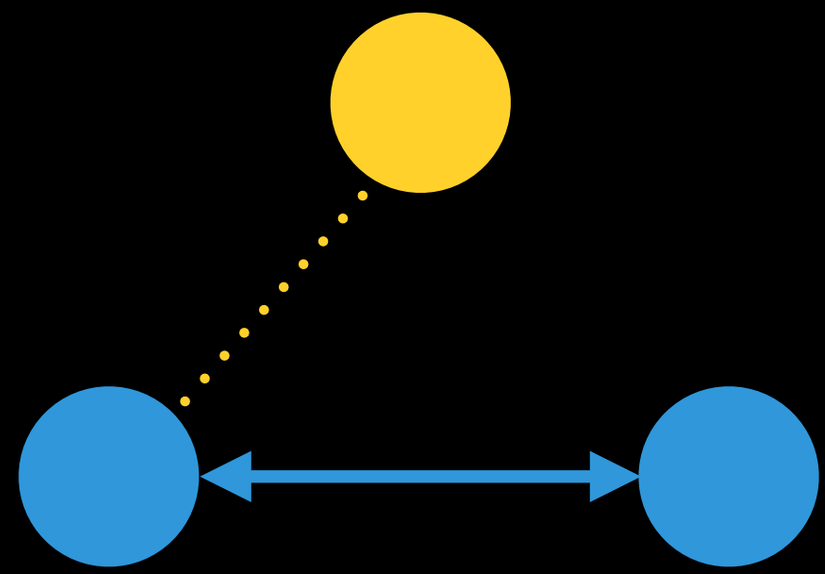

Limited to the nodes we **control**

Latency between **arbitrary** nodes

King technique

[Gummadi et al, 2002]

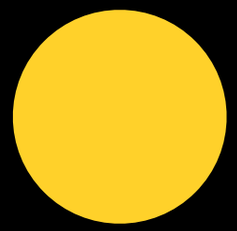

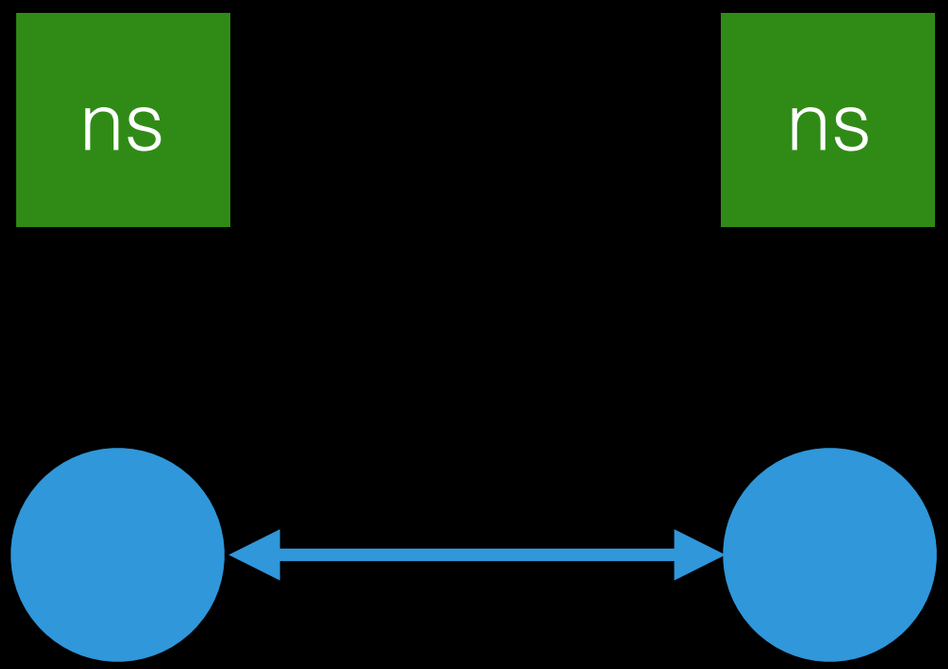

King technique

[Gummadi et al, 2002]

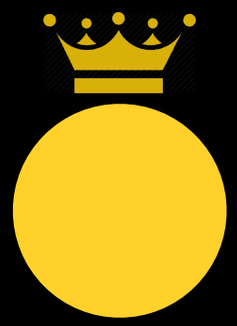

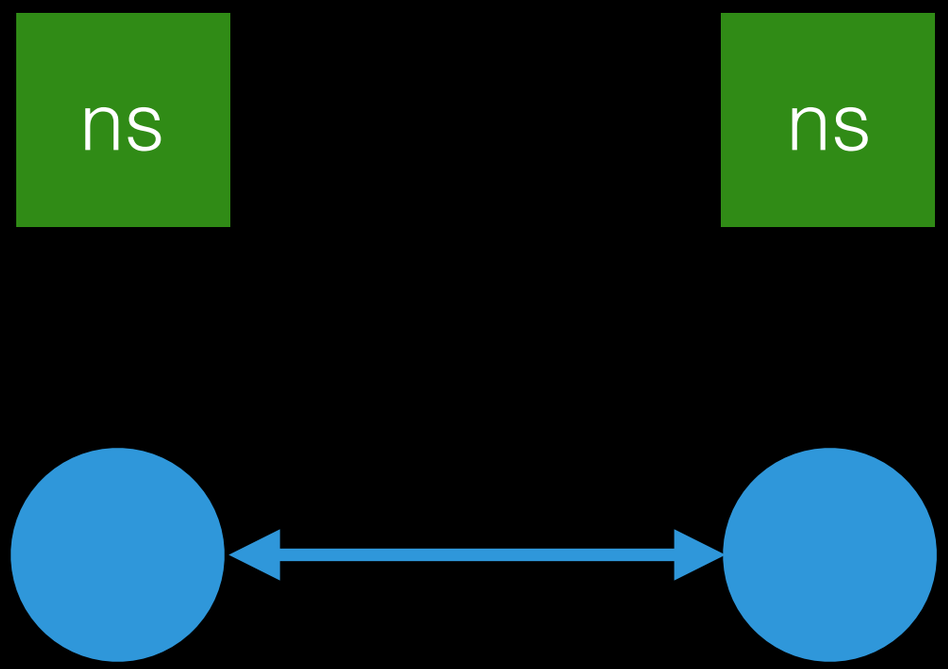

King technique

[Gummadi et al, 2002]

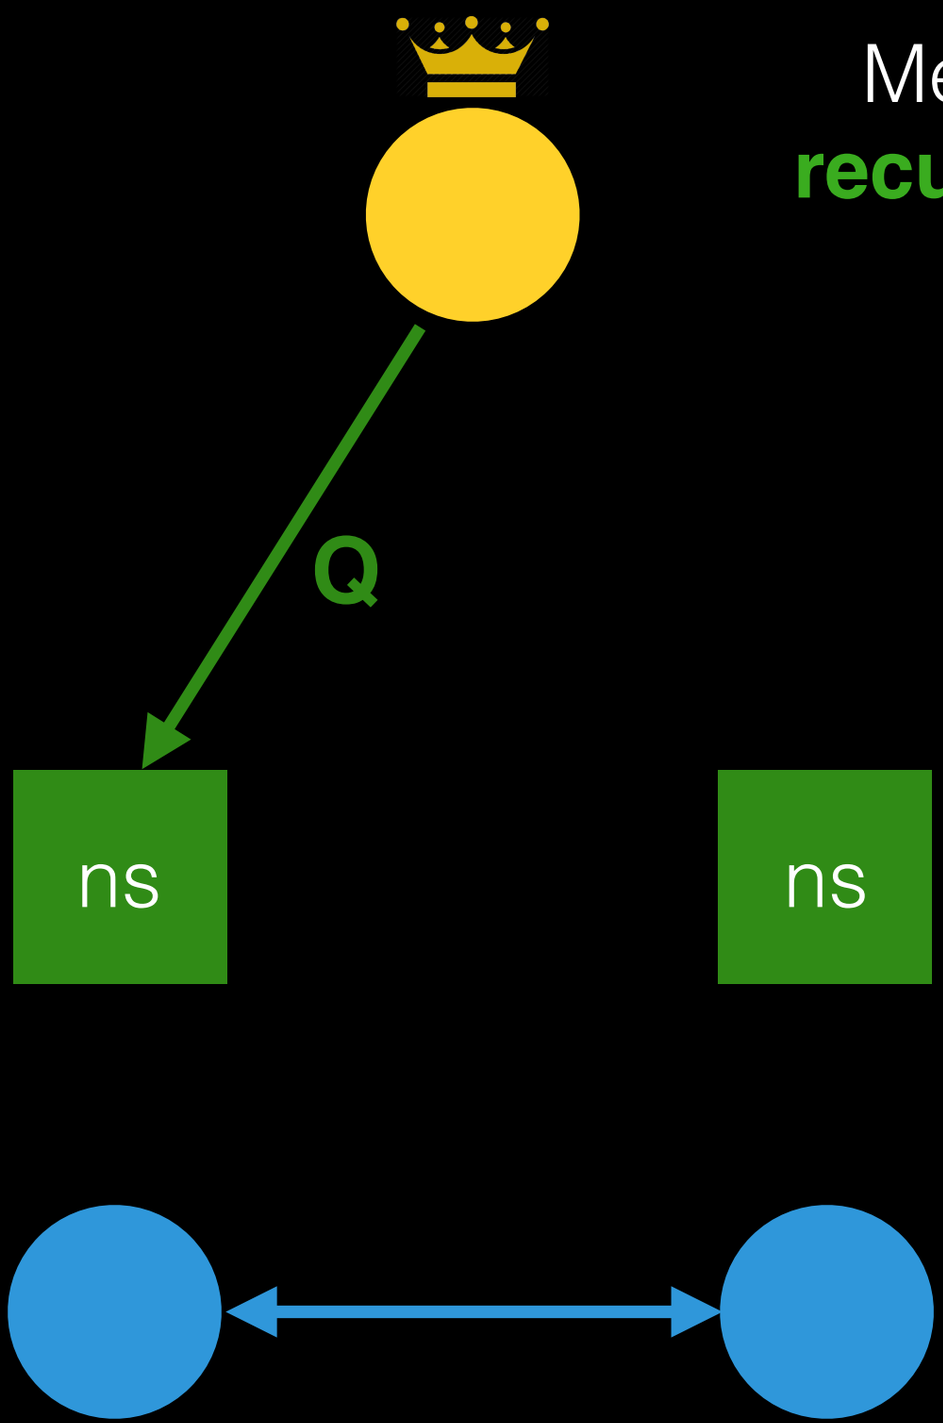

Measure the RTT of
recursive DNS queries

King technique

[Gummadi et al, 2002]

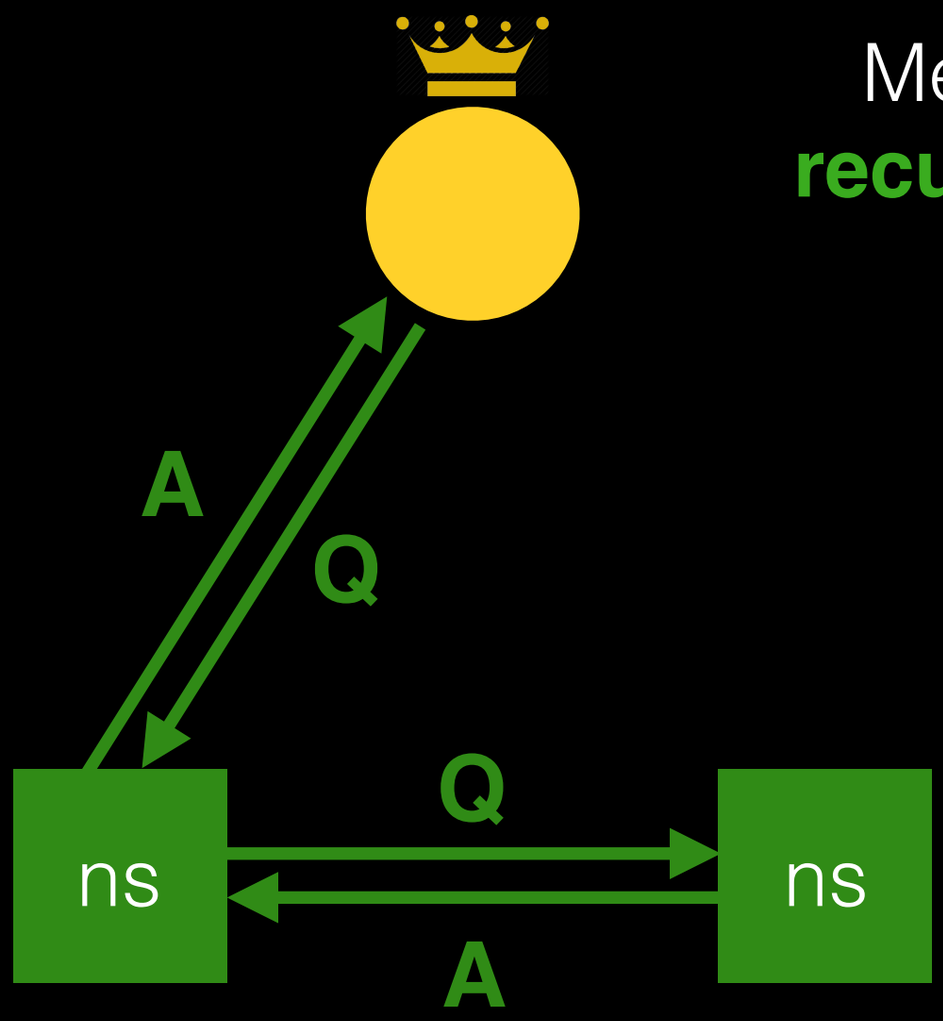

Measure the RTT of
recursive DNS queries

King technique

[Gummadi et al, 2002]

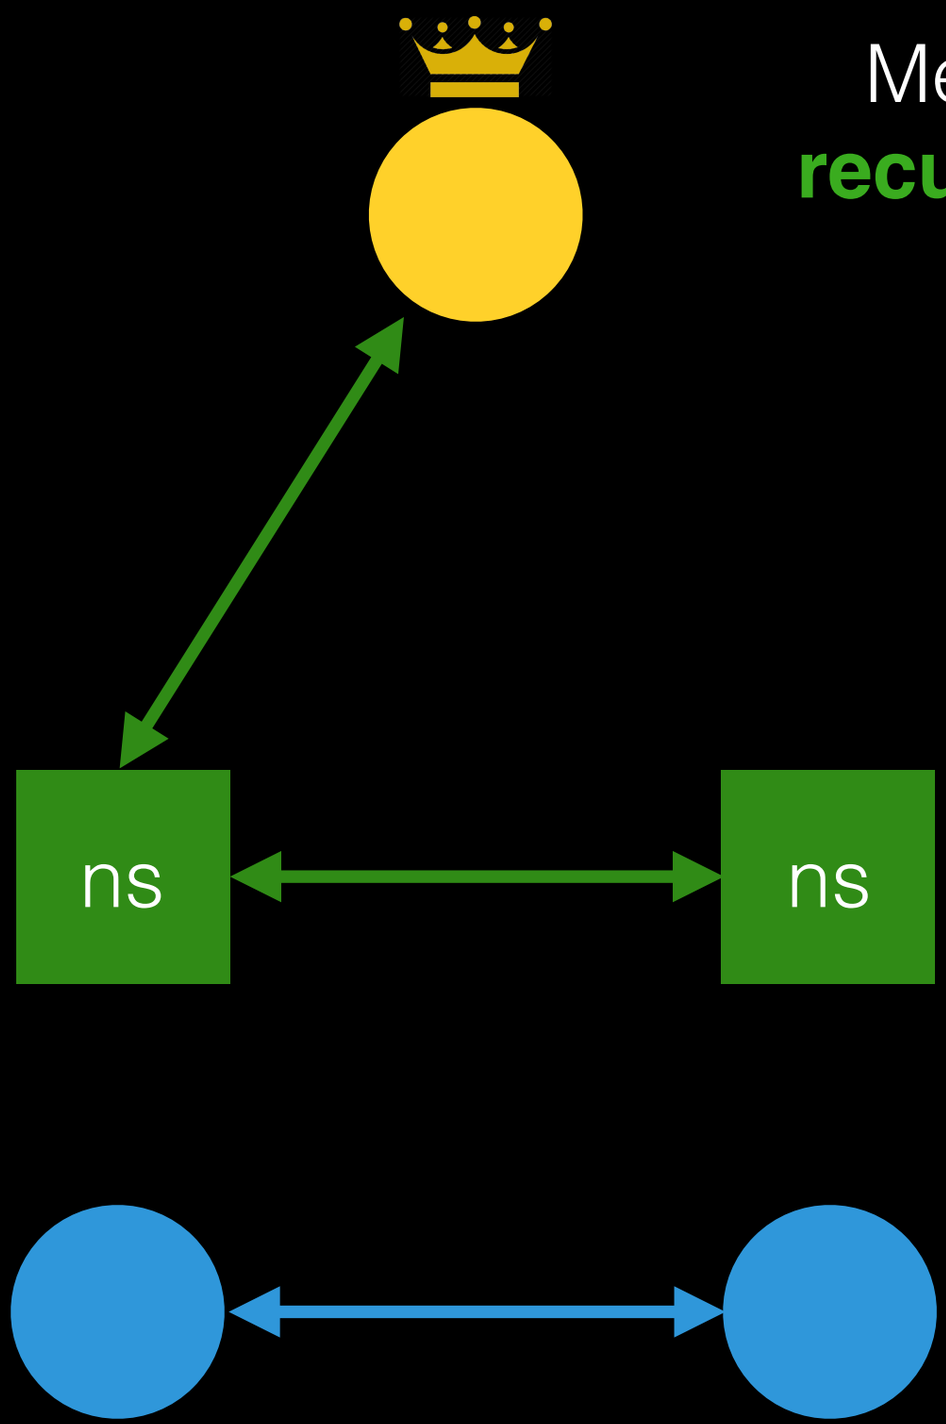

Measure the RTT of
recursive DNS queries

King technique

[Gummadi et al, 2002]

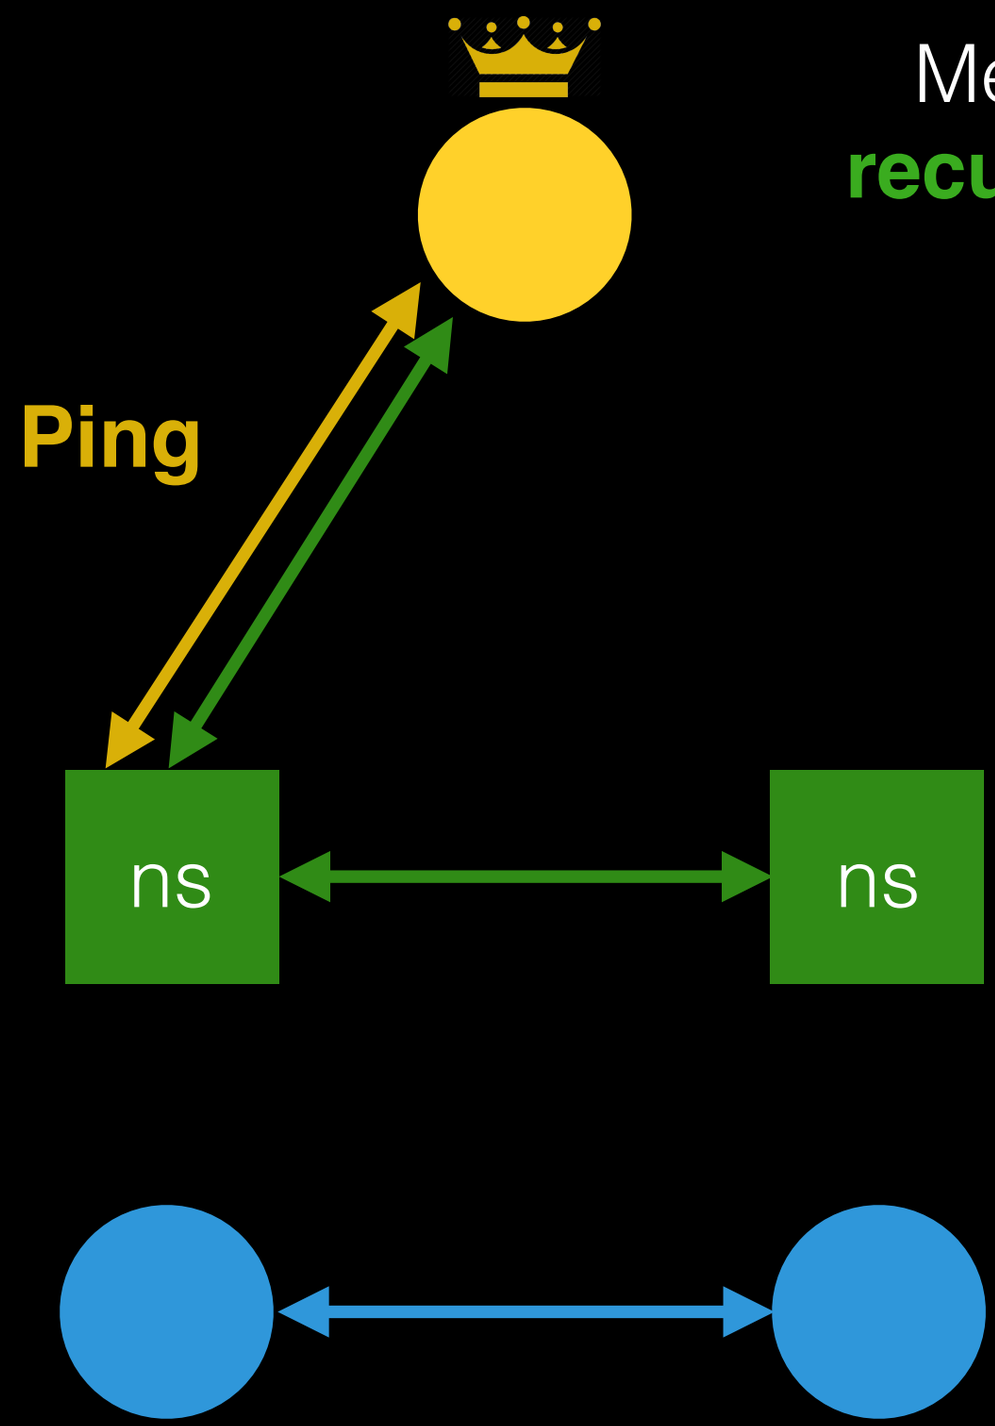

Measure the RTT of
recursive DNS queries

King technique

[Gummadi et al, 2002]

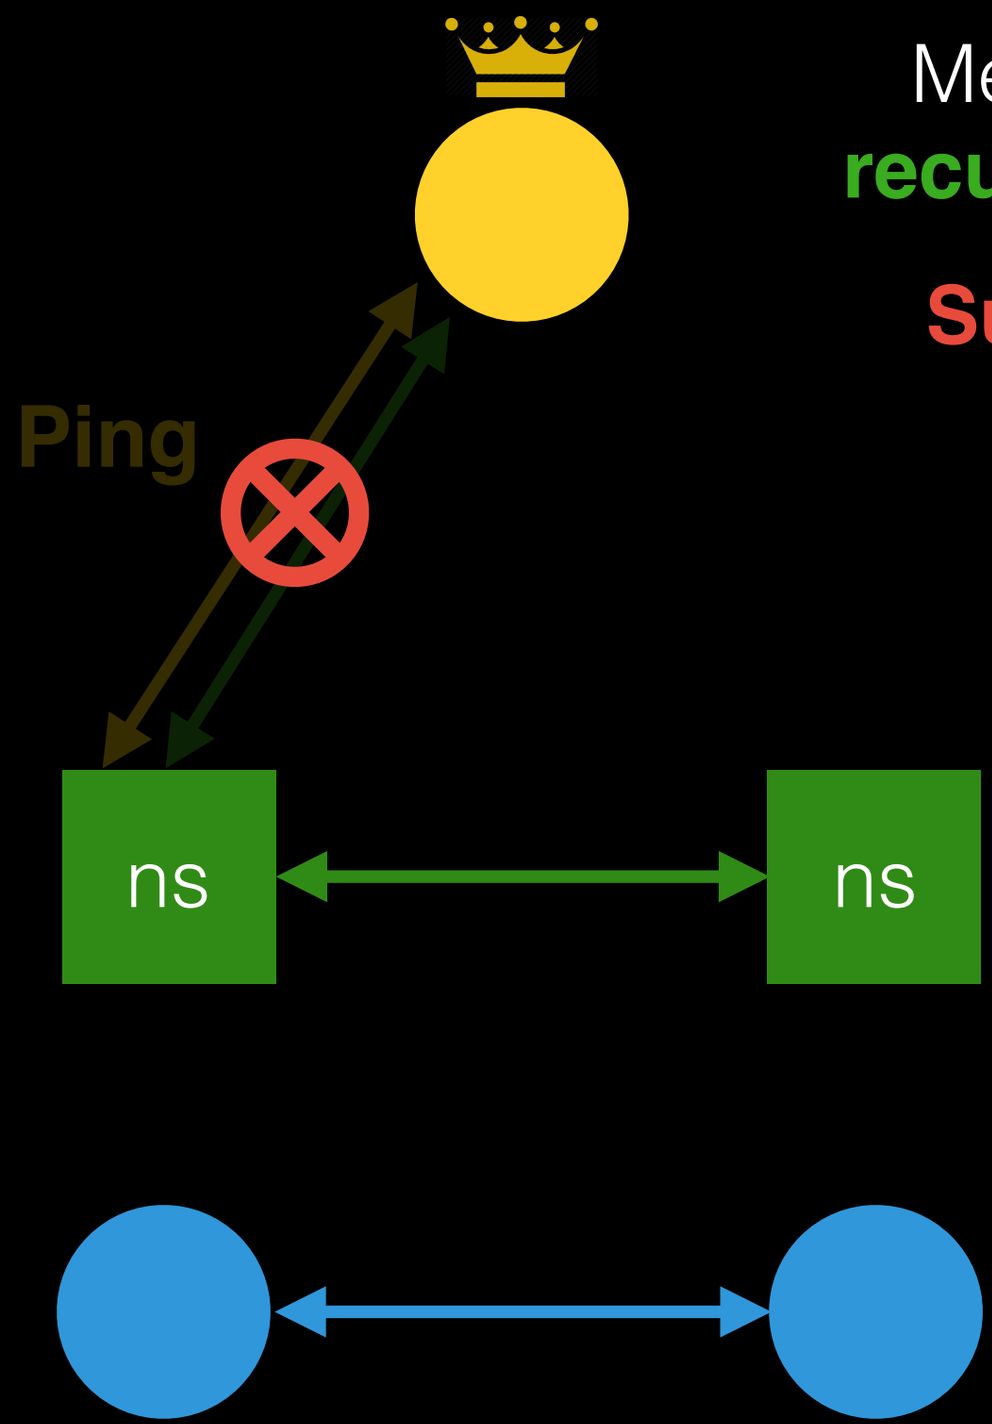

Measure the RTT of **recursive** DNS queries

Subtract latencies

King technique

[Gummadi et al, 2002]

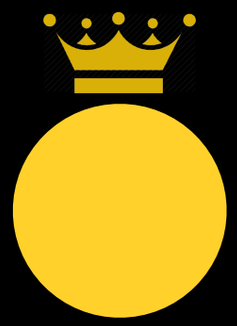

Measure the RTT of **recursive** DNS queries

Subtract latencies

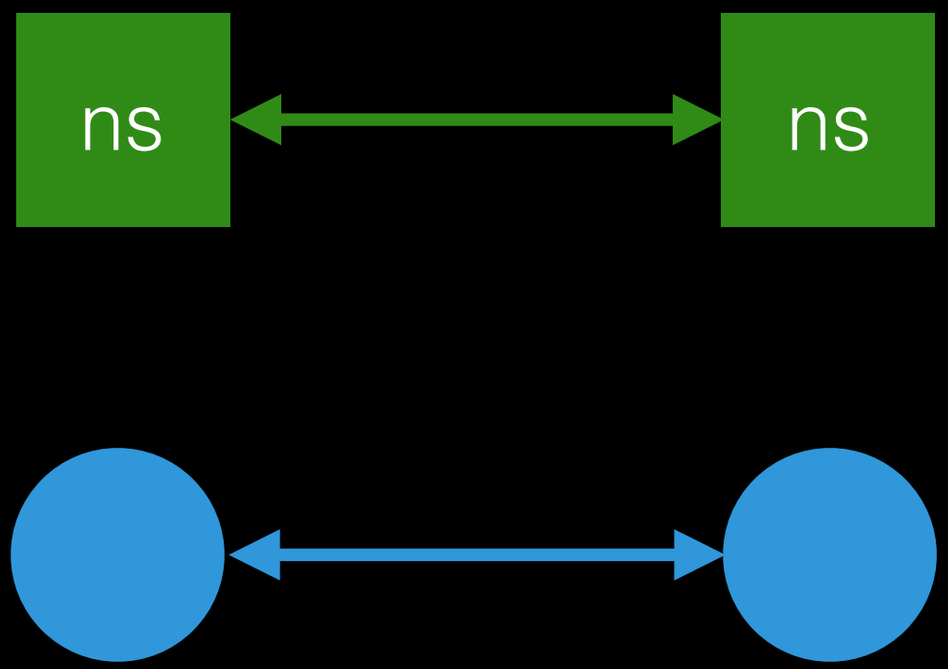

King technique

[Gummadi et al, 2002]

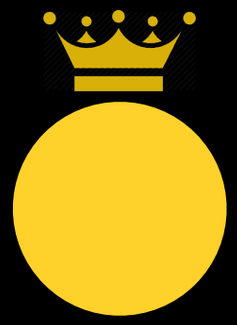

Measure the RTT of **recursive** DNS queries

Subtract latencies

King does not *directly* measure the path between two hosts

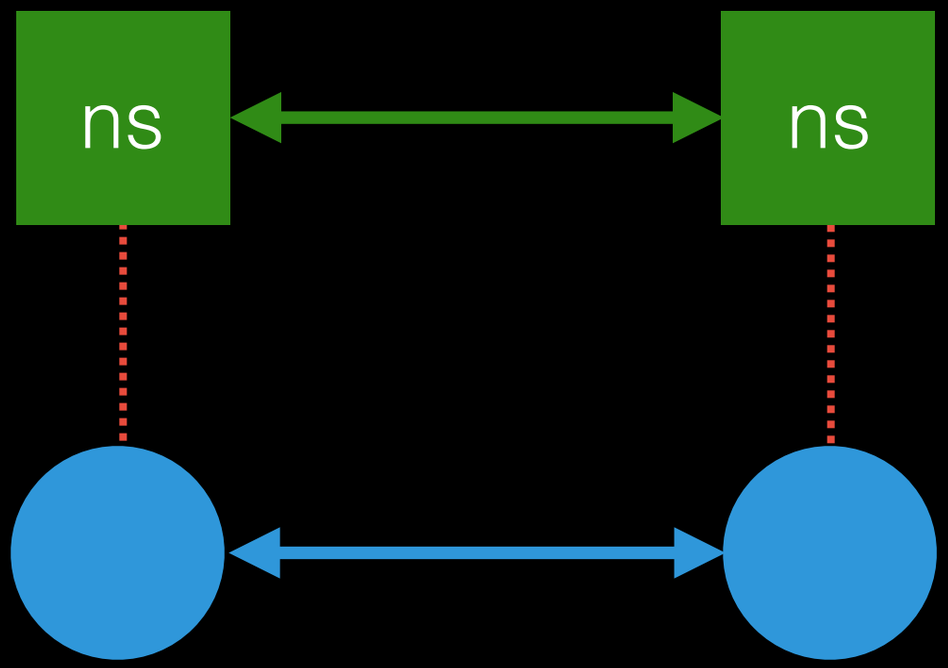

King technique

[Gummadi et al, 2002]

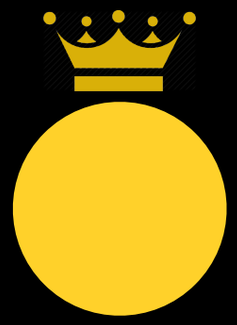

Measure the RTT of **recursive** DNS queries

Subtract latencies

King does not *directly* measure the path between two hosts

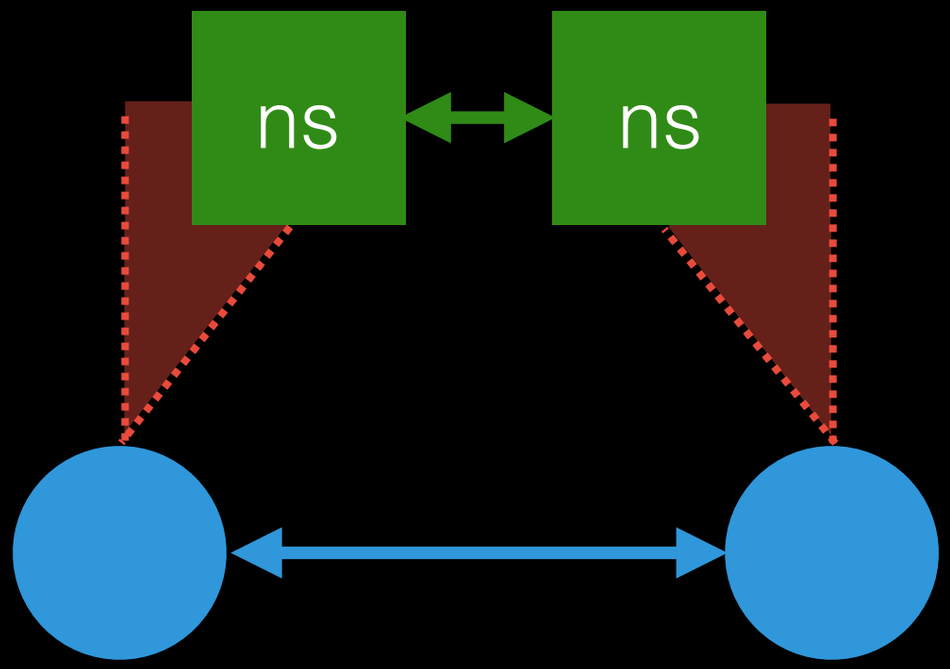

King technique

[Gummadi et al, 2002]

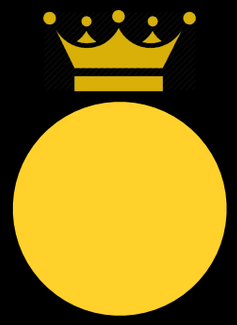

Measure the RTT of
recursive DNS queries

Subtract latencies

King does not *directly* measure
the path between two hosts

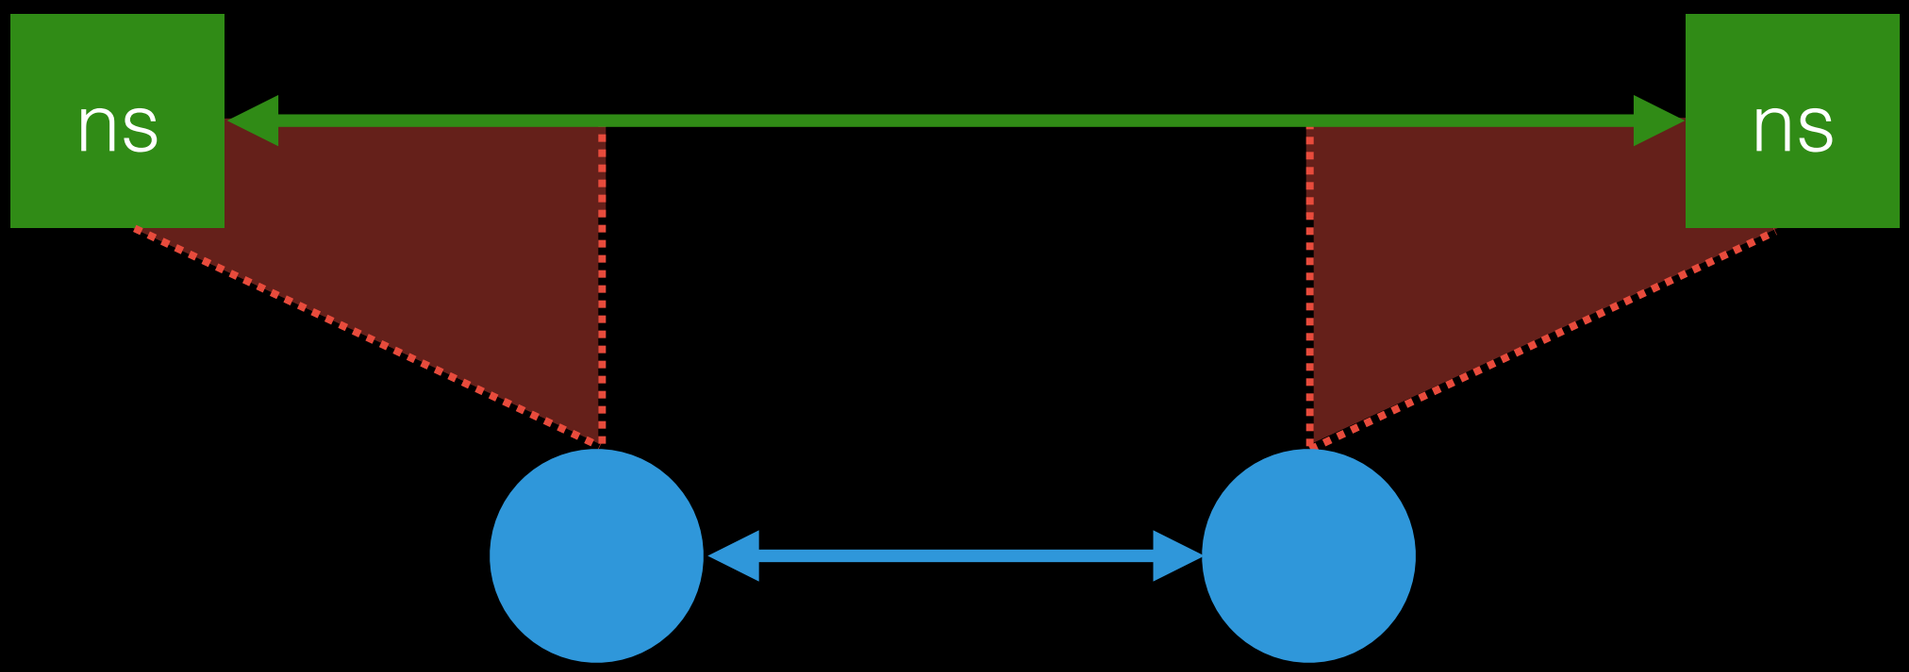

King technique

[Gummadi et al, 2002]

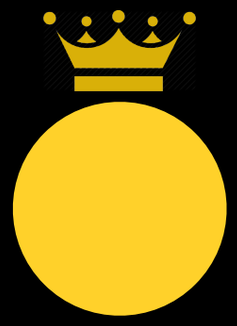

Measure the RTT of
recursive DNS queries

Subtract latencies

King does not *directly* measure
the path between two hosts

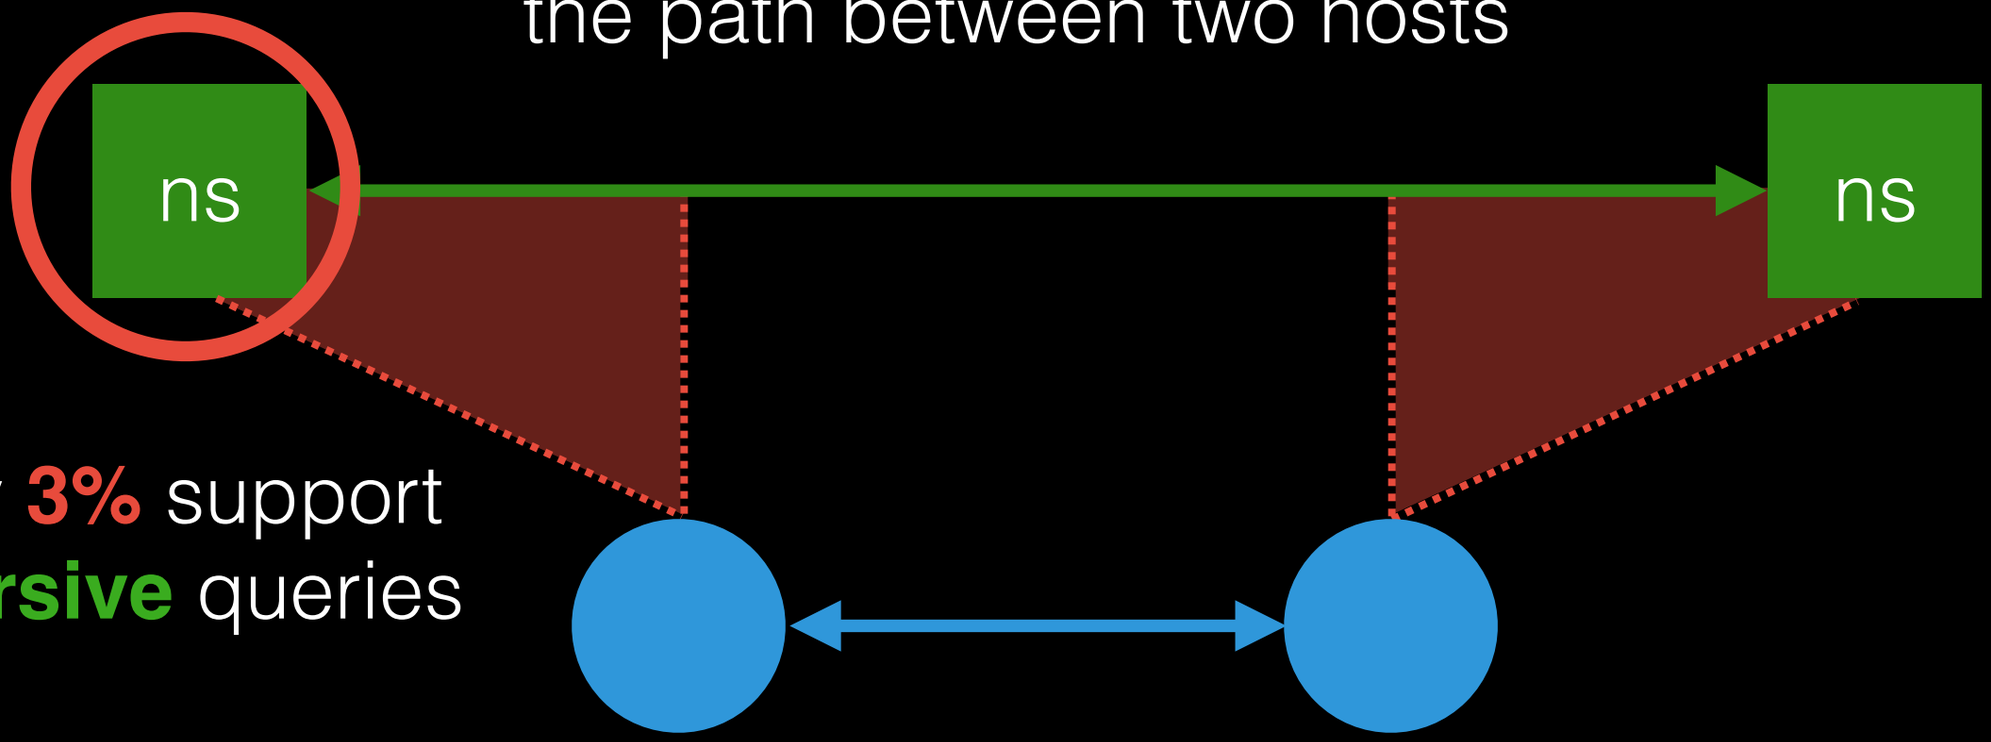

Only **3%** support
recursive queries

TING a tool for measuring latency between arbitrary Tor nodes

1. **Accurate** — measures the full path between end hosts
2. **Practical** — does not require modification of end hosts

What is Tor?

Anonymity-enabling **overlay** network

Packets routed through series of **relays**, called a **circuit**

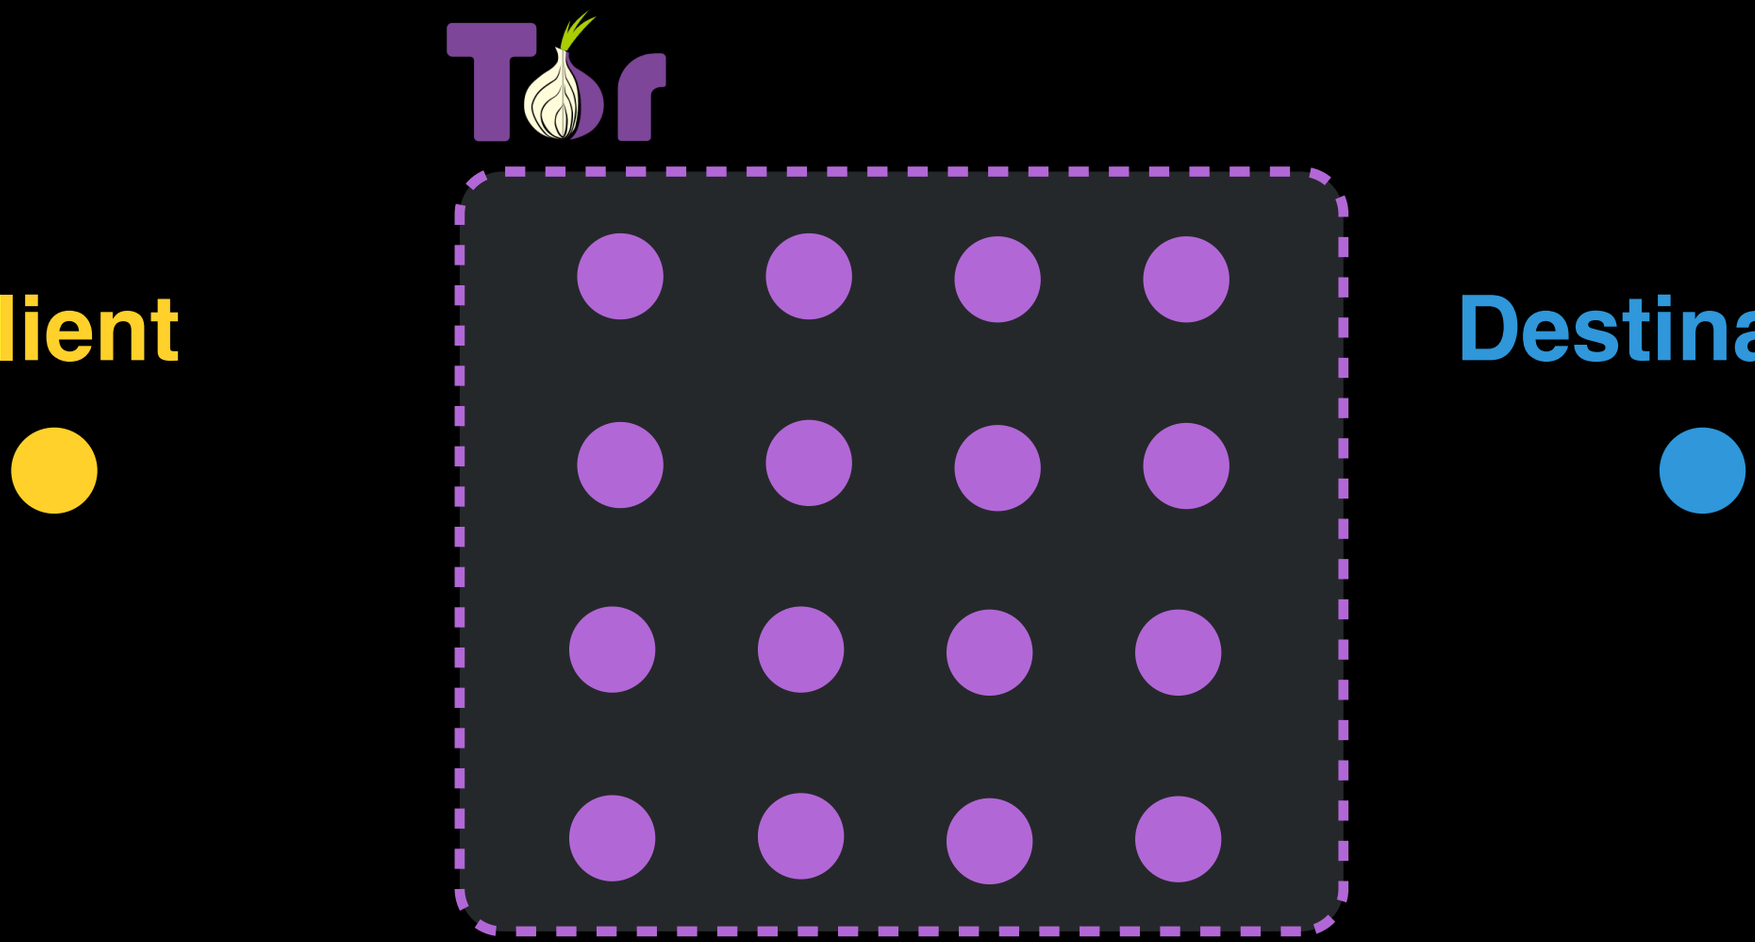

Clients choose their own circuits

What is Tor?

Anonymity-enabling **overlay** network

Packets routed through series of **relays**, called a **circuit**

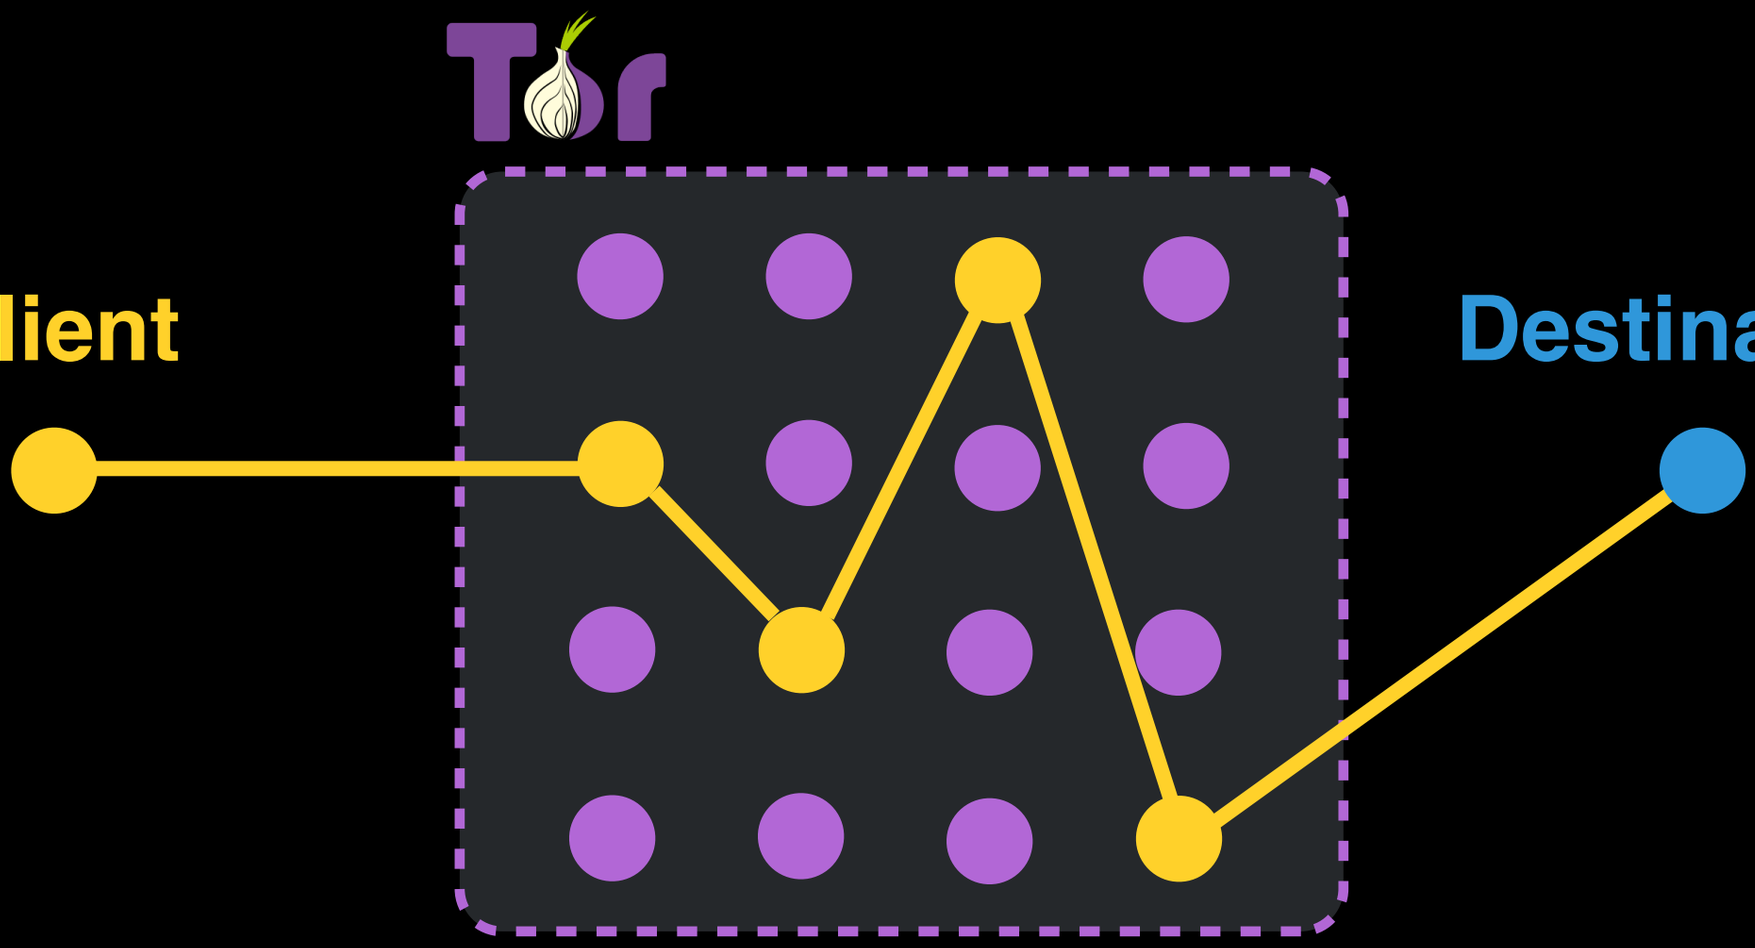

Clients choose their own circuits

Why Tor?

Large — **6,536** relays in **77** countries (on October 29, 2015)

Diverse — **5,520** /24s, **~61%** residential networks

Growing

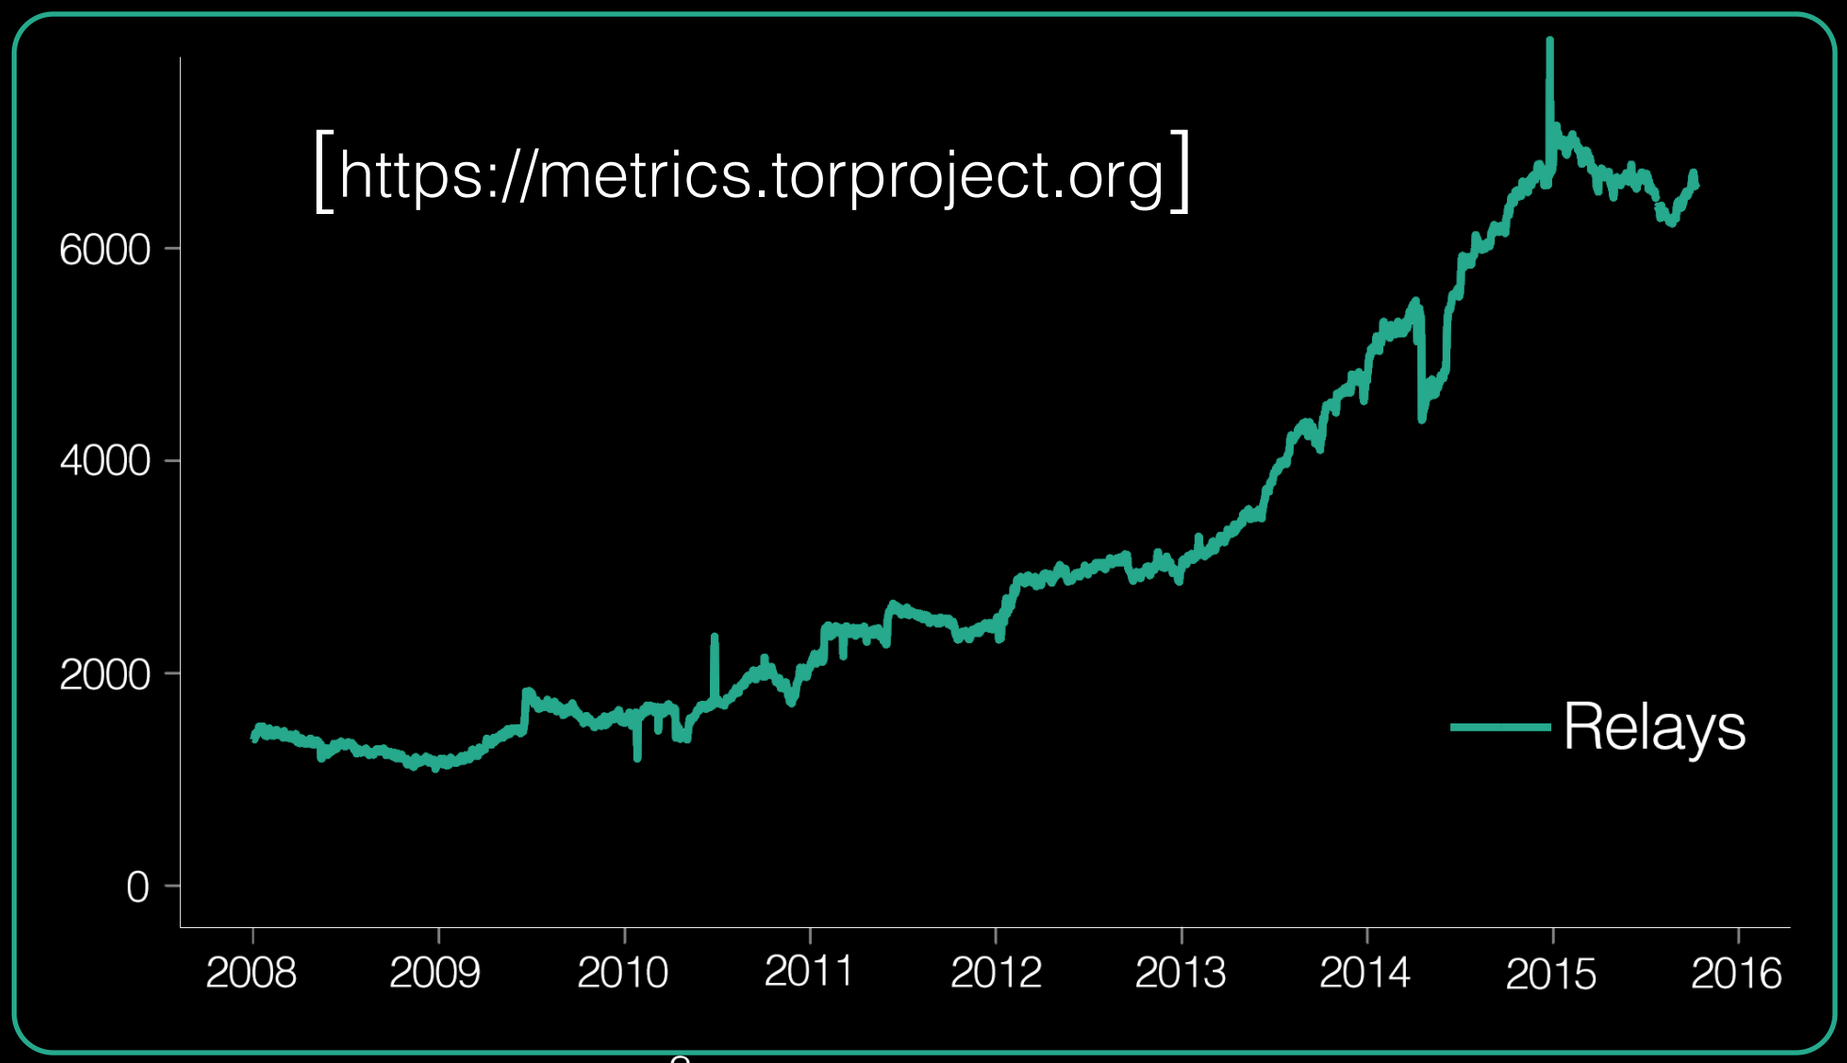

Tor-specific constraints

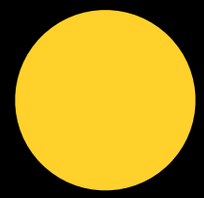

x

y

Tor-specific constraints

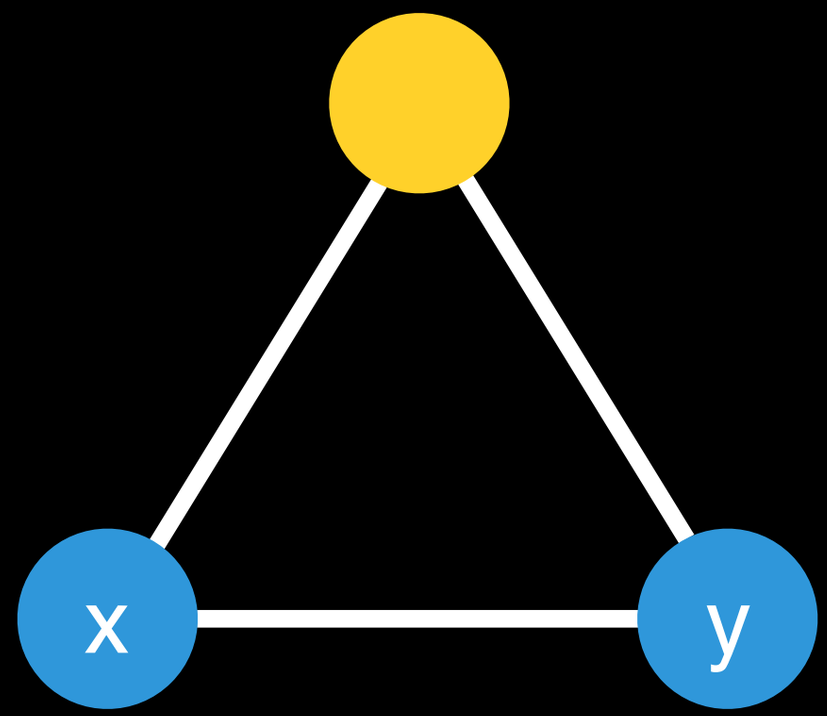

Tor-specific constraints

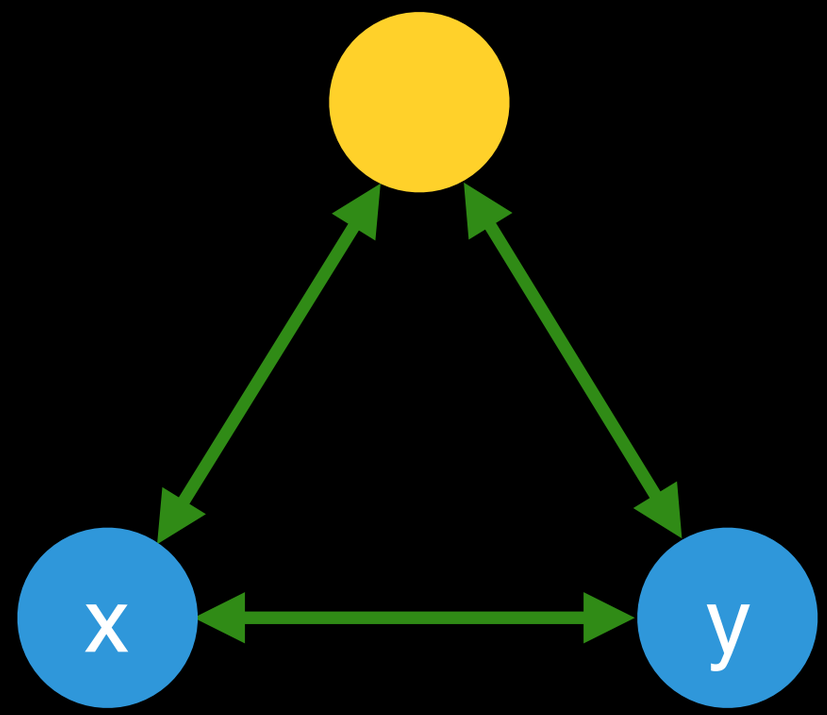

Tor-specific constraints

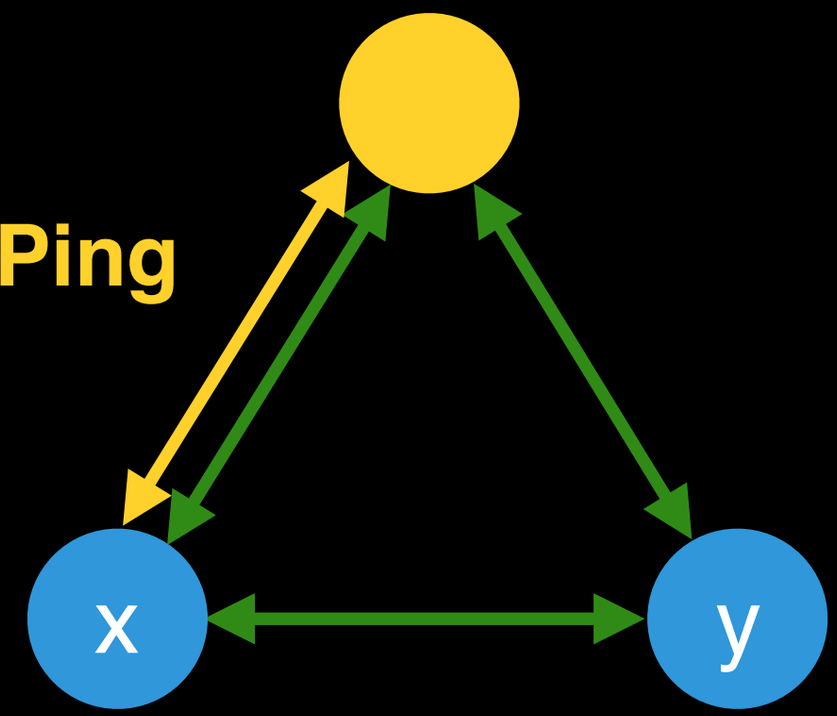

Tor-specific constraints

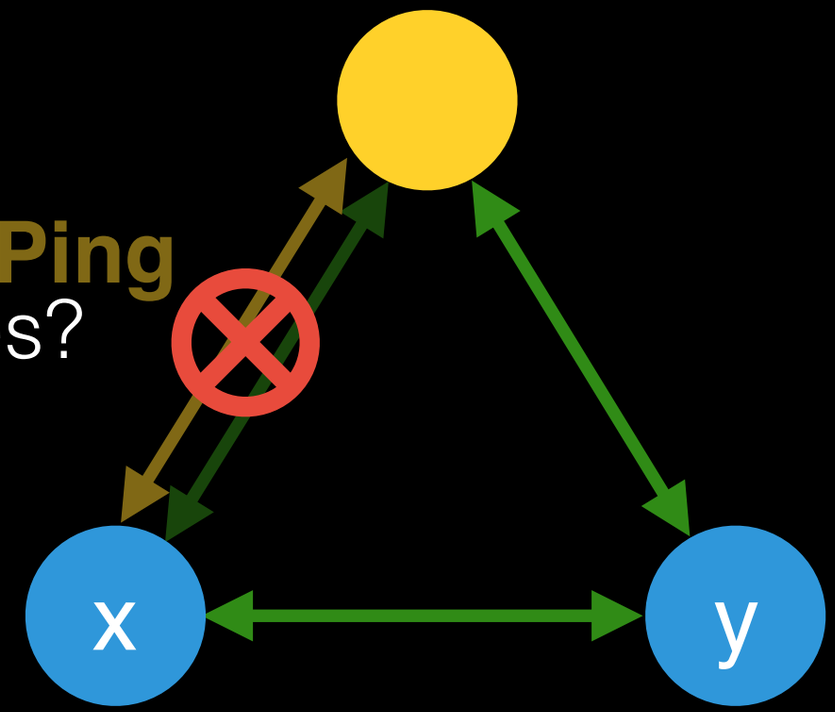

Tor-specific constraints

Tor traffic may be treated differently

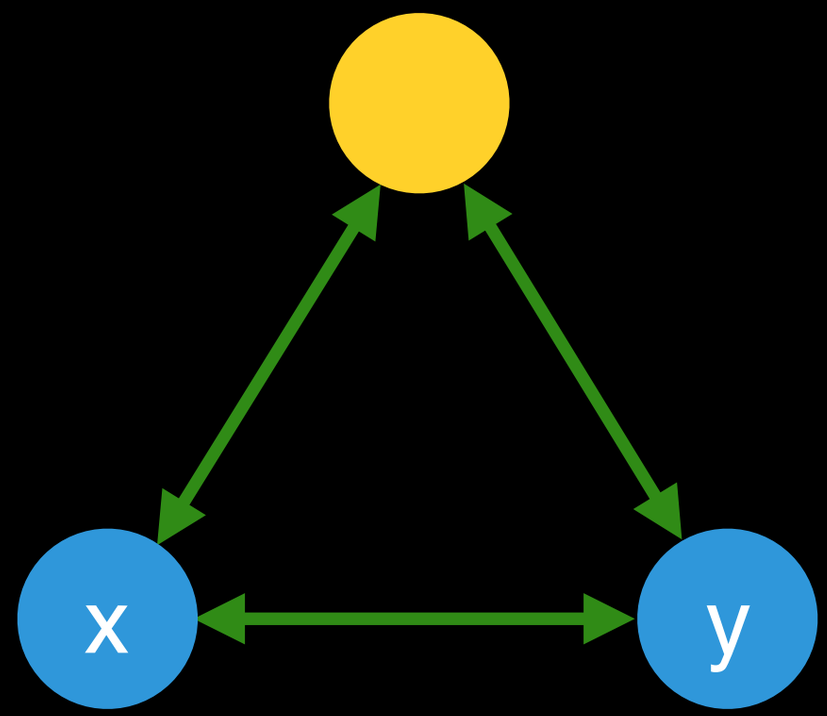

Tor-specific constraints

Tor traffic may be treated differently

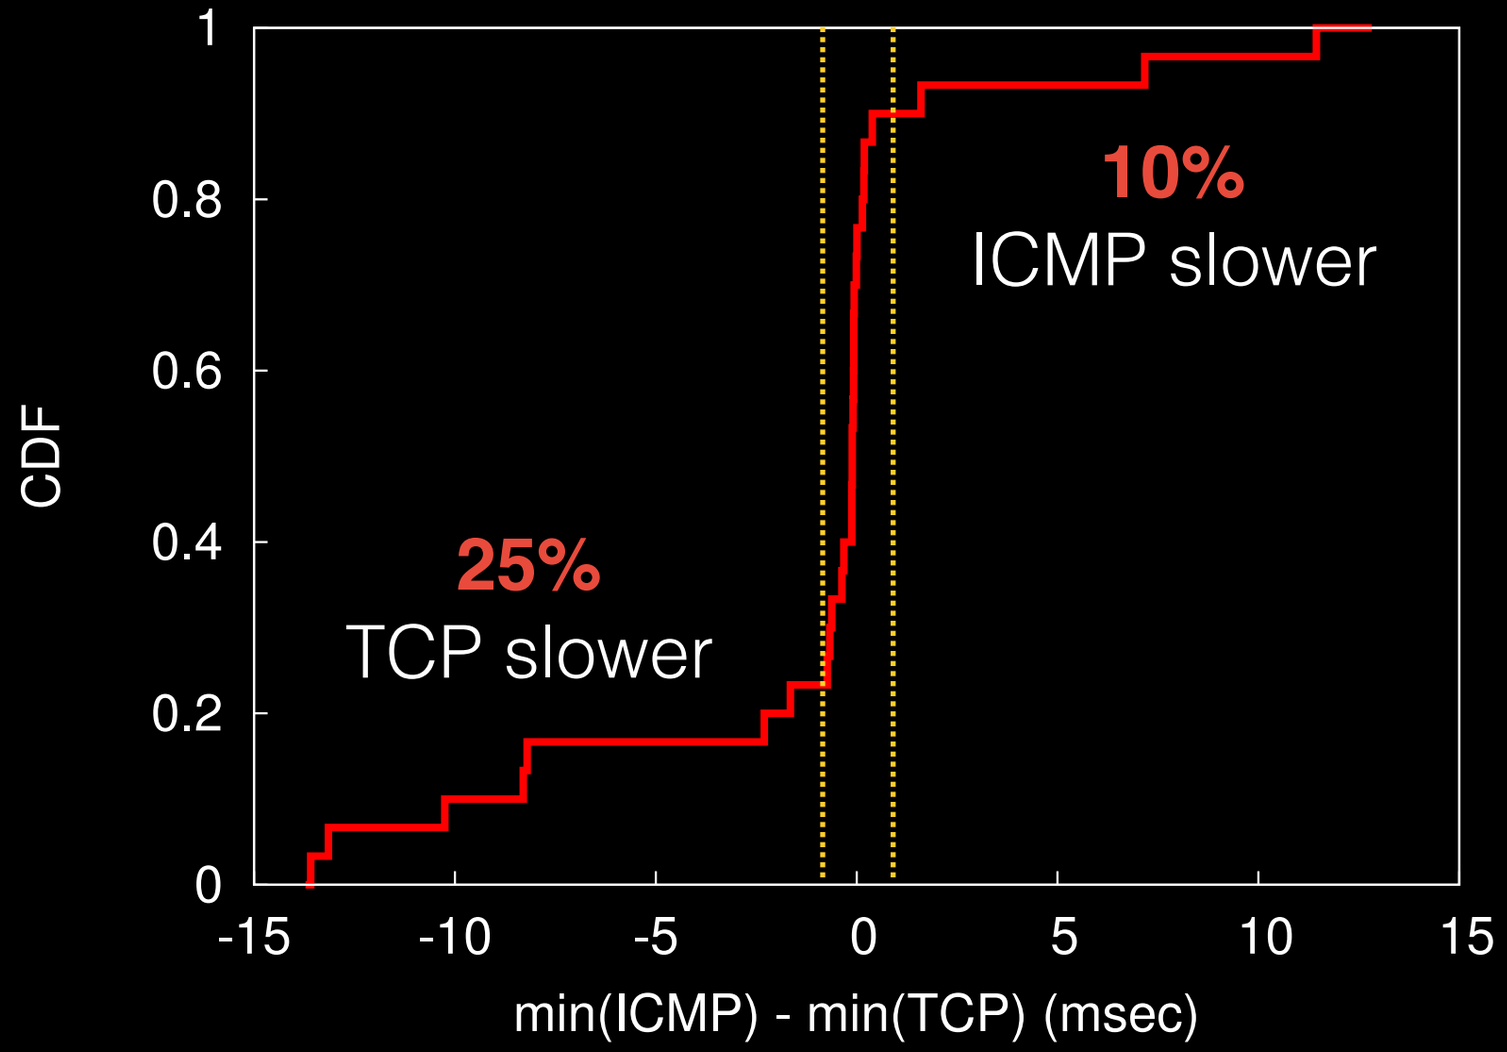

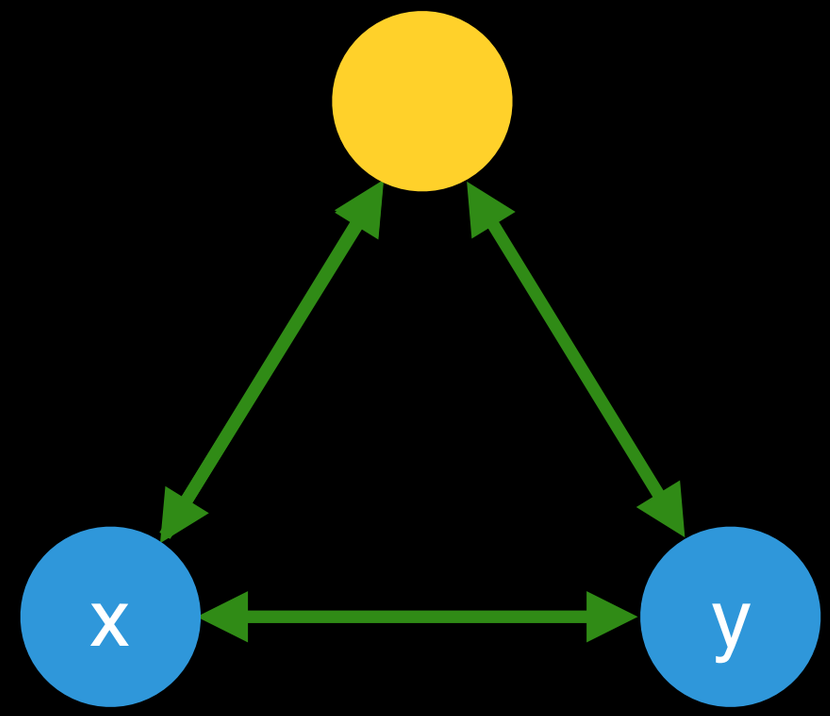

Tor-specific constraints

Tor traffic may be treated differently

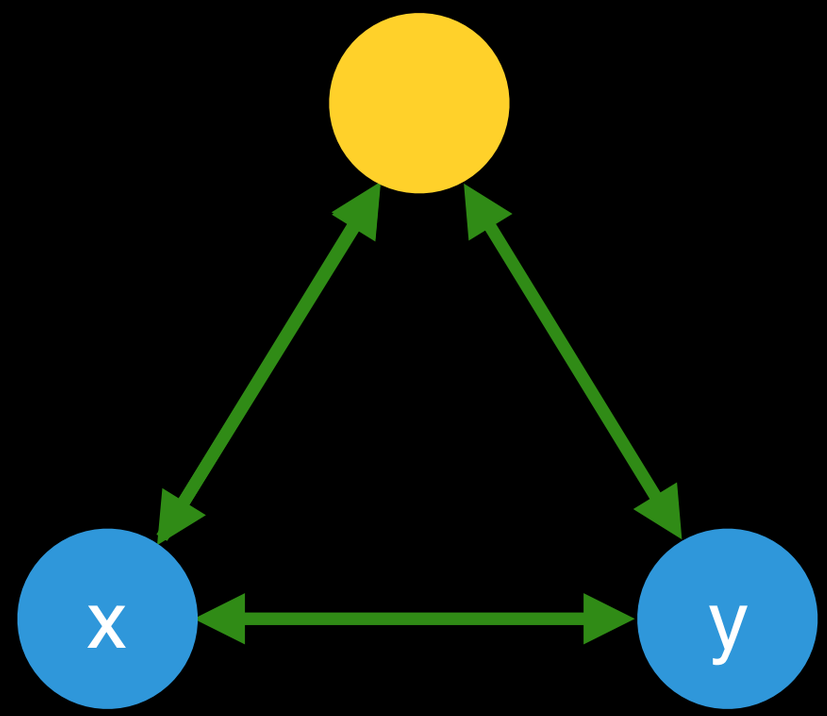

Tor-specific constraints

Tor traffic may be treated differently

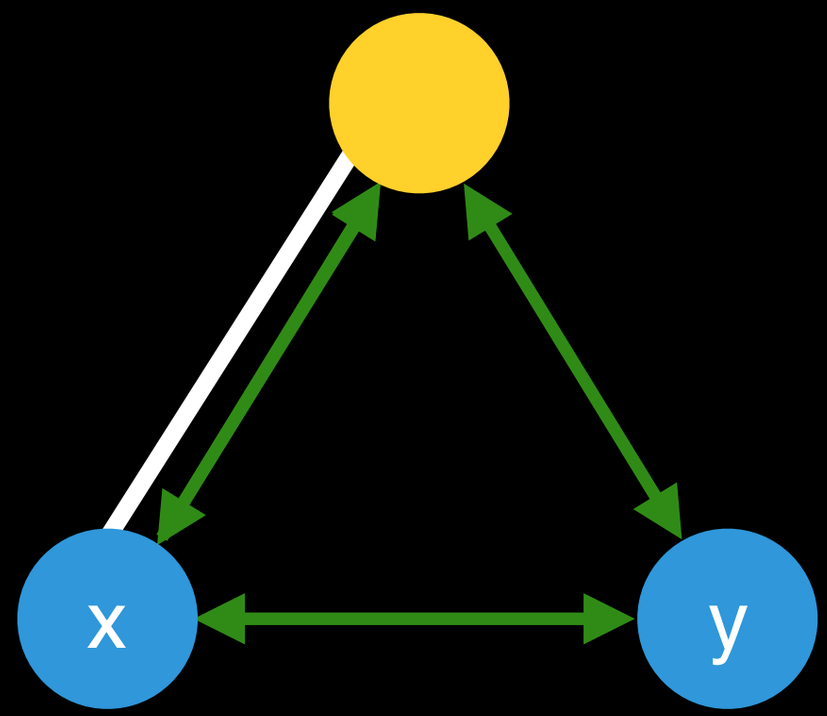

Tor-specific constraints

Tor traffic may be treated differently

Cannot create one-hop circuits

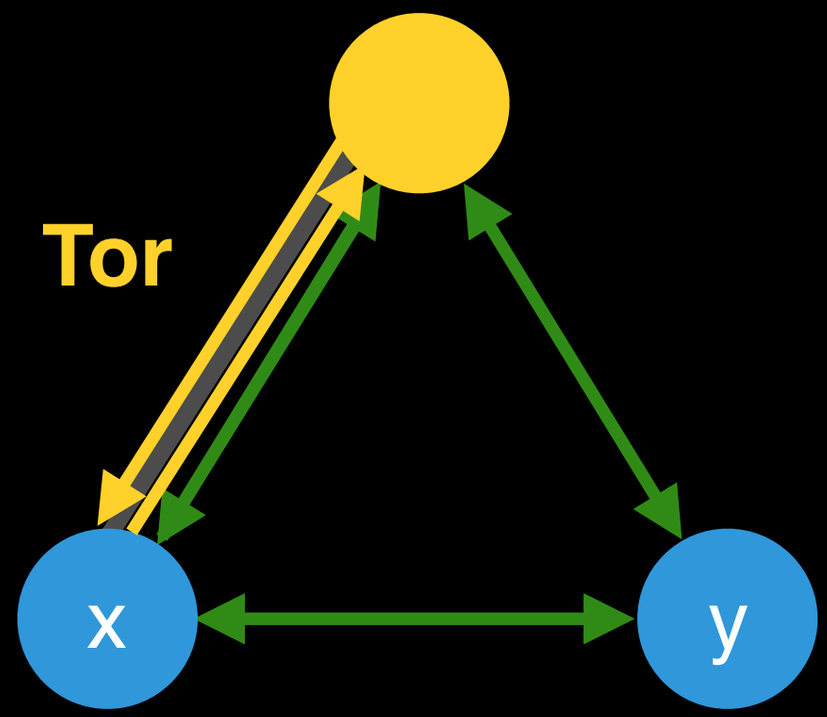

Tor-specific constraints

Tor traffic may be treated differently

Cannot create one-hop circuits

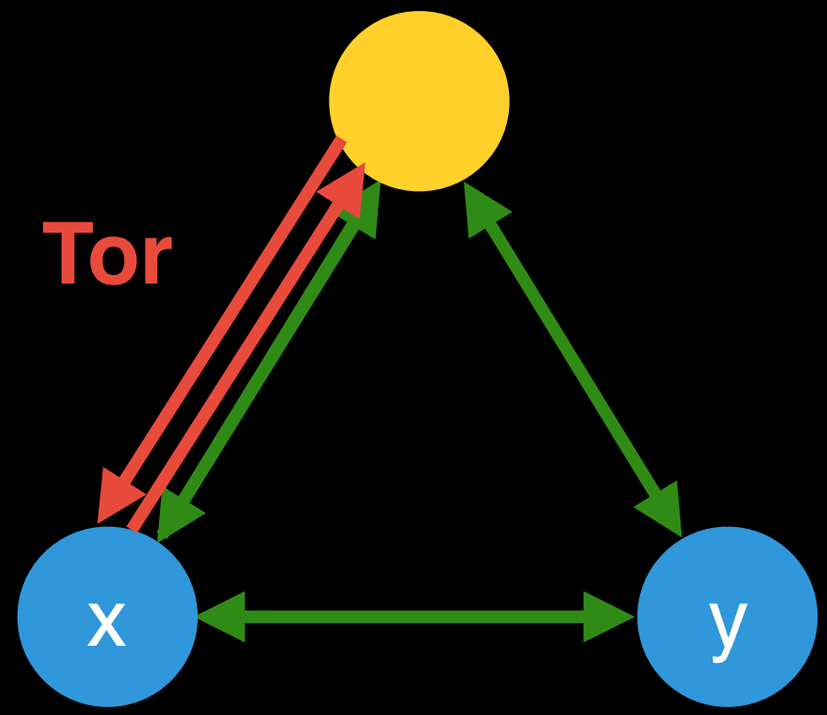

Tor-specific constraints

Tor traffic may be treated differently

Cannot create one-hop circuits

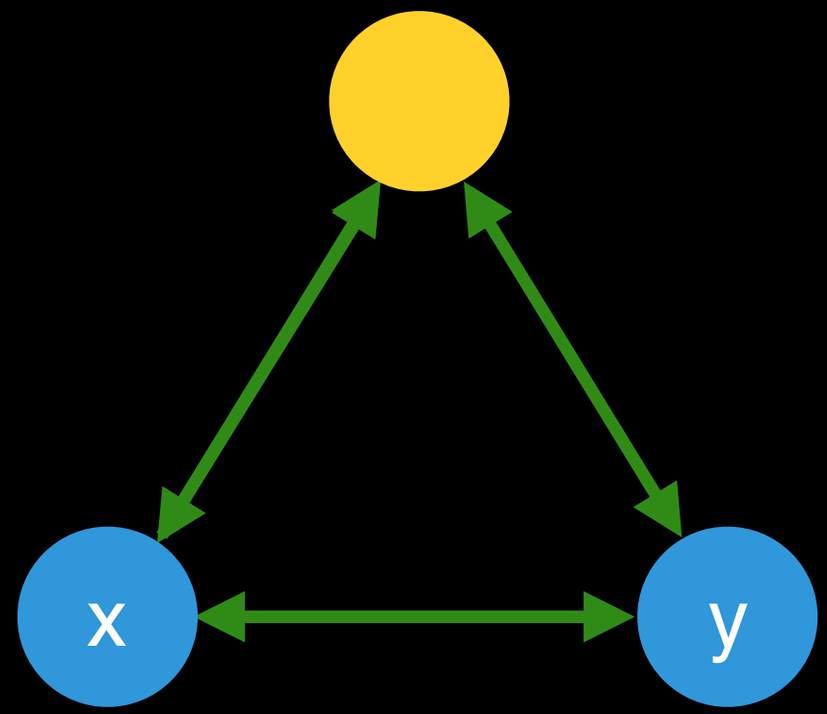

Tor-specific constraints

Tor traffic may be treated differently

Cannot create one-hop circuits

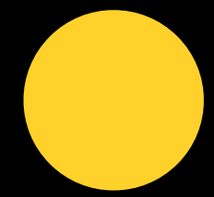

Tor-specific constraints

Tor traffic may be treated differently

Cannot create one-hop circuits

Must account for forwarding delays

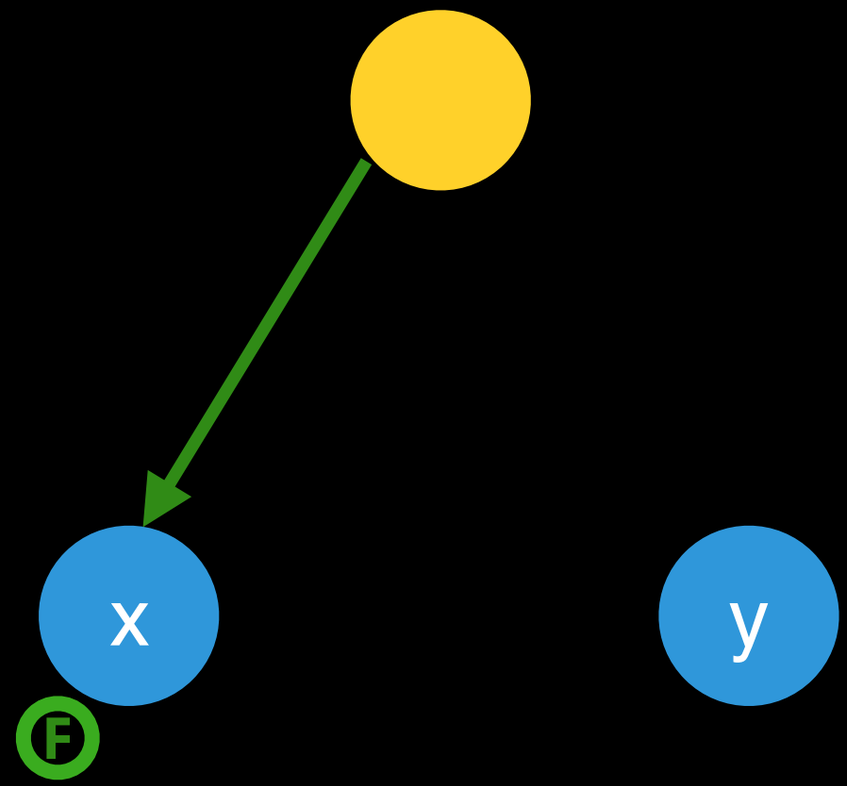

Tor-specific constraints

Tor traffic may be treated differently

Cannot create one-hop circuits

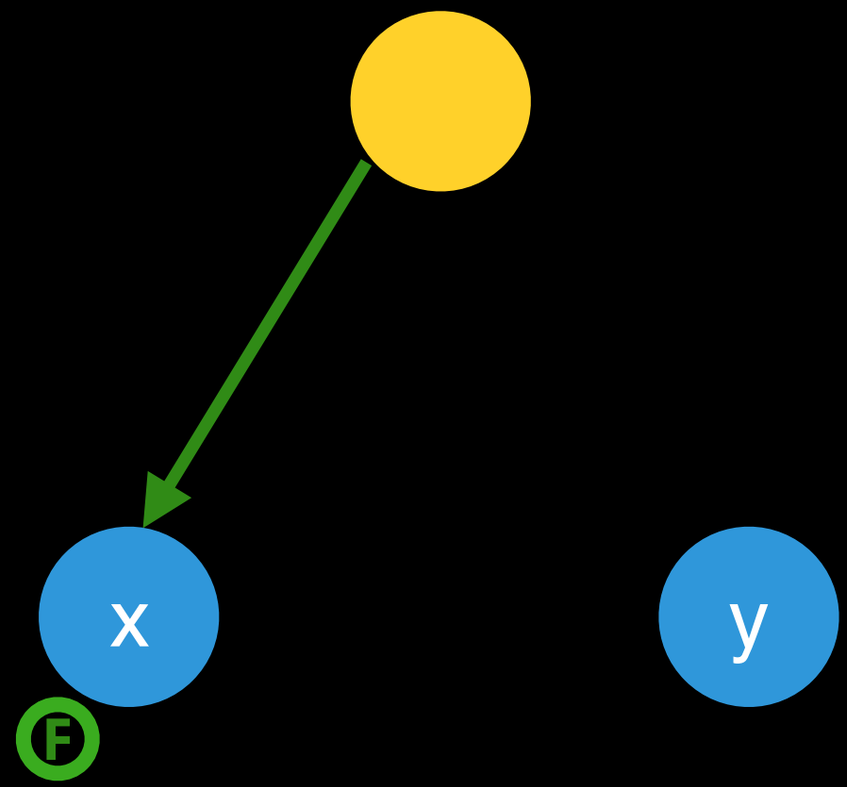

Must account for forwarding delays

- Queuing / Scheduling
- Encryption & Decryption
- Context Switches

Tor-specific constraints

Tor traffic may be treated differently

Cannot create one-hop circuits

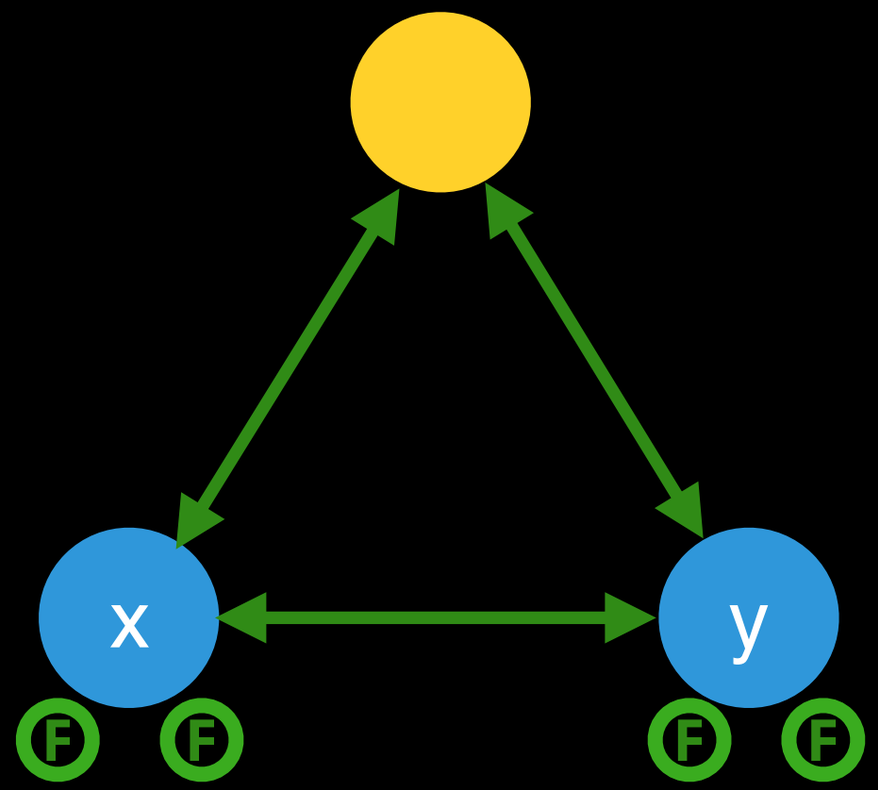

Must account for forwarding delays

- Queuing / Scheduling
- Encryption & Decryption
- Context Switches

Tor-specific constraints

Tor traffic may be treated differently

Cannot create one-hop circuits

Must account for forwarding delays

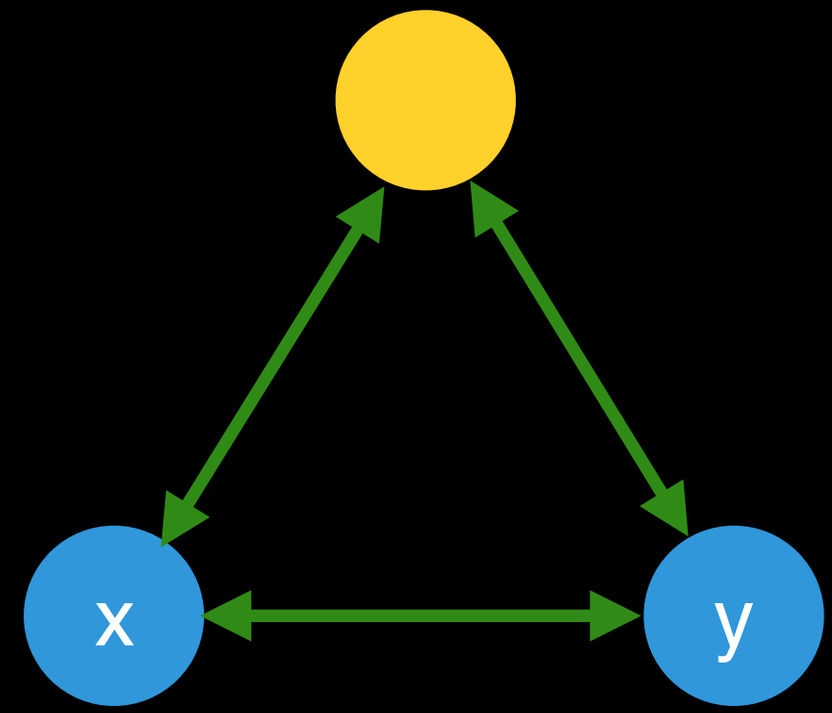

Ting technique

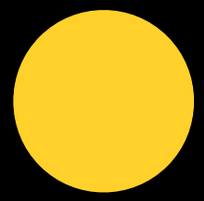

x

y

Measurement Host

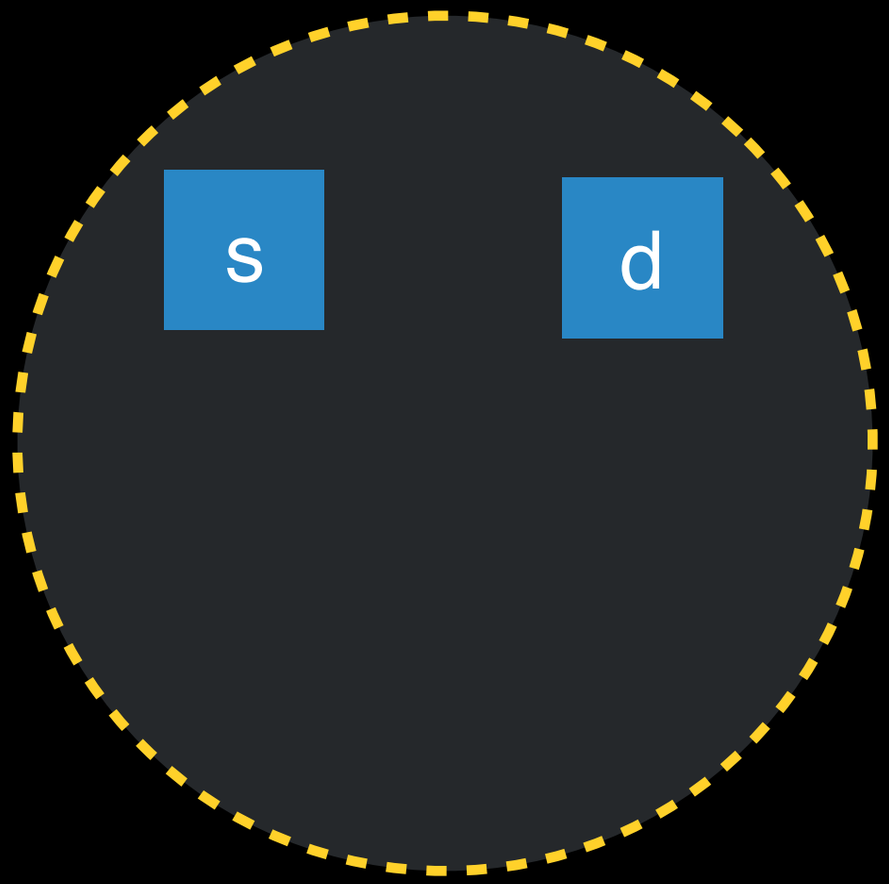

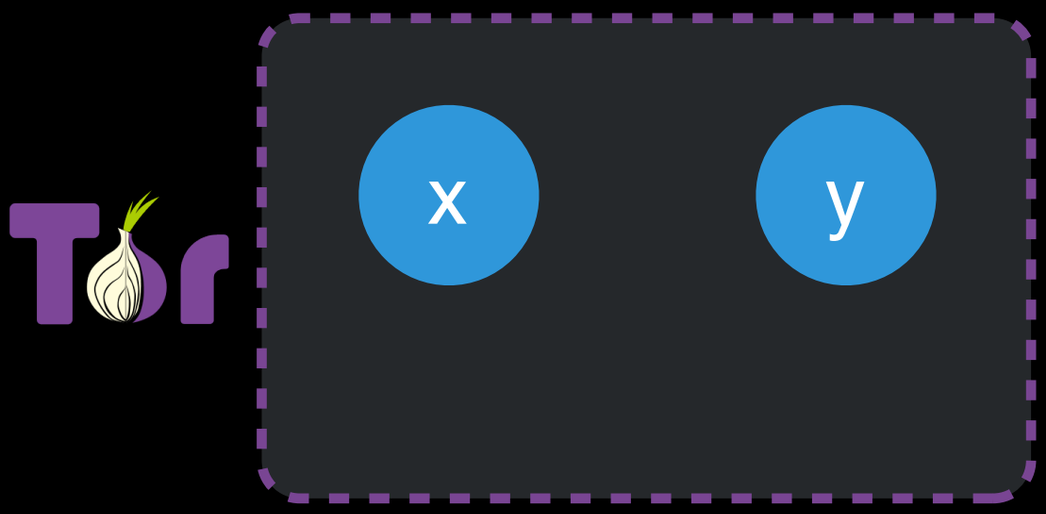

1 2 3

Summary

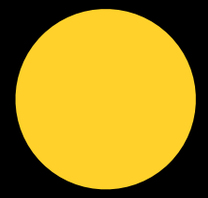

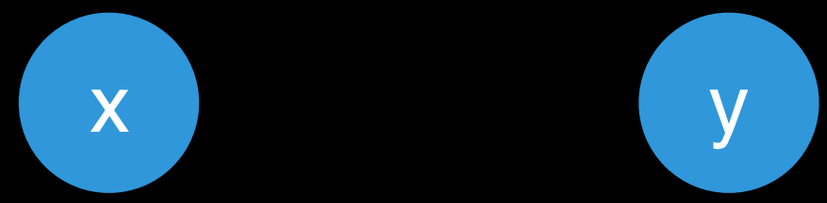

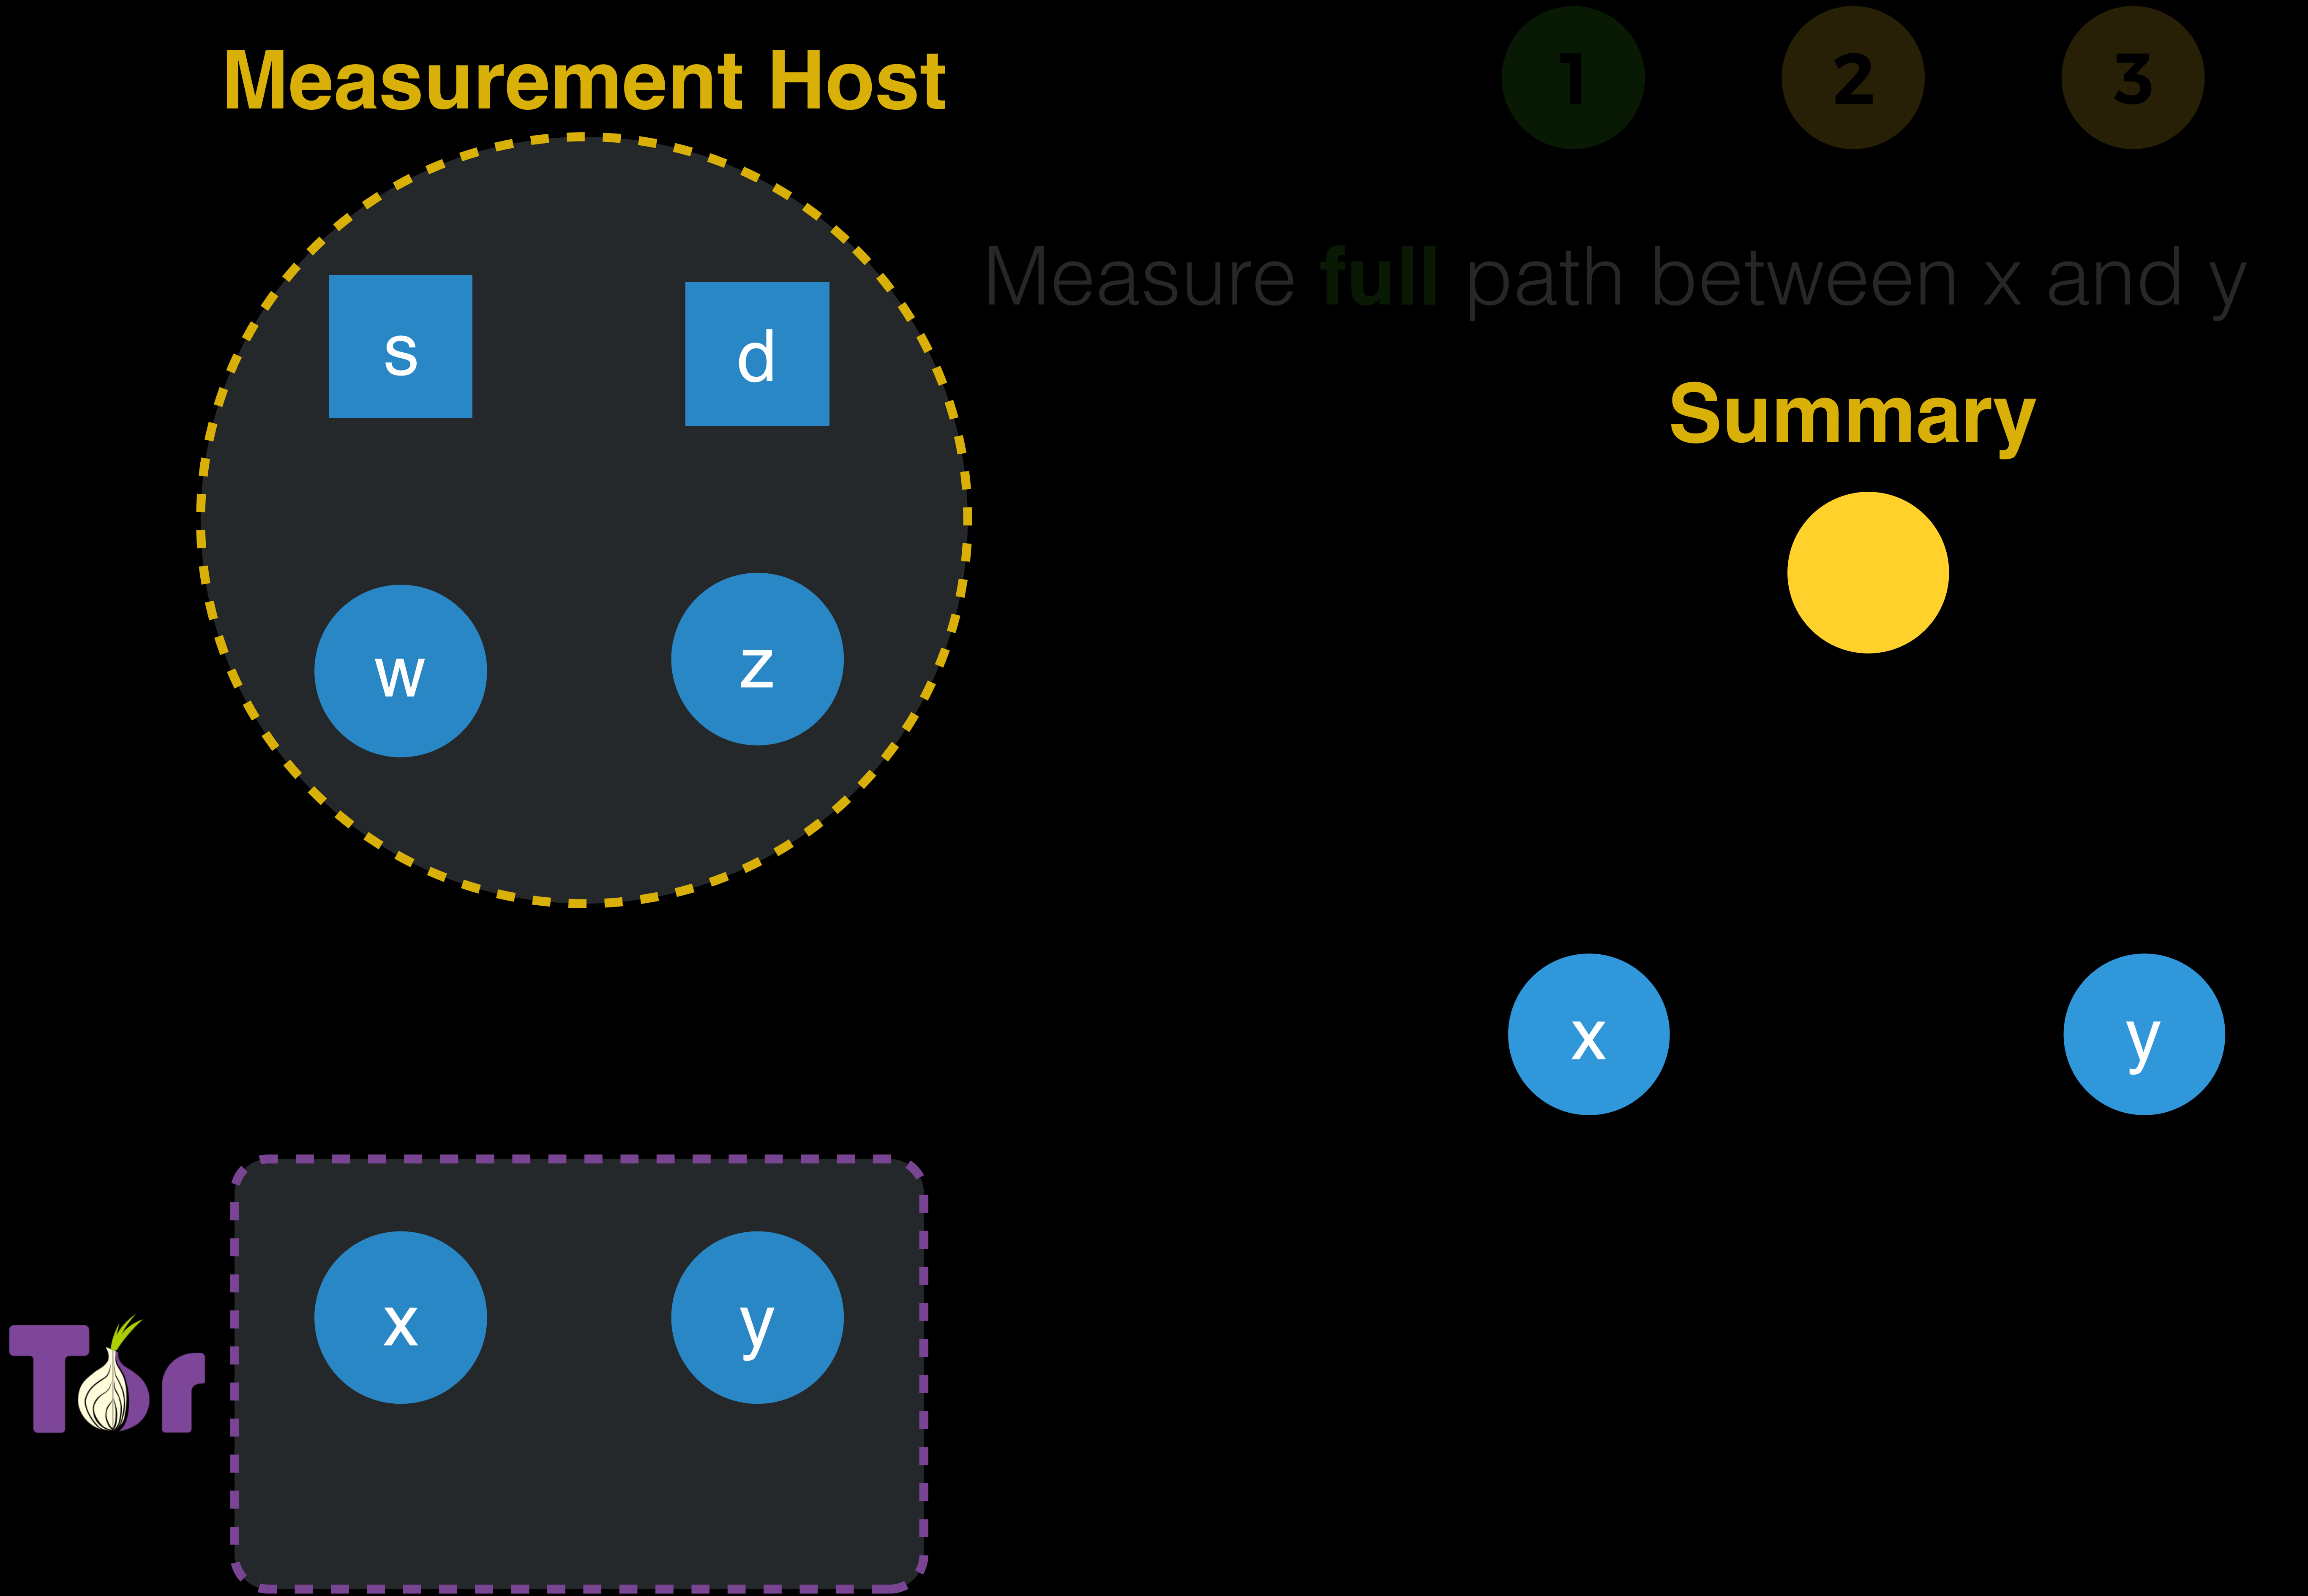
Measurement Host
1
2
3
Measure full path between x and y
Summary
s
d
w
z
x
y
x
y

Measurement Host

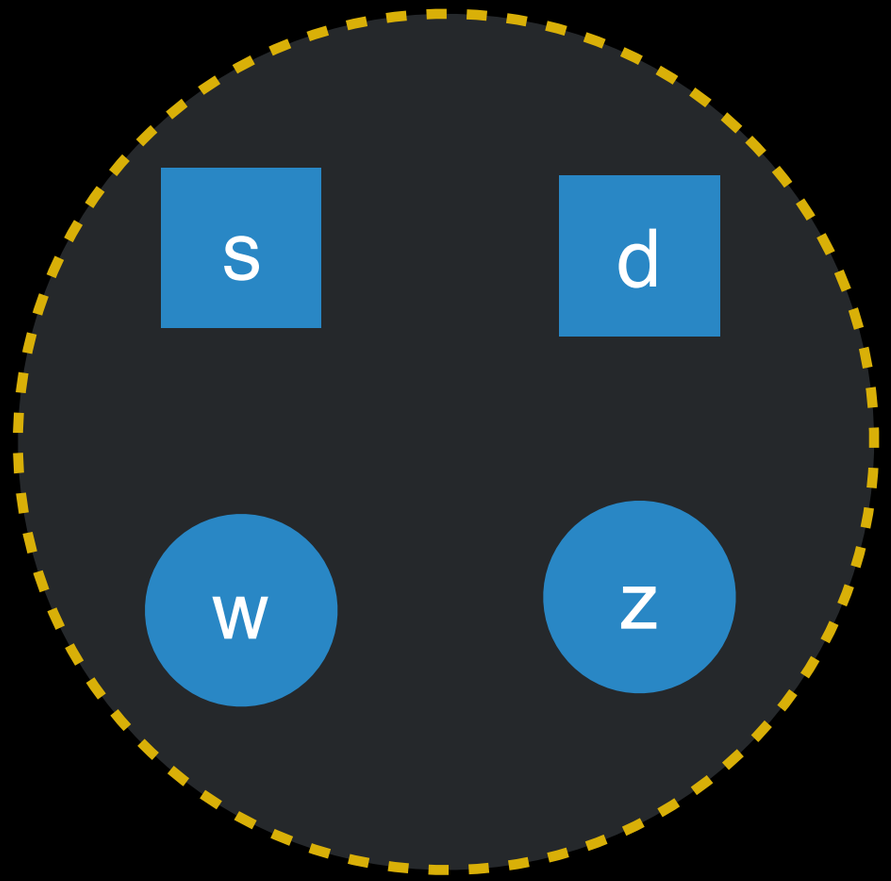

1 2 3

Measure **full** path between x and y

Summary

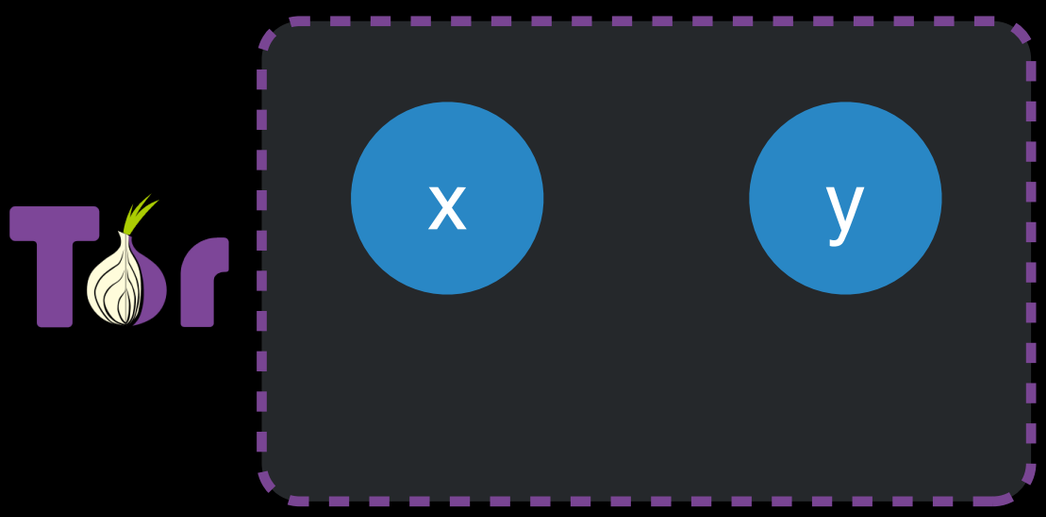

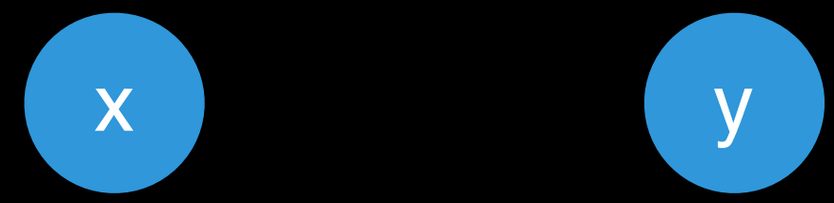

Measurement Host

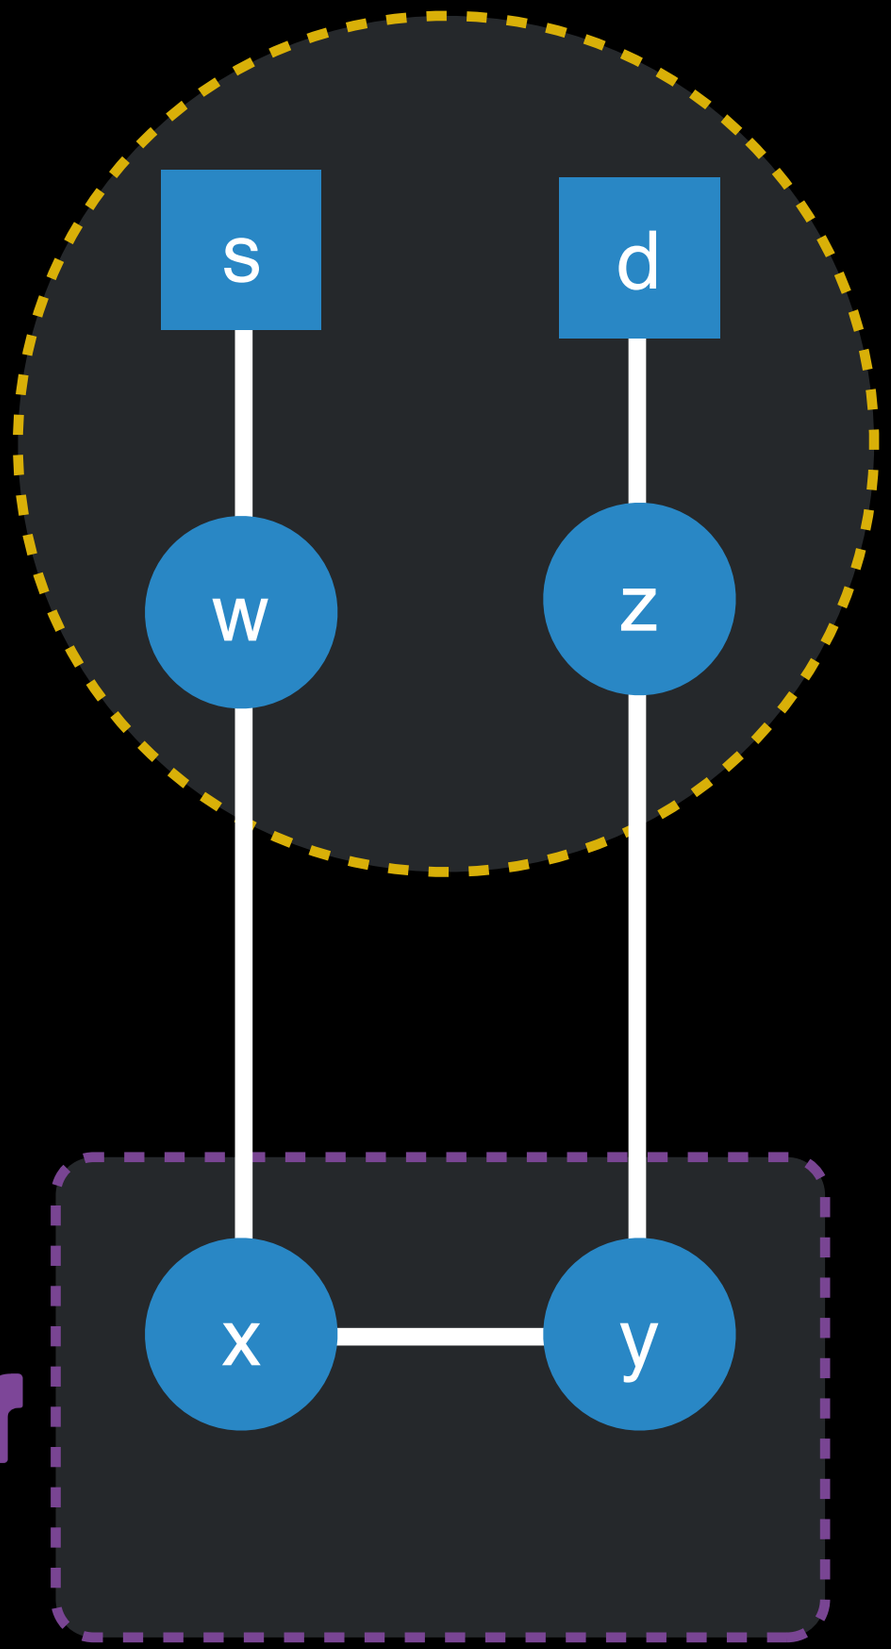

1 2 3

Measure **full** path between x and y

Summary

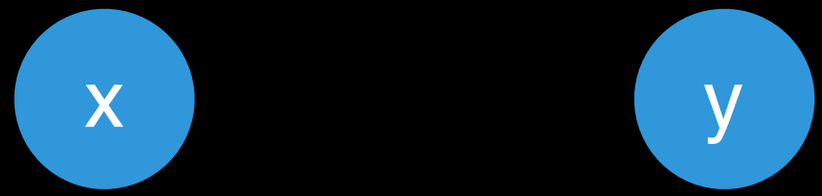

Measurement Host

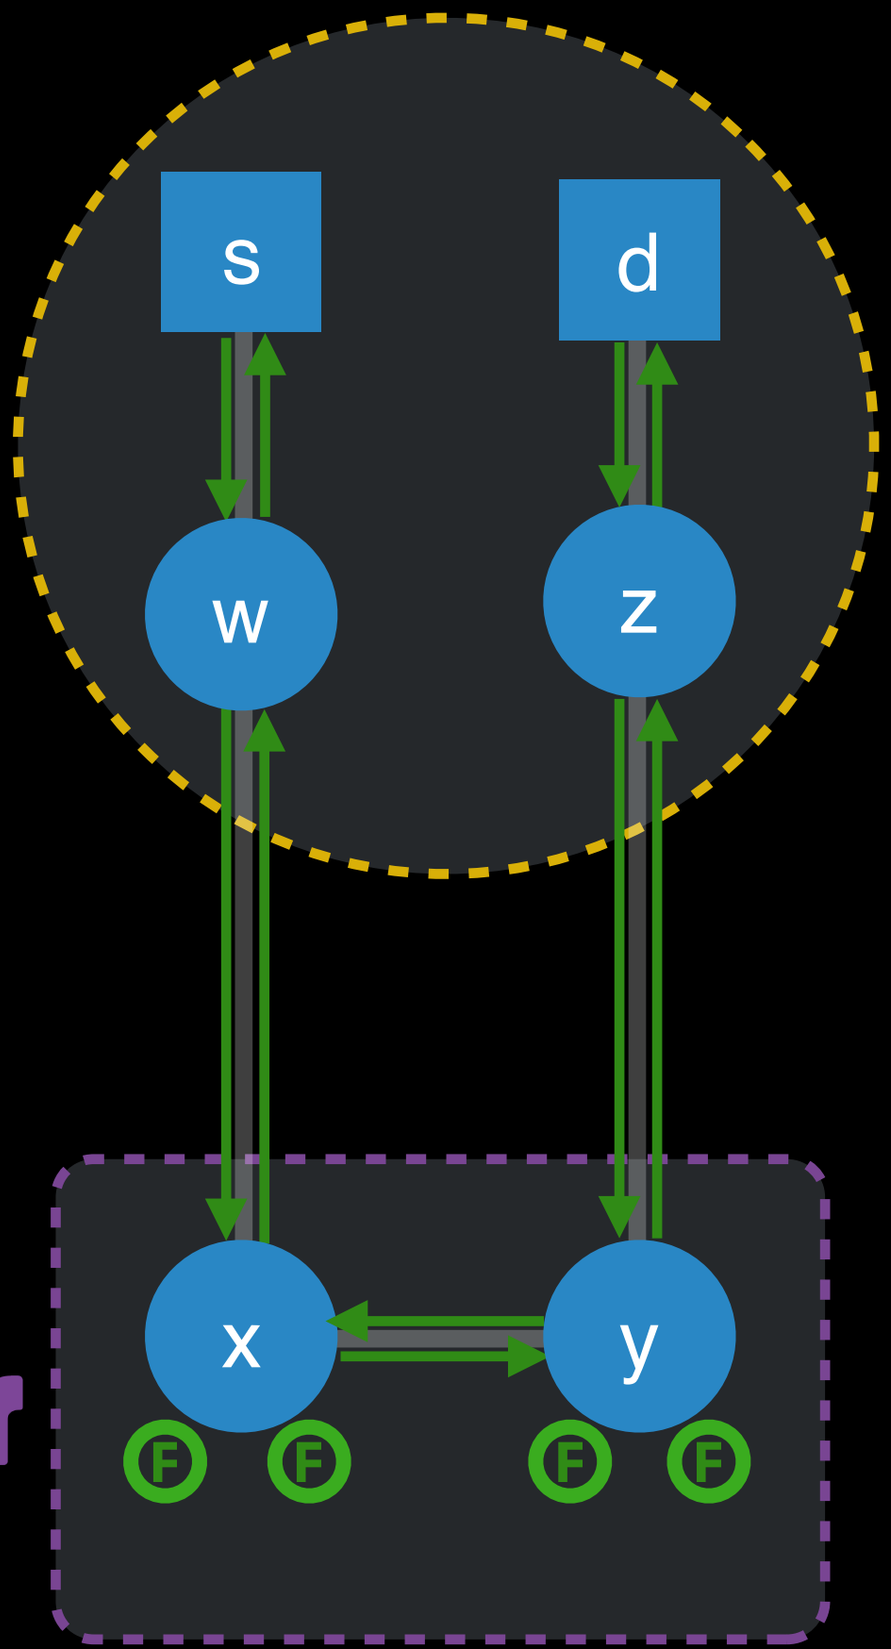

1 2 3

Measure **full** path between x and y

Summary

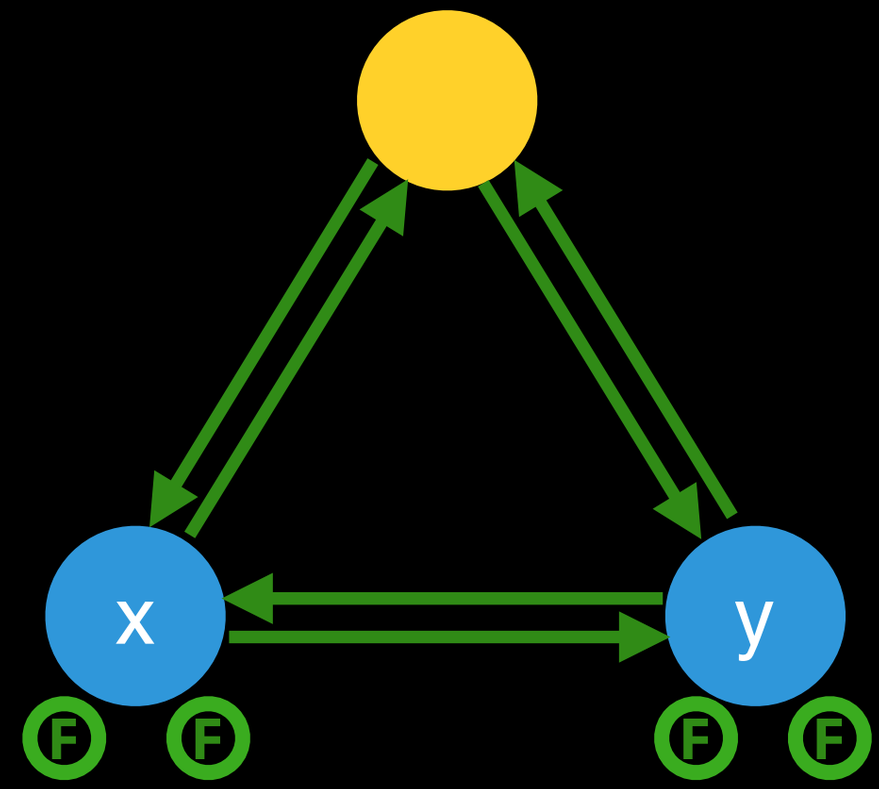

Measurement Host

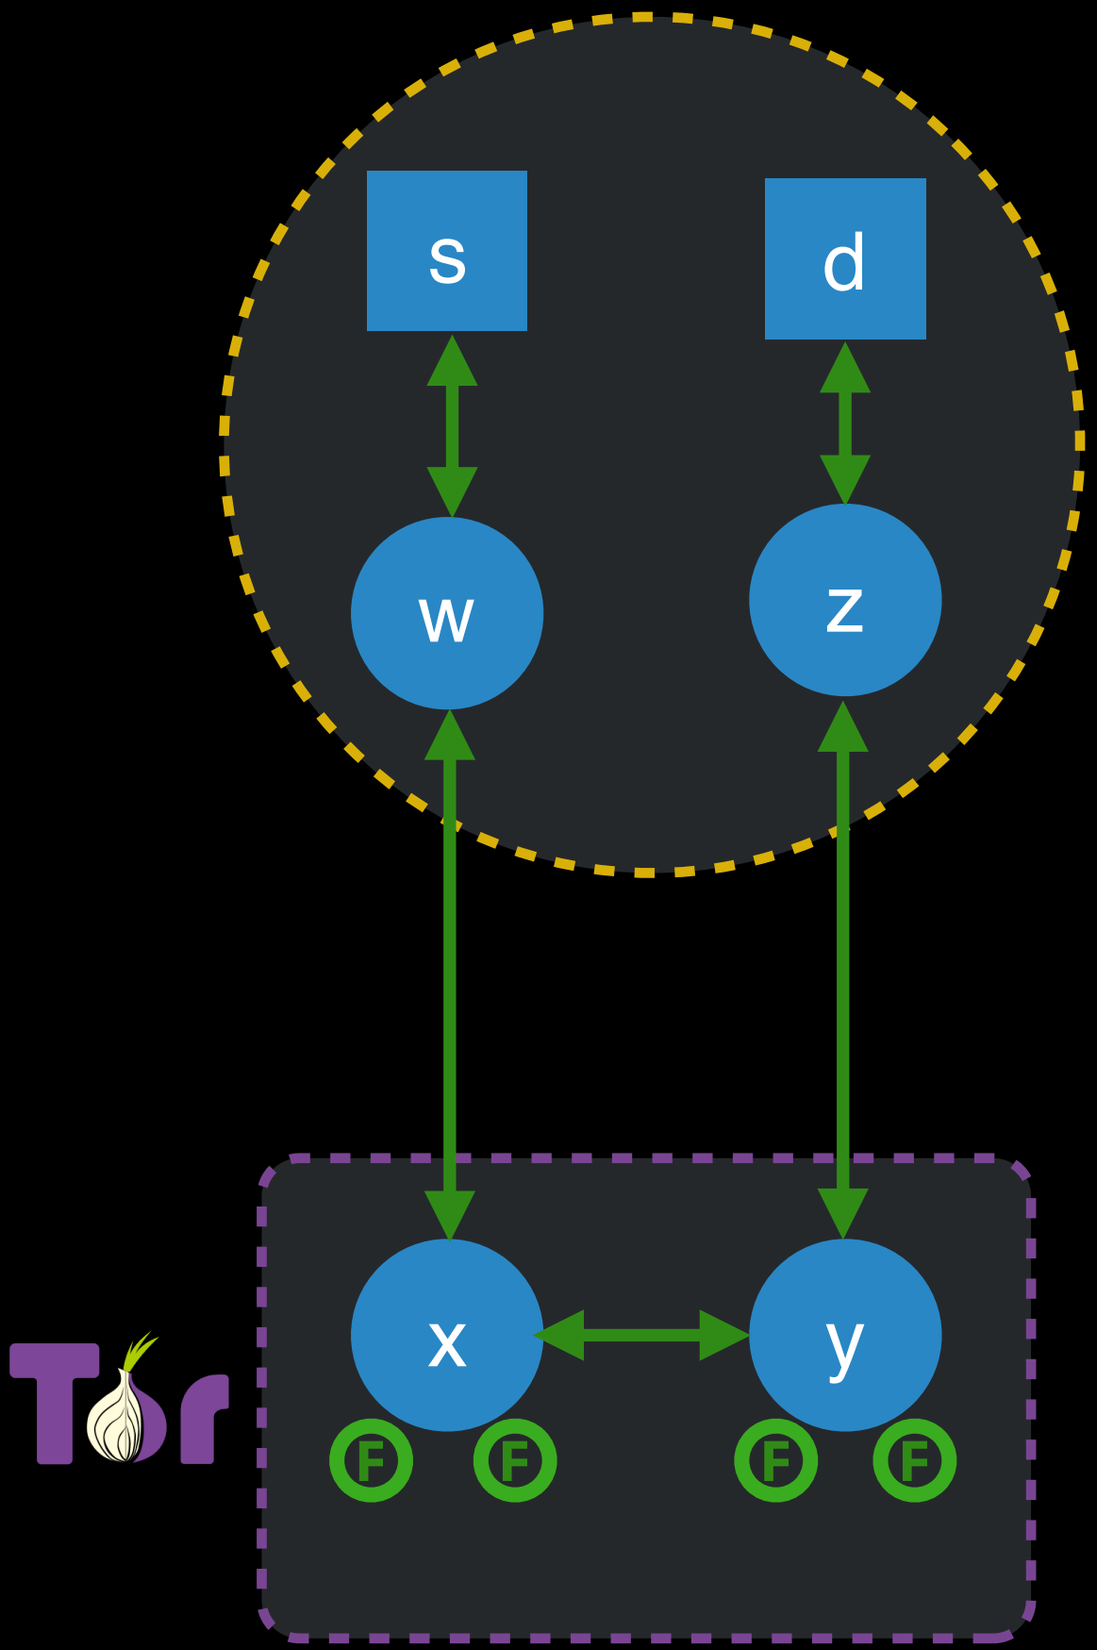

1 2 3

Measure **full** path between x and y

Summary

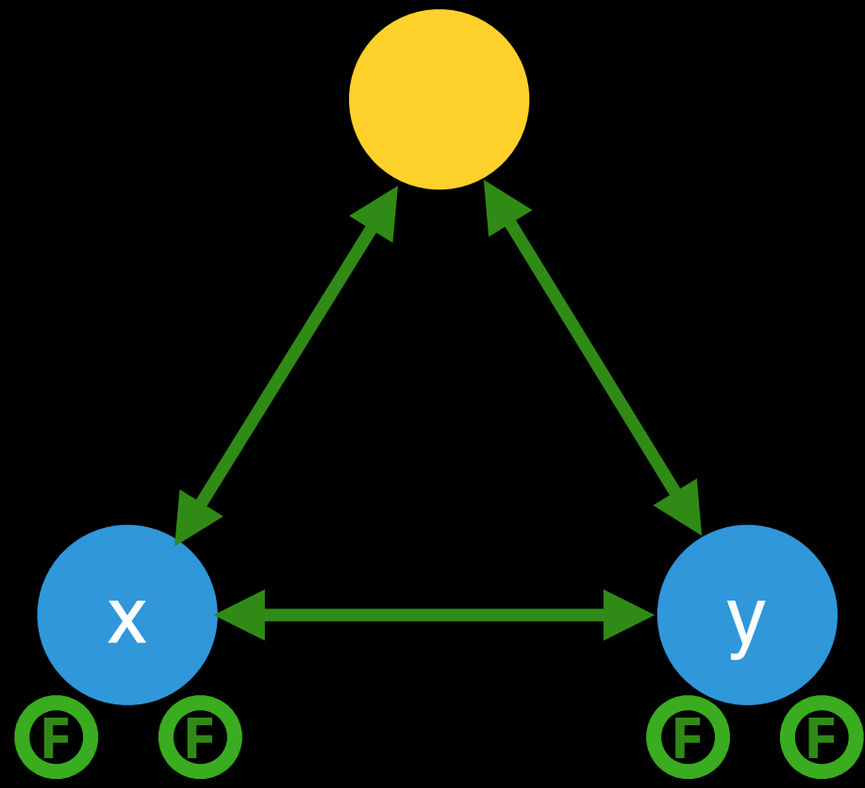

Measurement Host

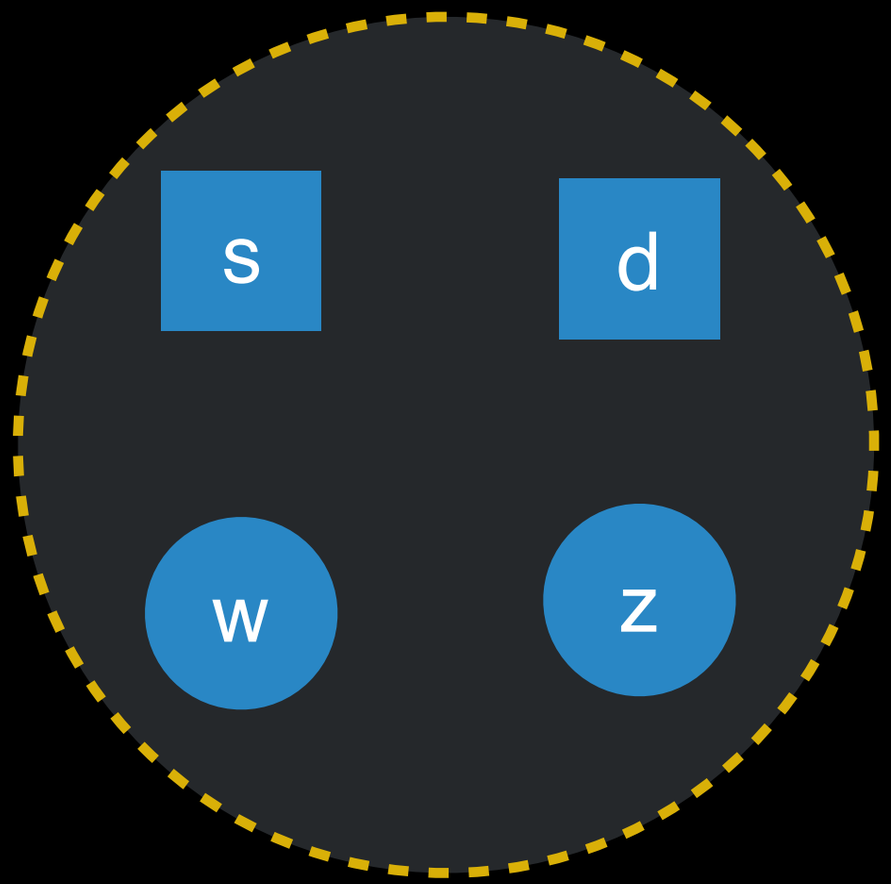

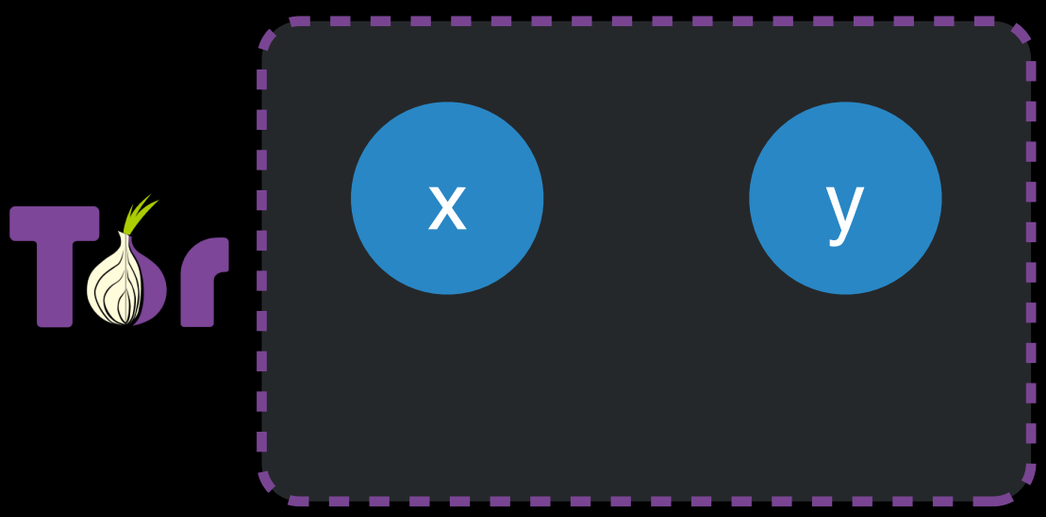

1 2 3

Isolate **RTT** between client and x

Summary

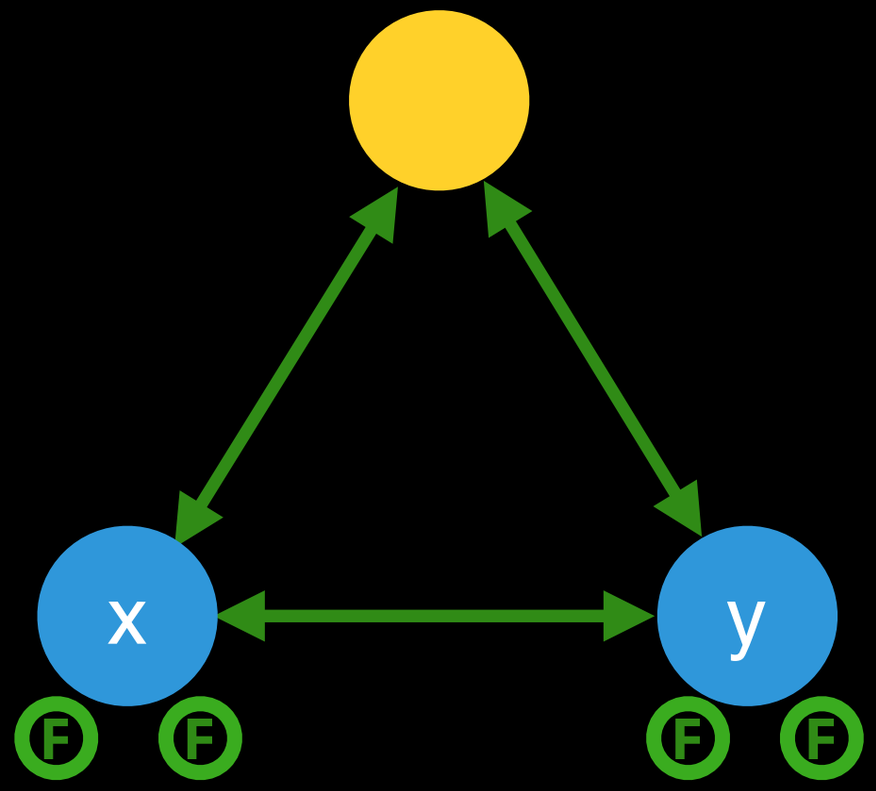

Measurement Host

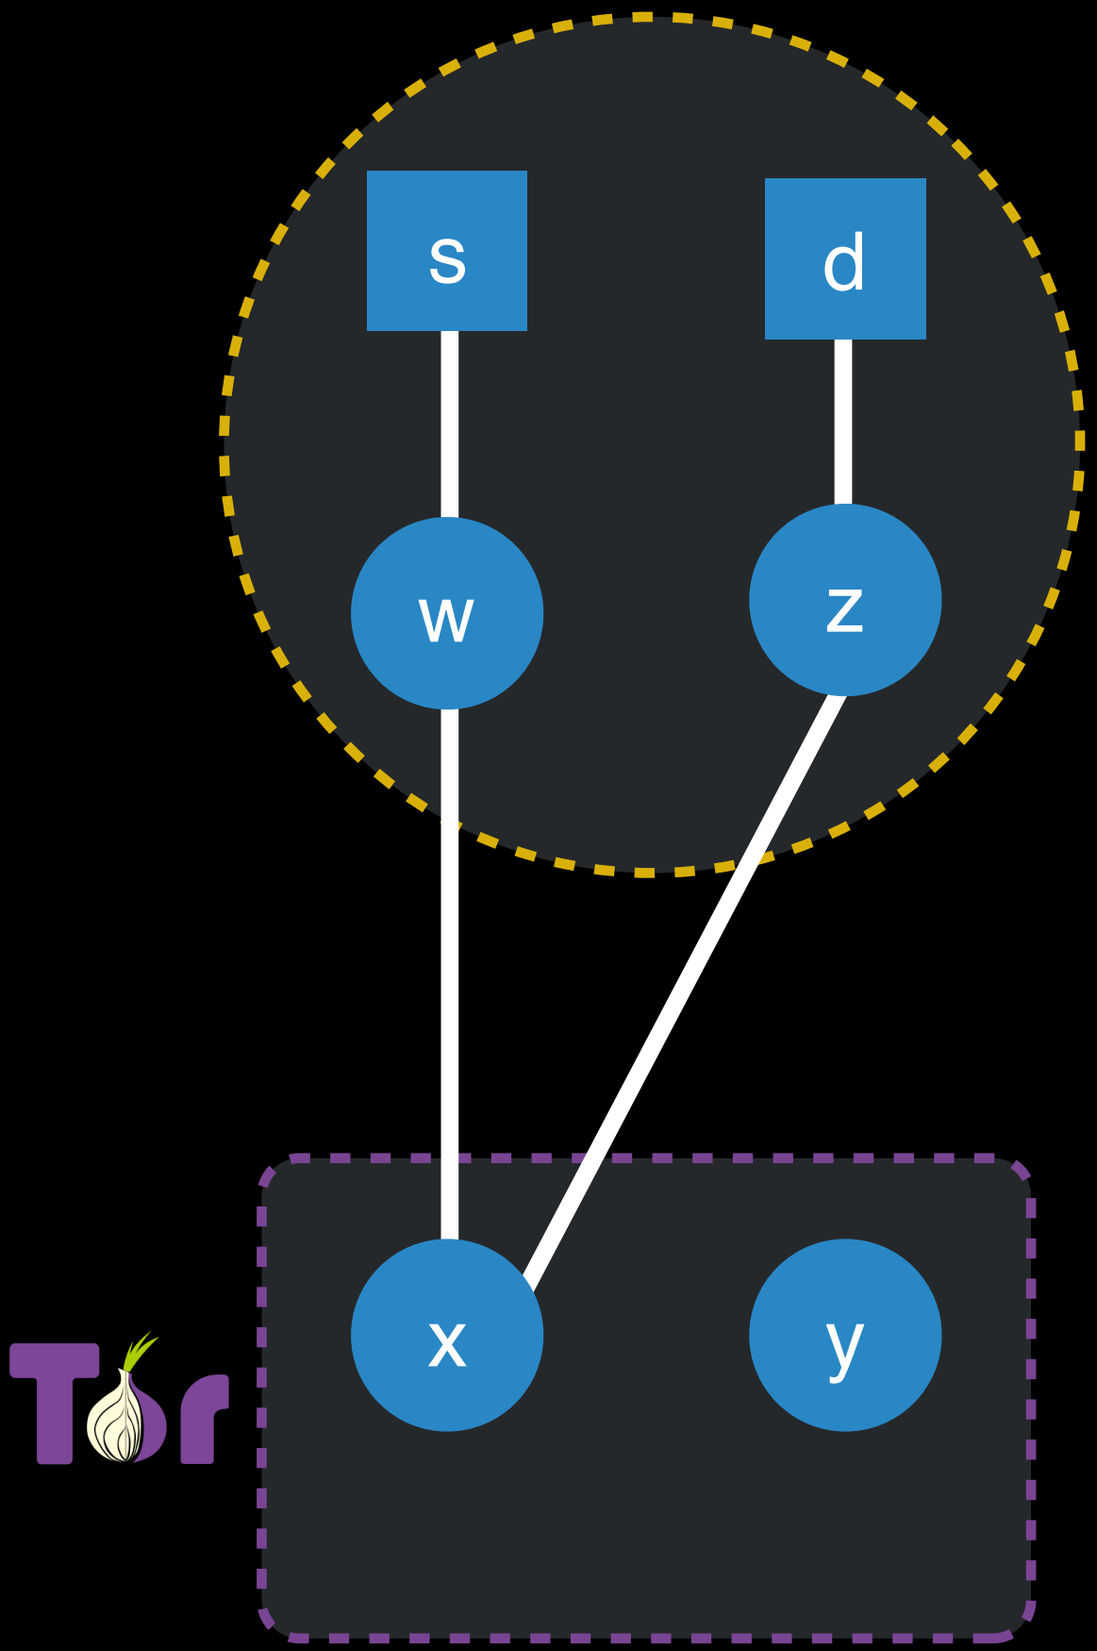

1 2 3

Isolate **RTT** between client and x

Summary

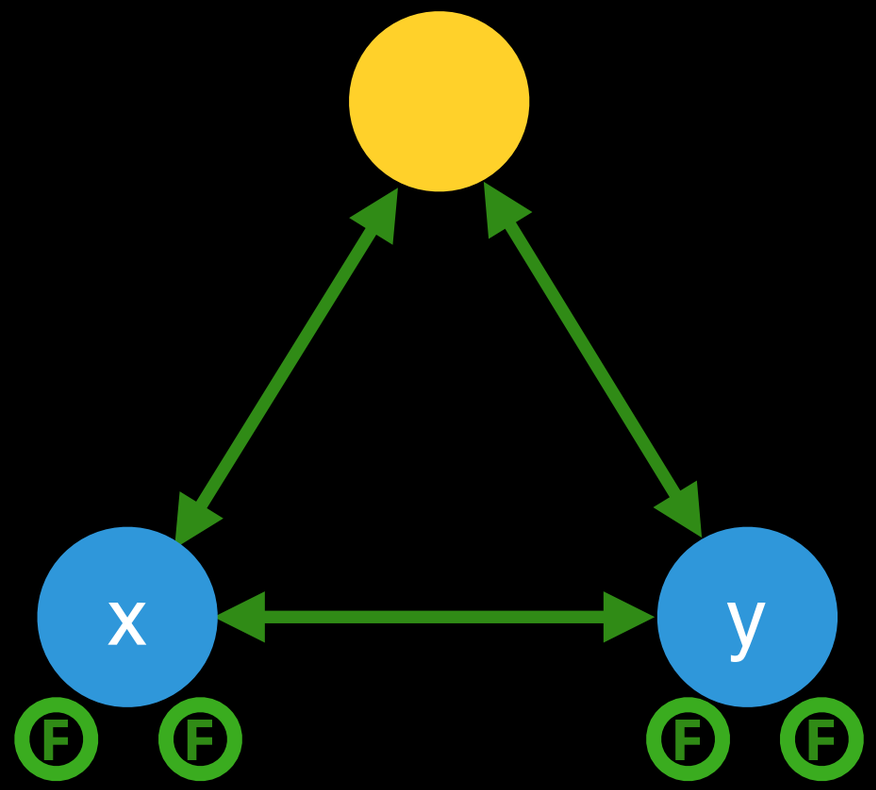

Measurement Host

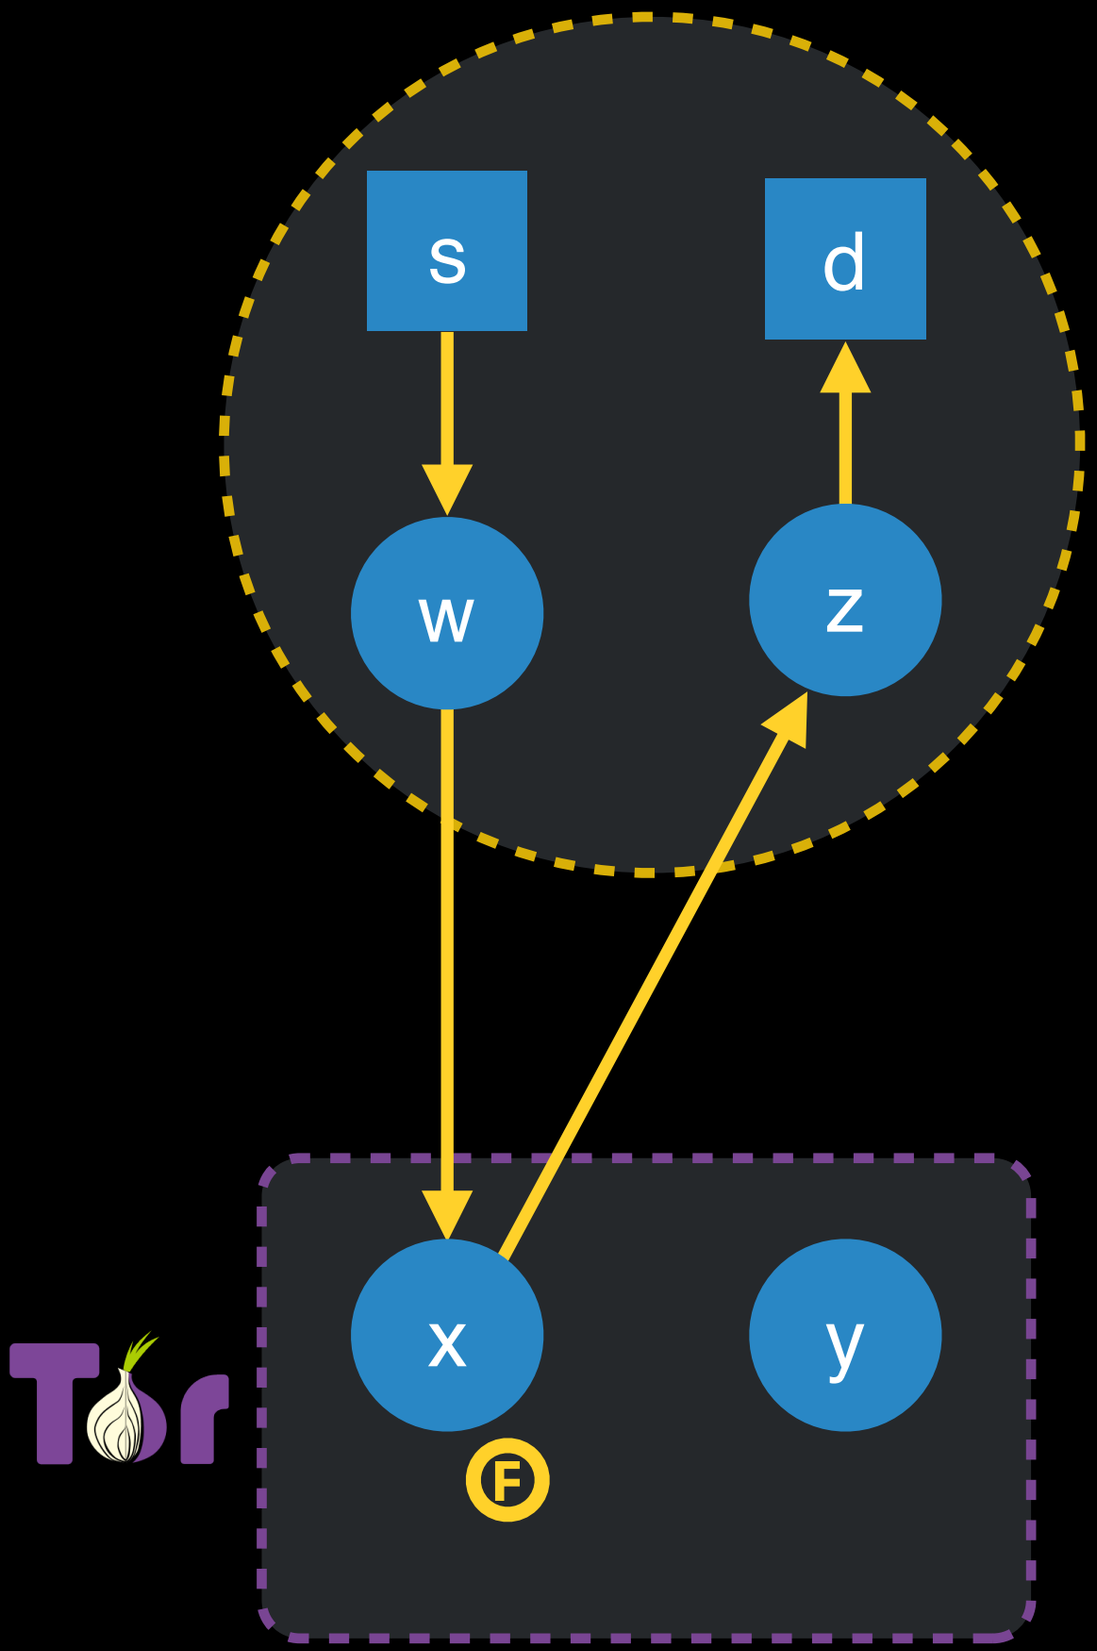

1 2 3

Isolate **RTT** between client and x

Summary

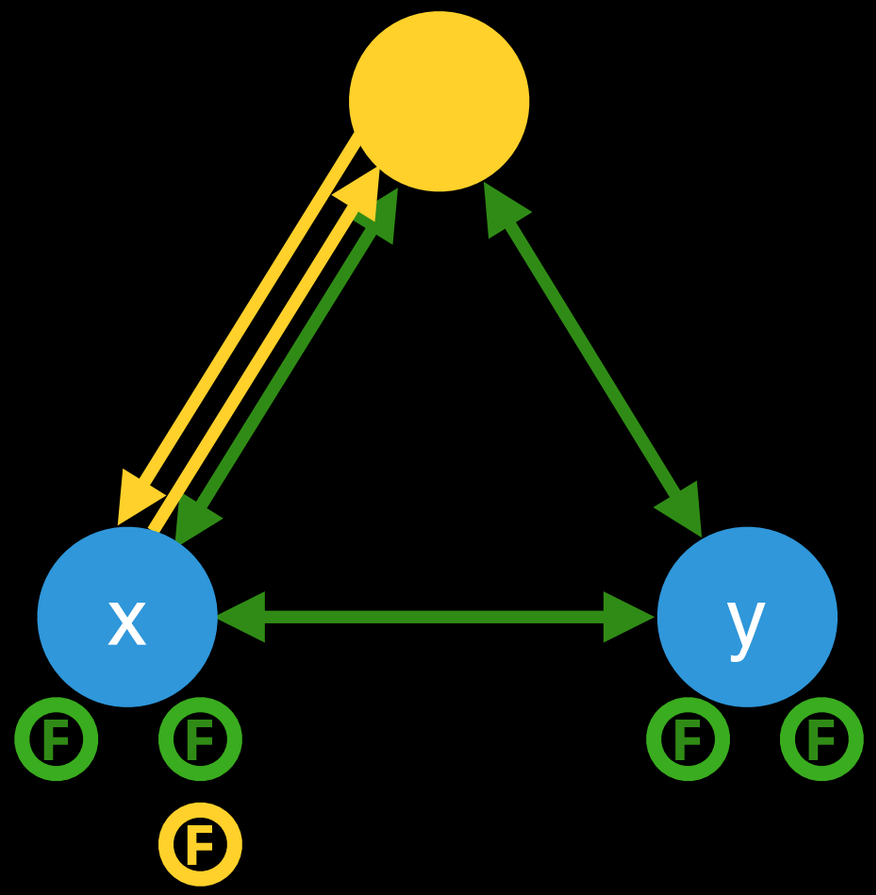

Measurement Host

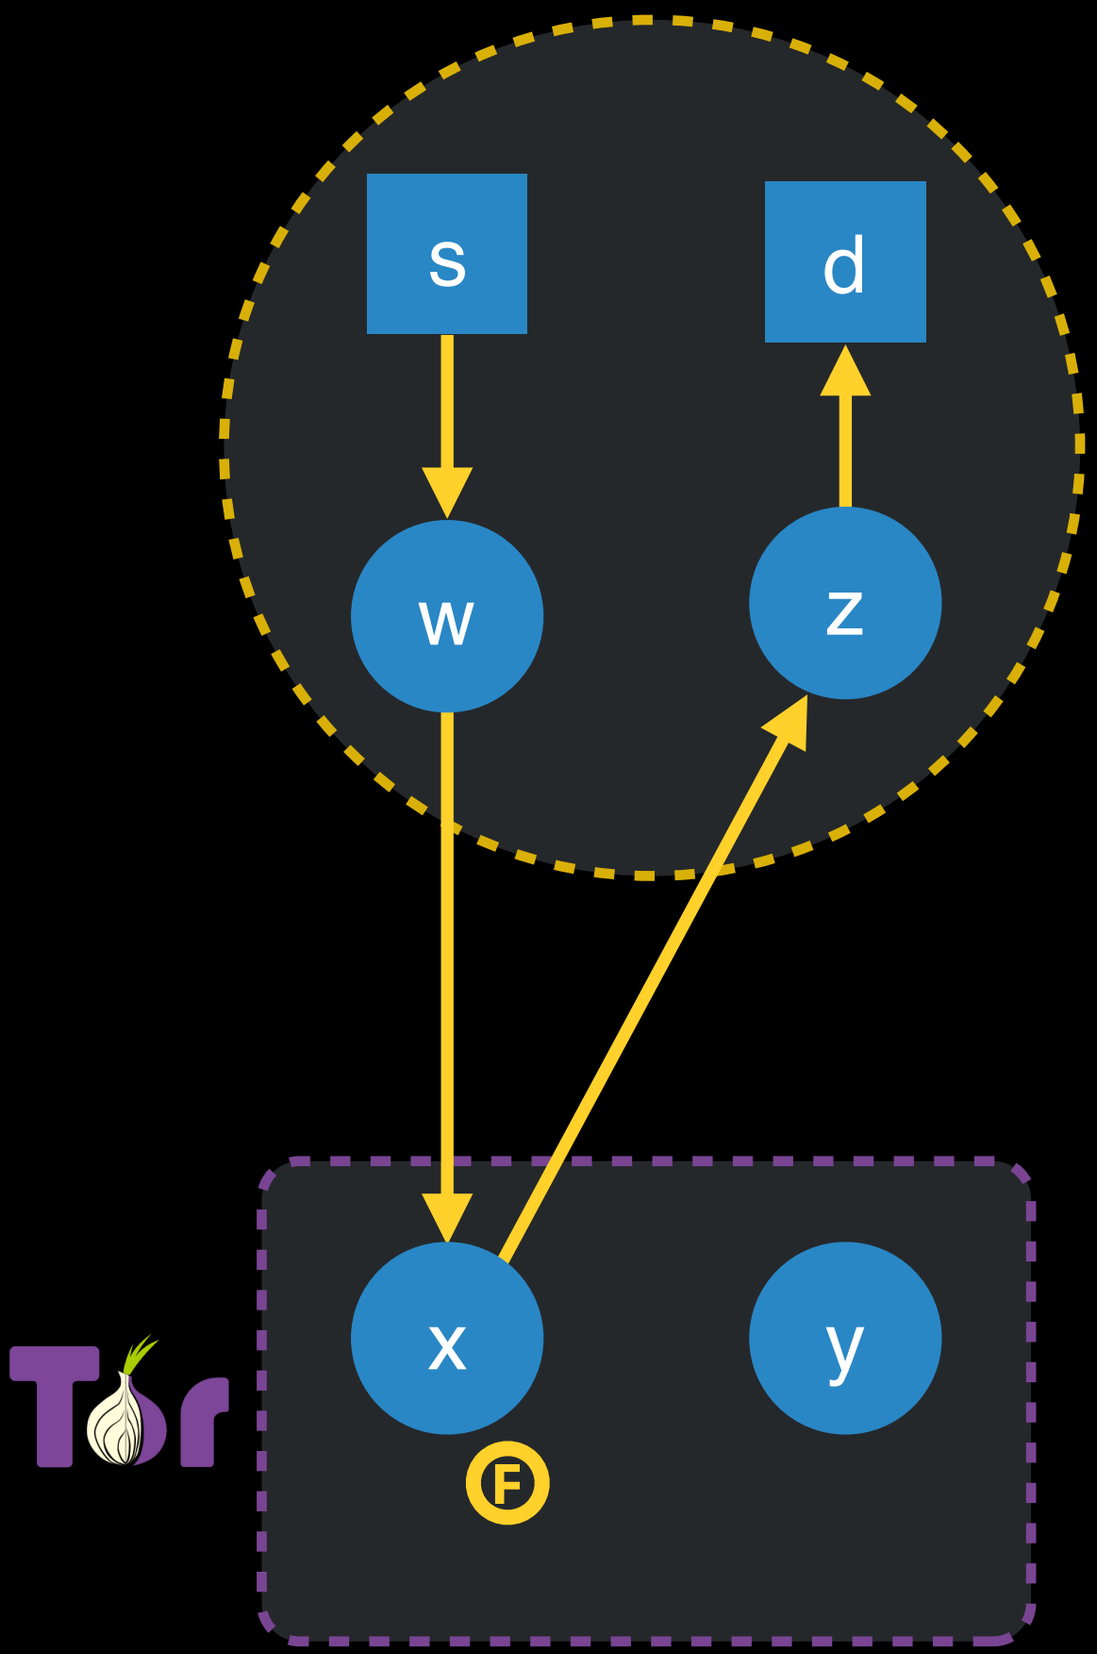

1 2 3

Isolate **RTT** between client and x

Summary

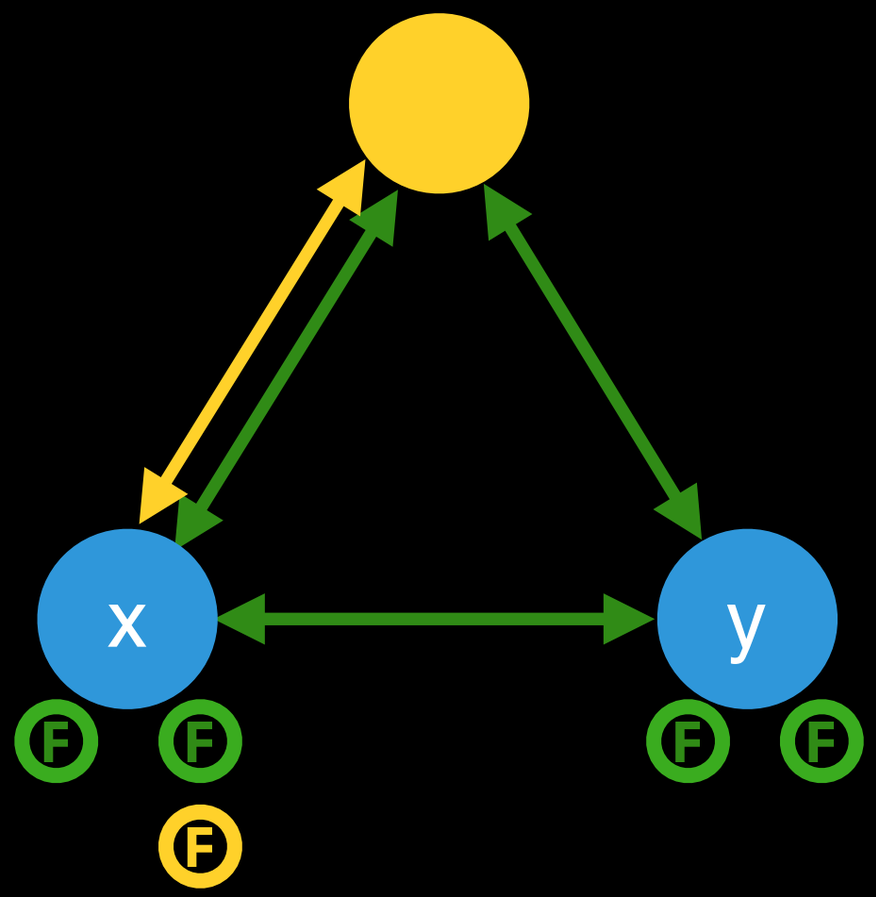

Measurement Host

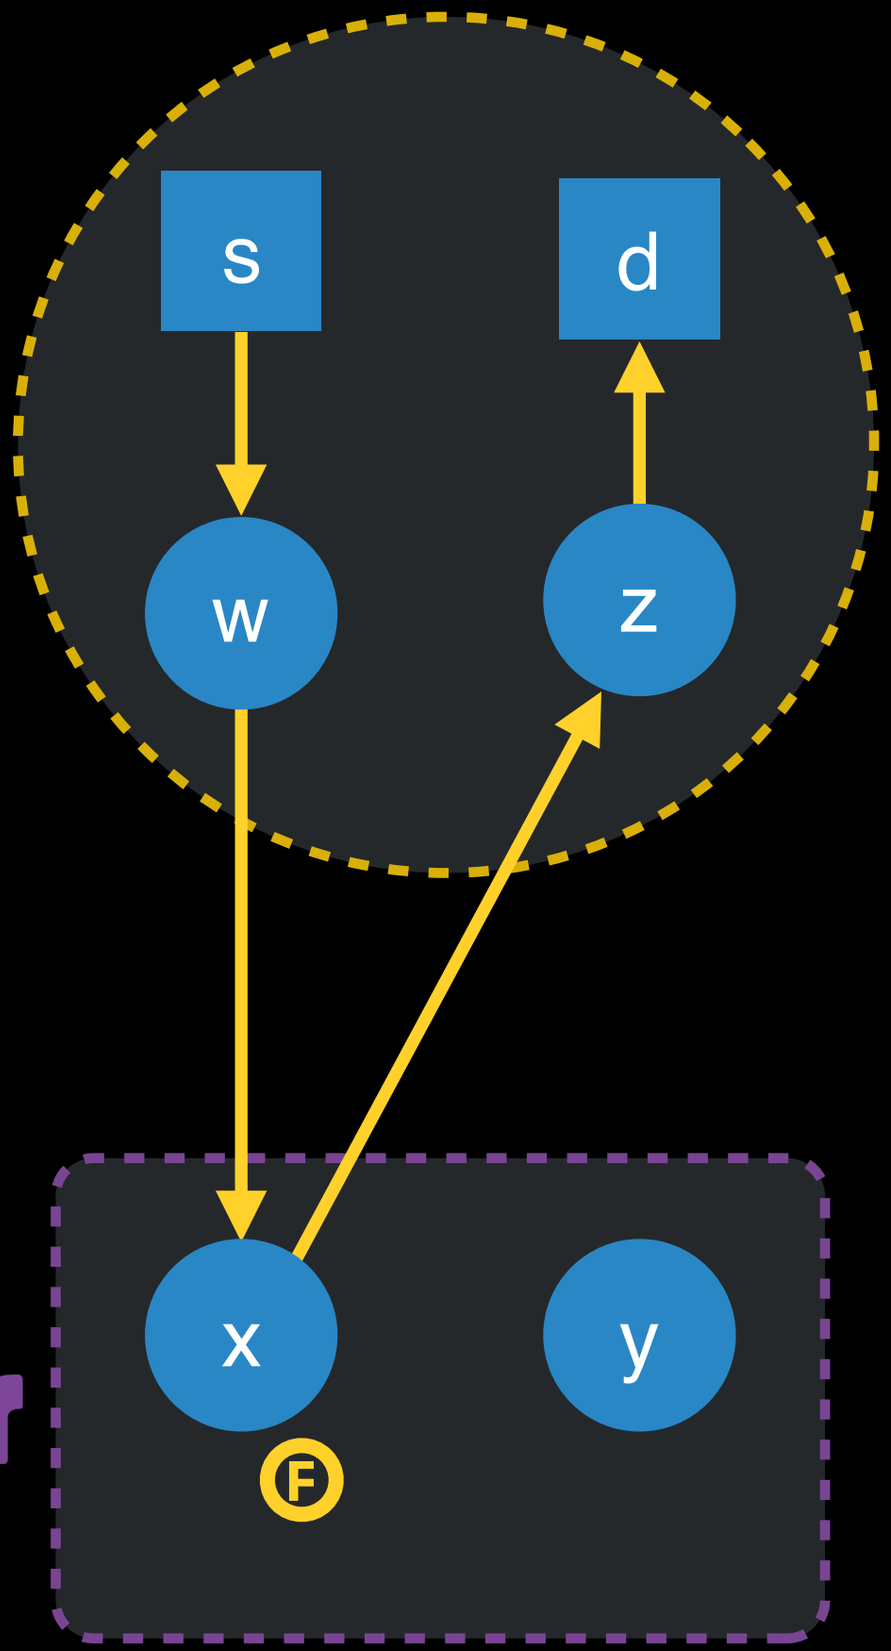

1 - 2 3

Isolate **RTT** between client and x

Summary

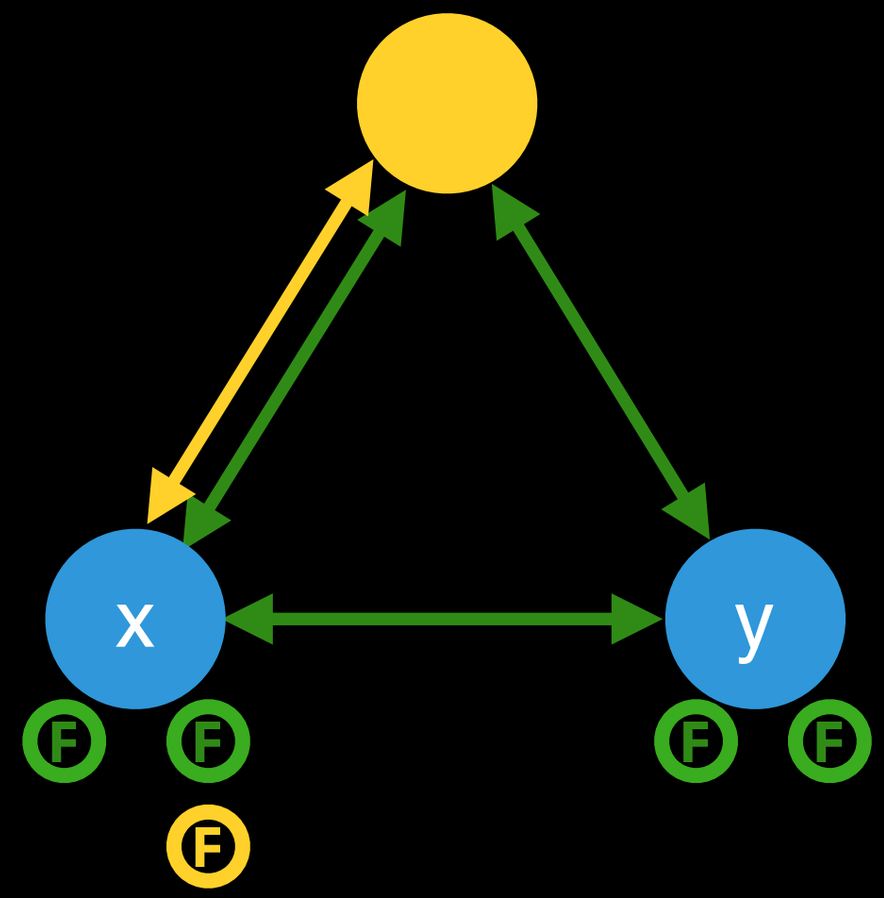

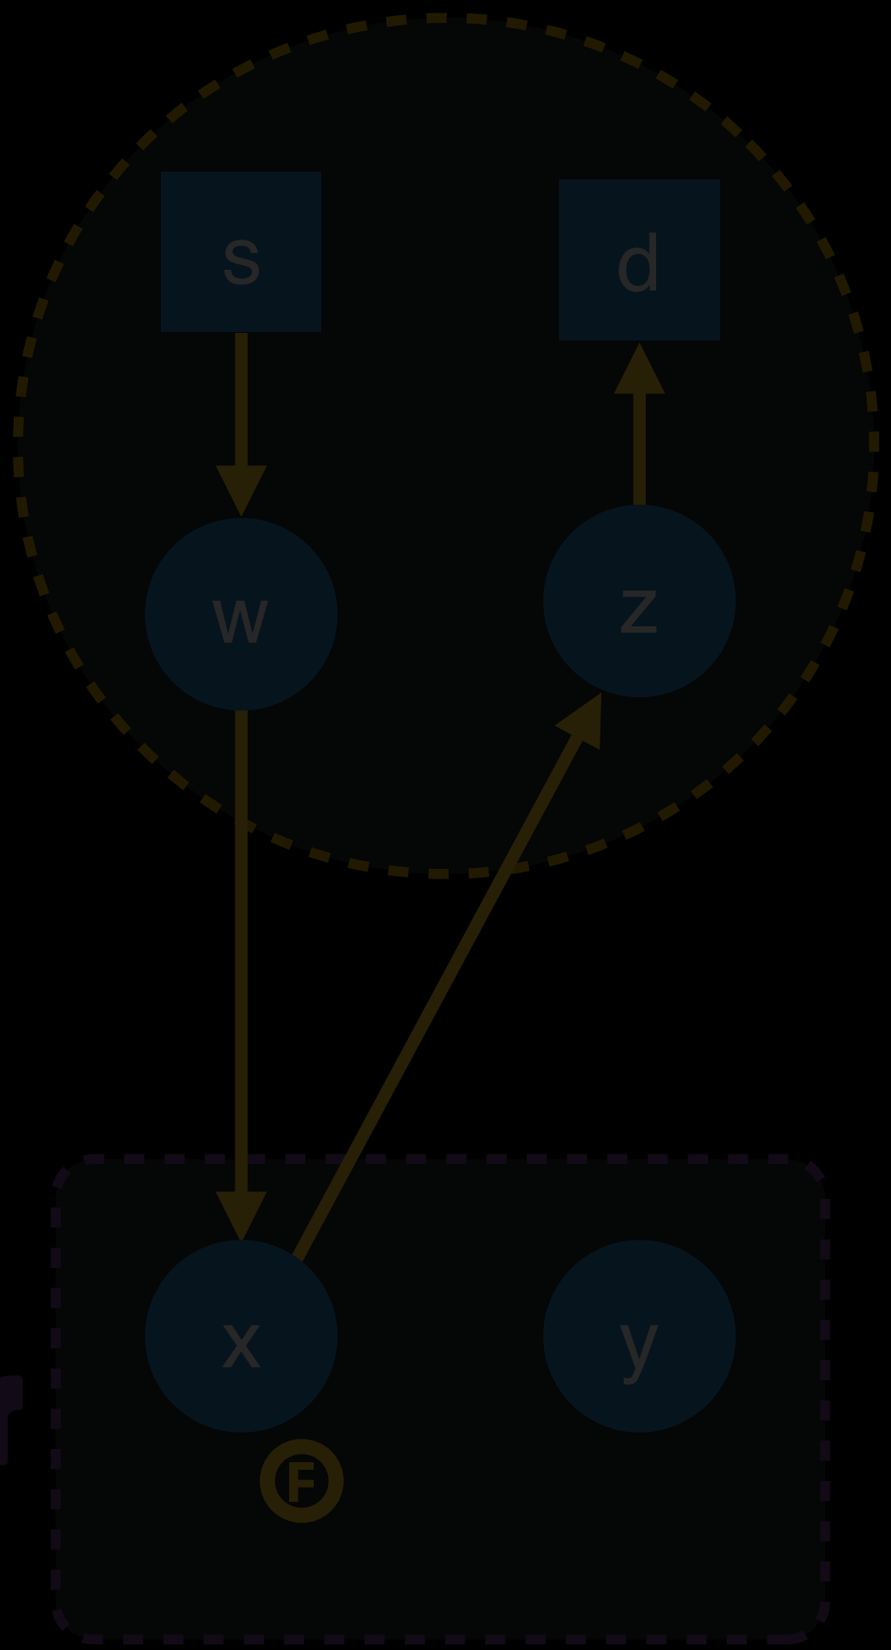

Isolate **RTT** between client and x

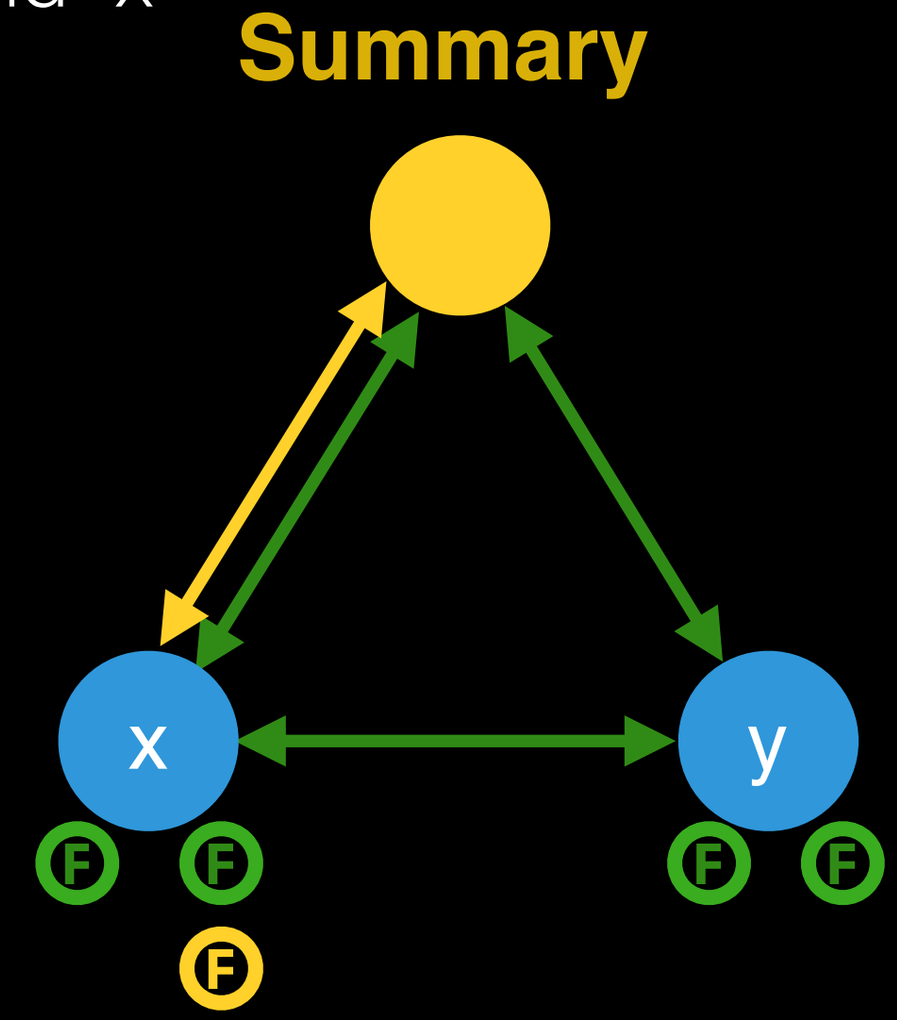

Measurement Host

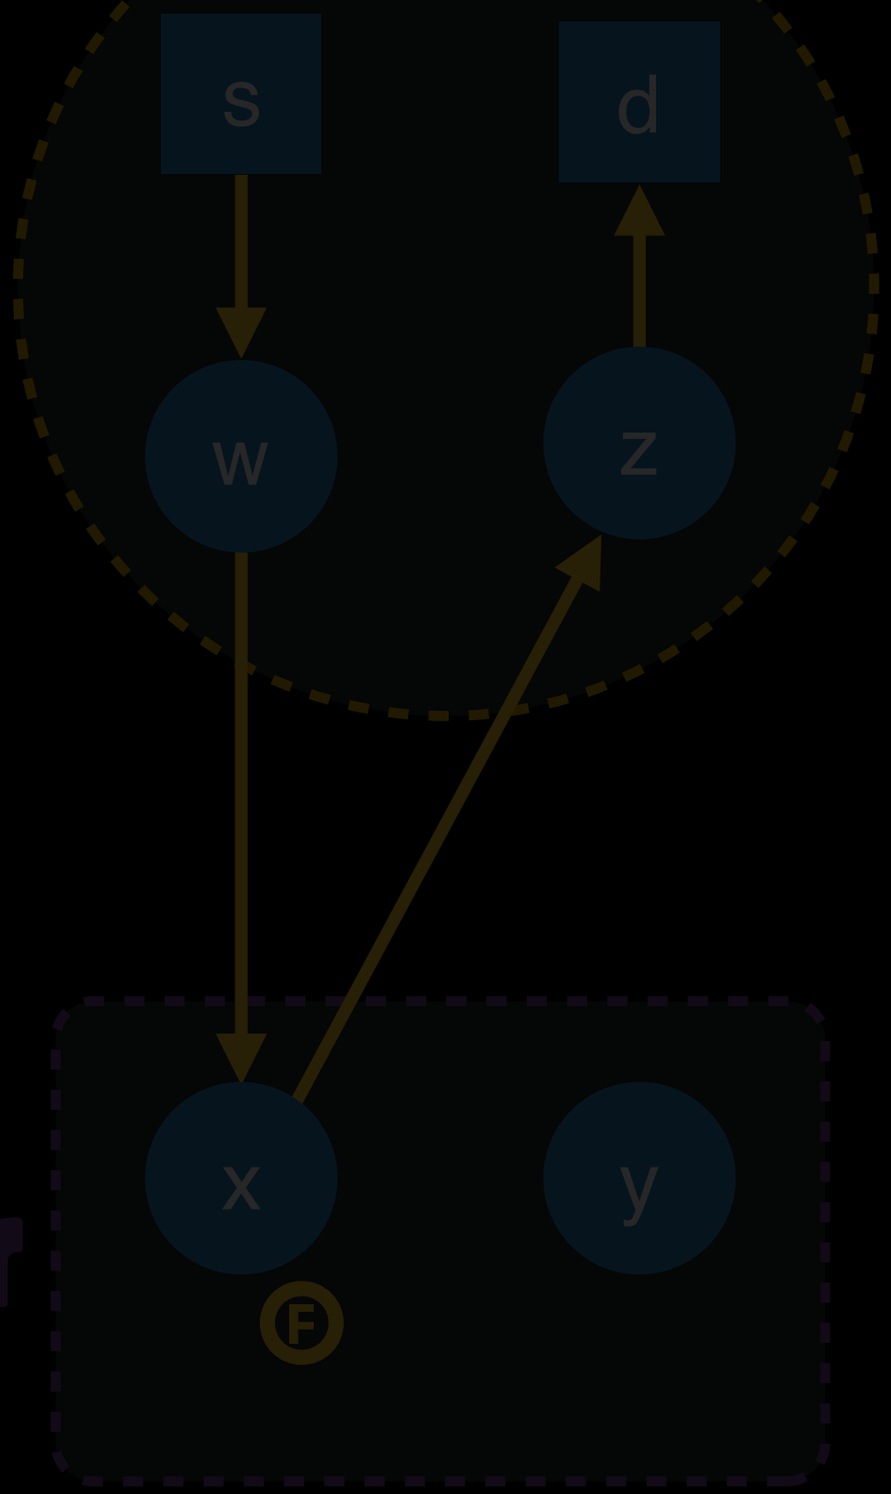

Isolate **RTT** between client and x

Summary

Subtract latencies

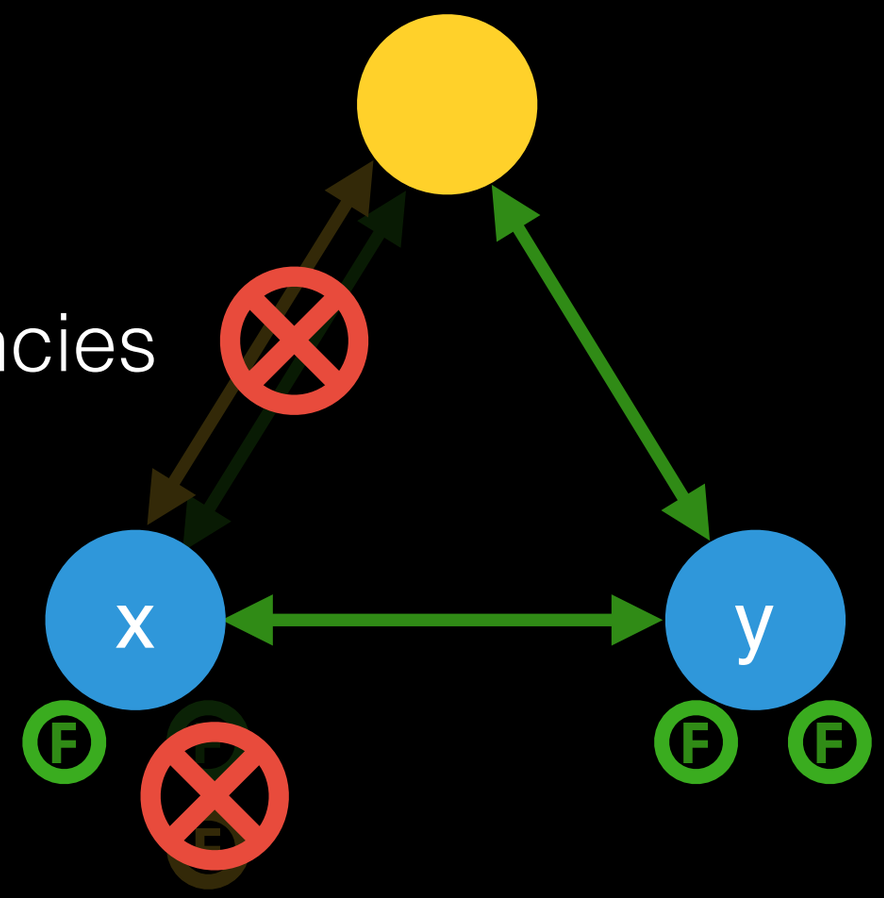

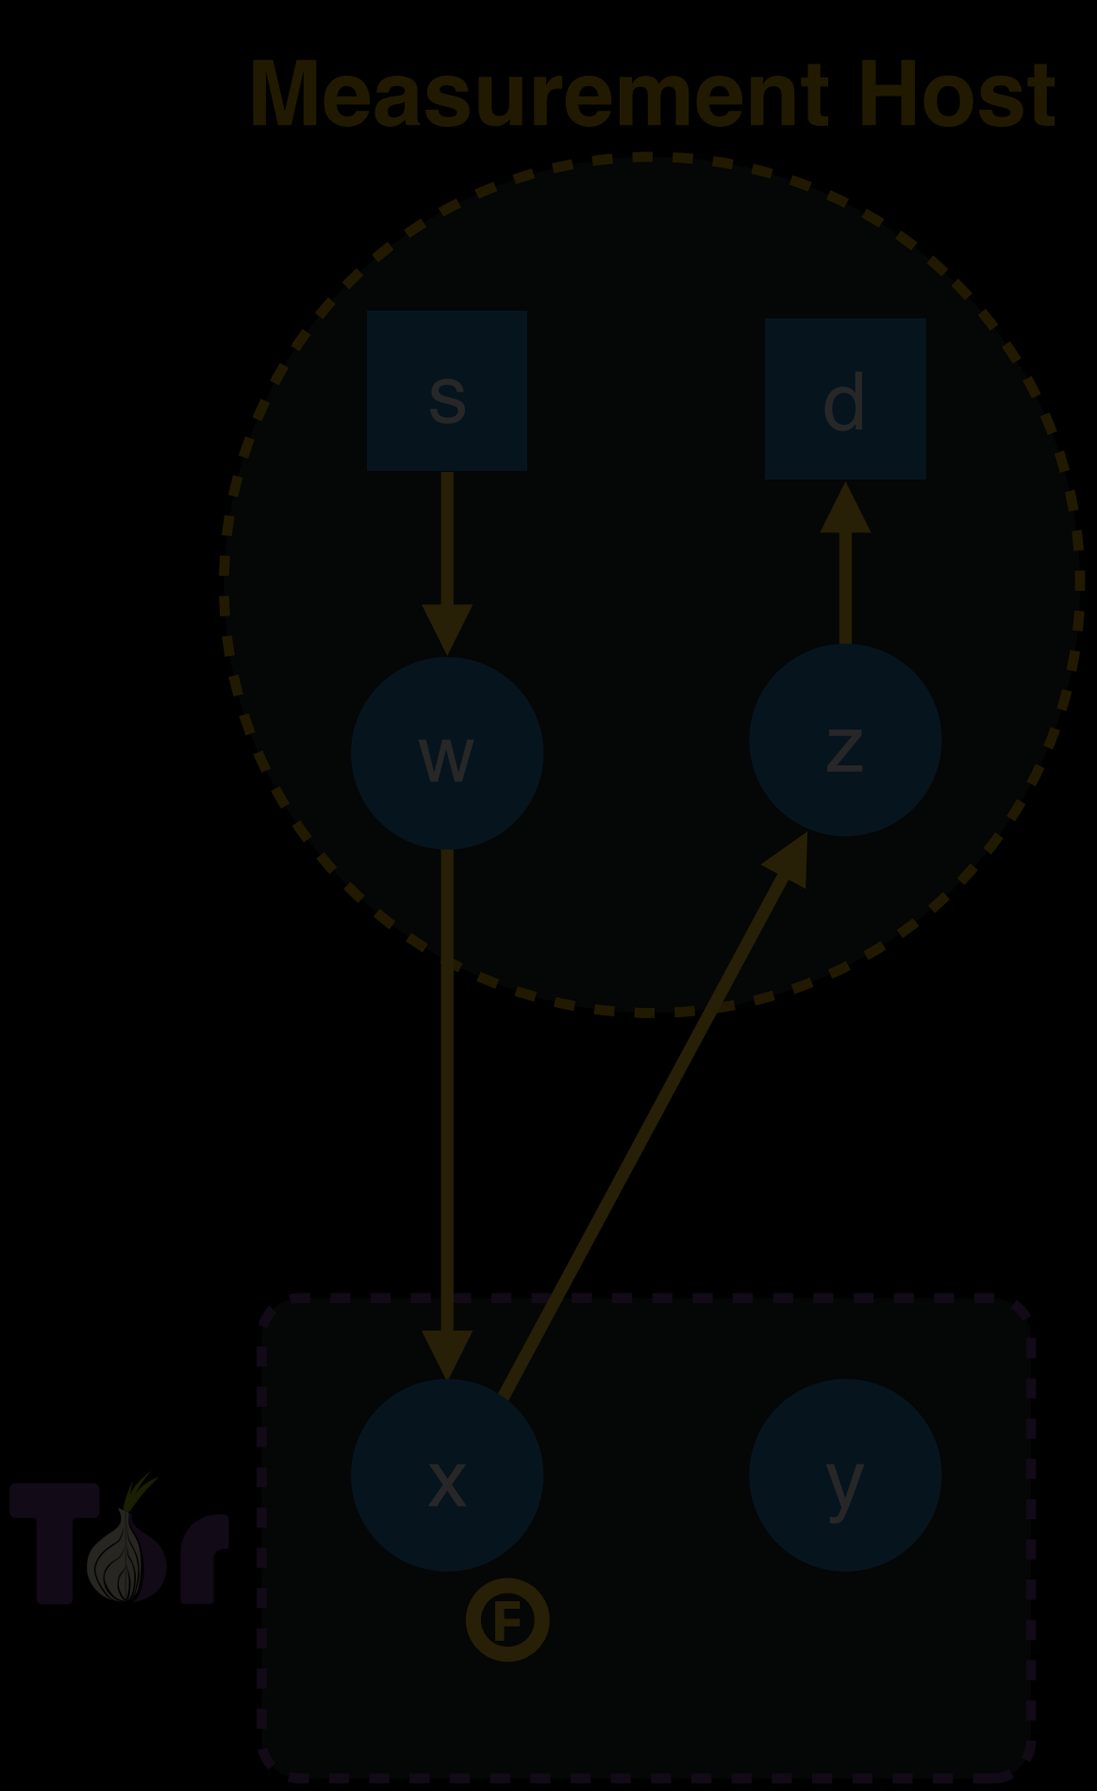

1 – 2 3

Isolate **RTT** between client and x

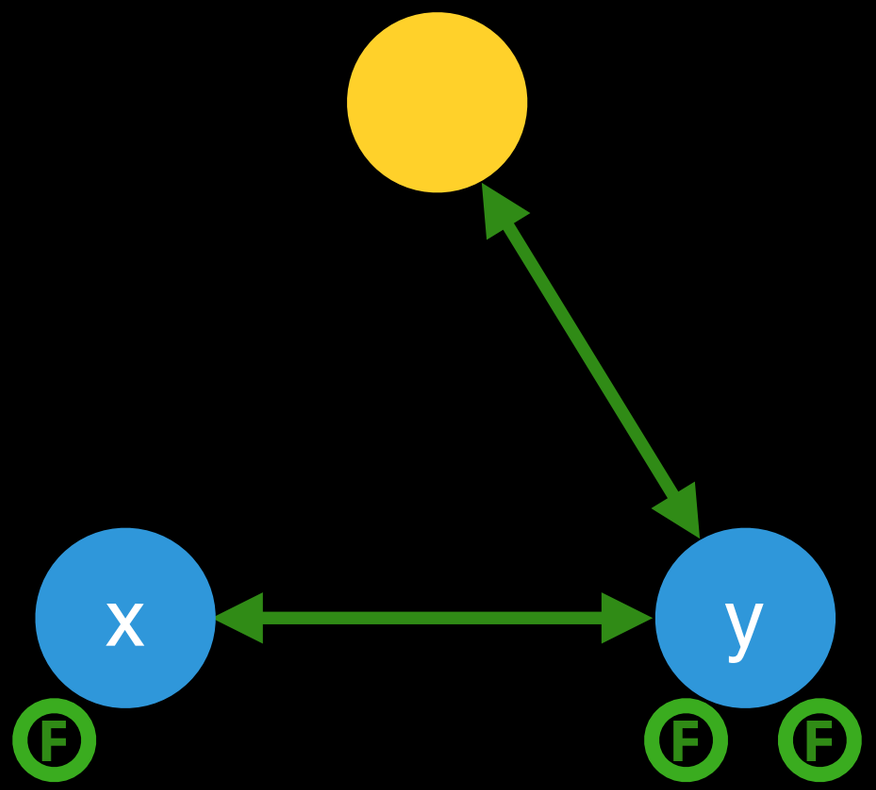

Measurement Host

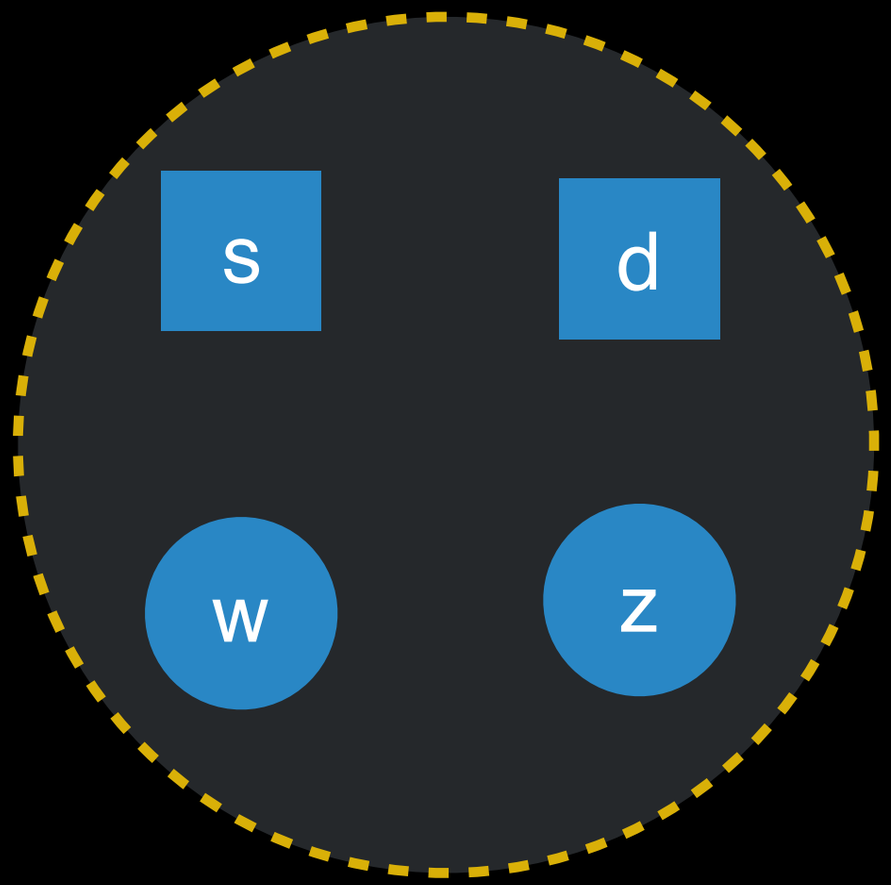

1 2 3

Isolate **RTT** between client and y

Summary

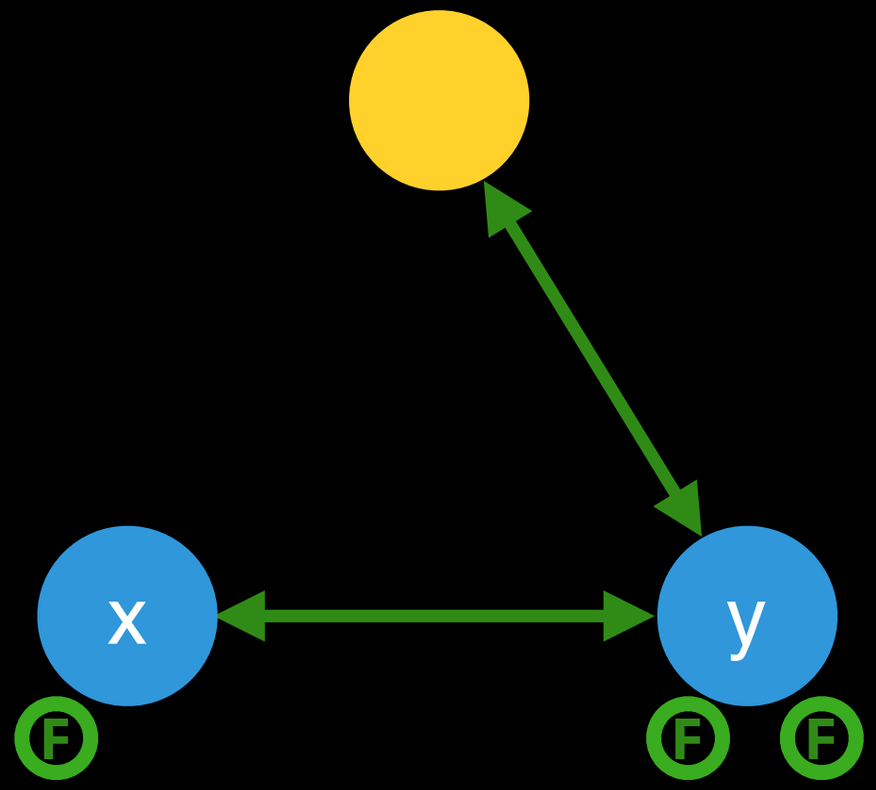

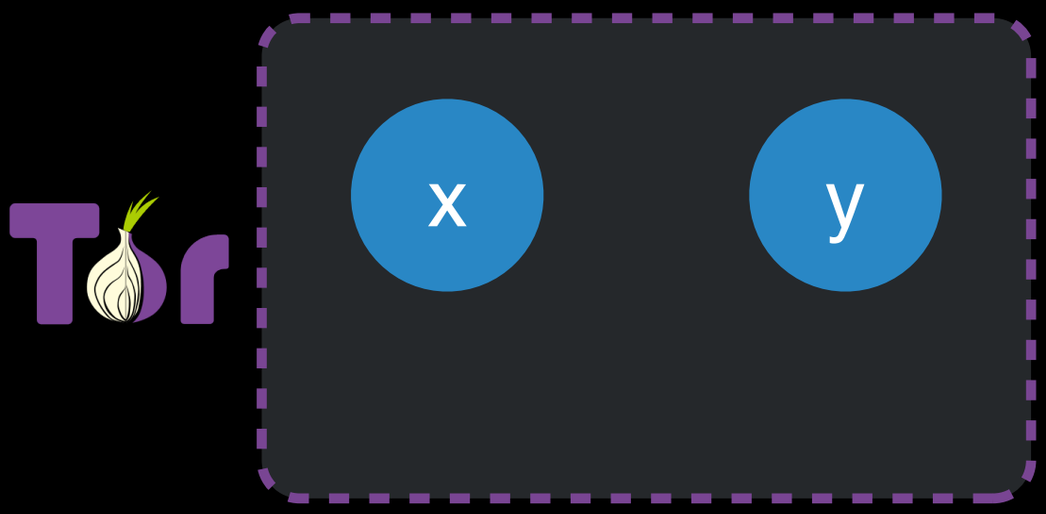

Measurement Host

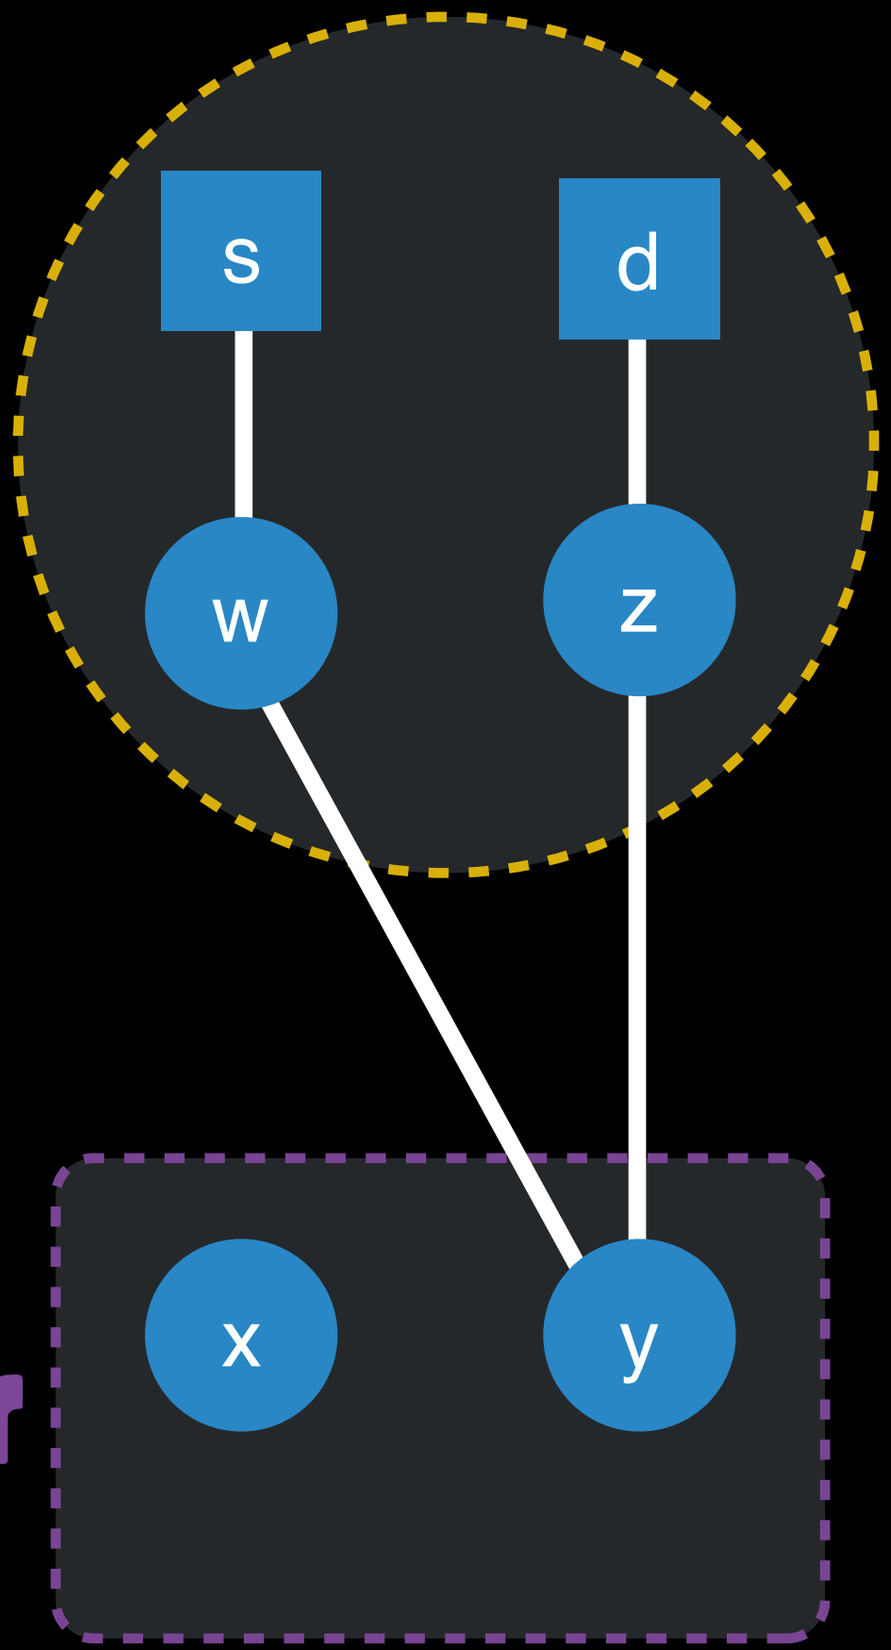

1 2 3

Isolate **RTT** between client and y

Summary

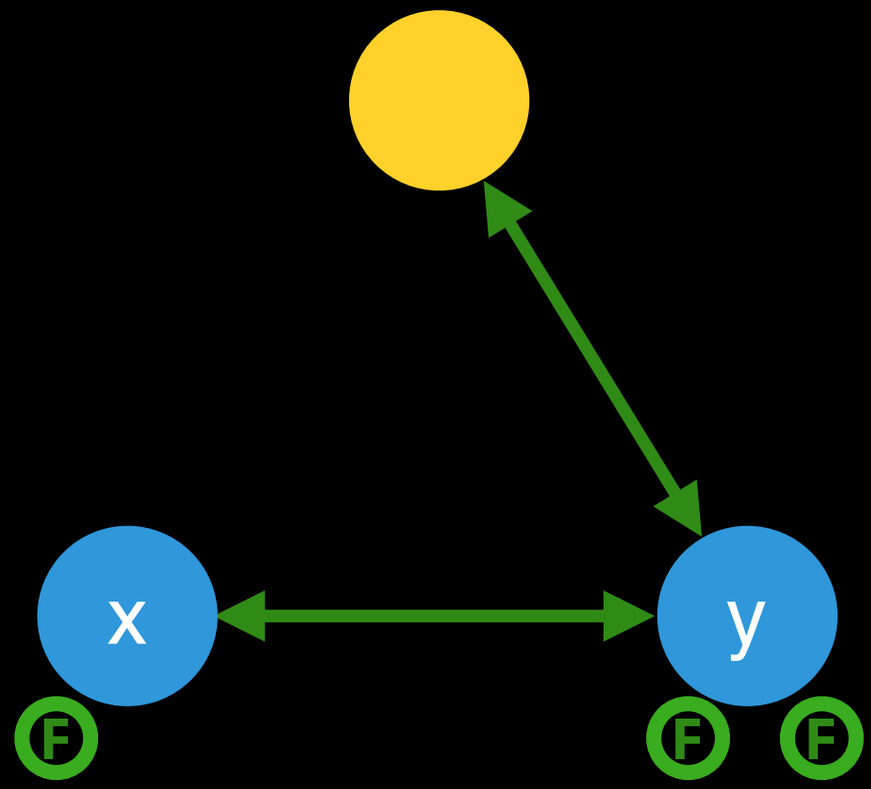

Measurement Host

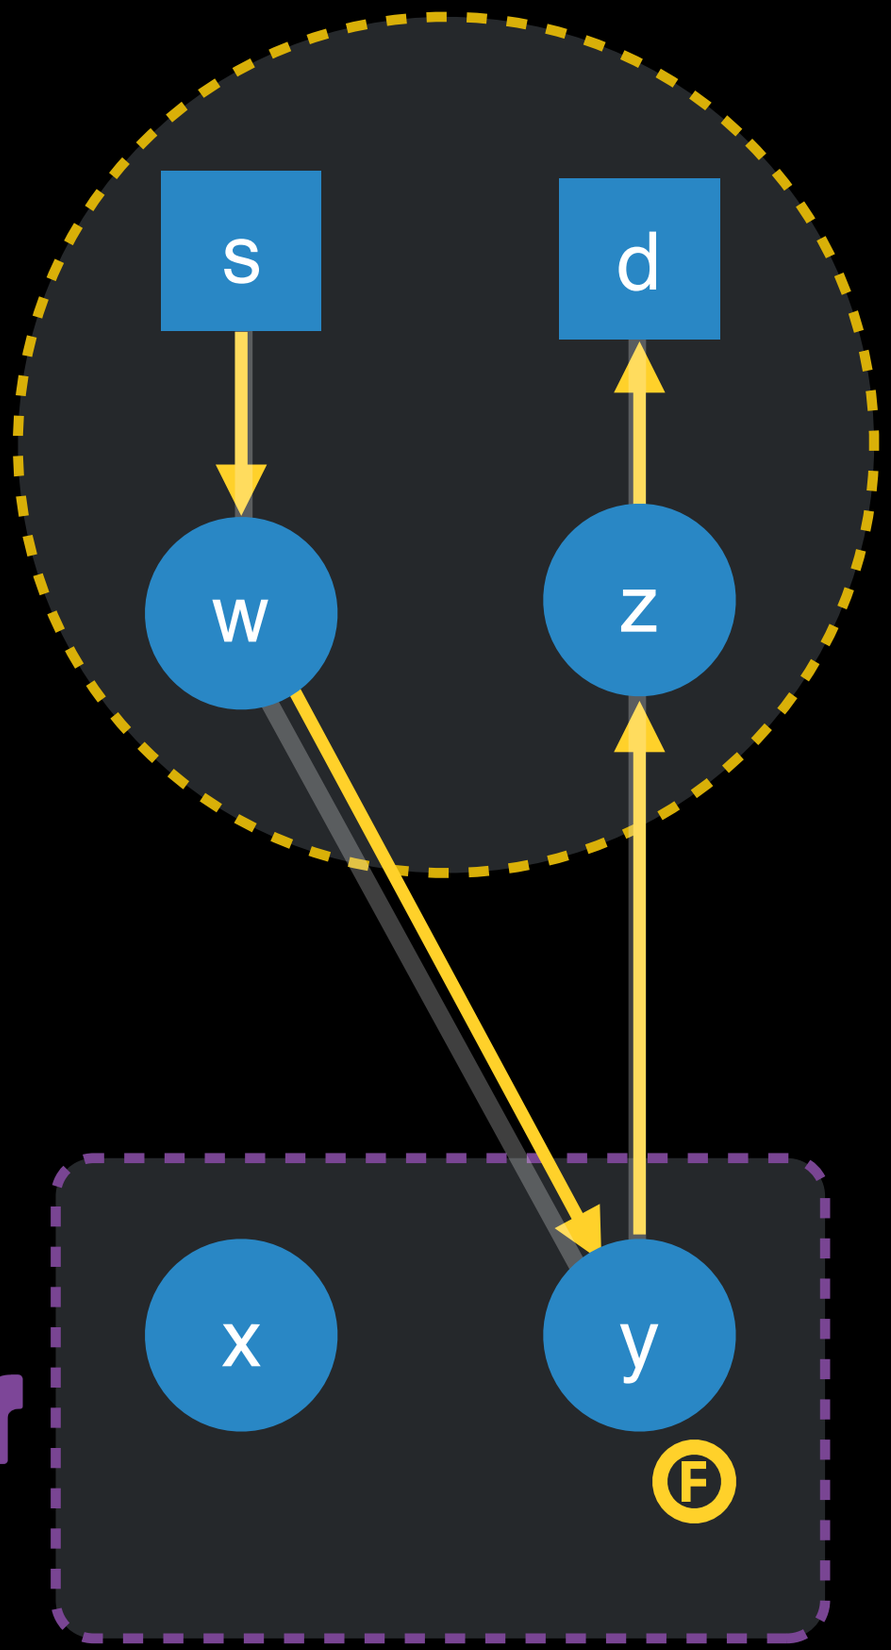

1 2 3

Isolate **RTT** between client and y

Summary

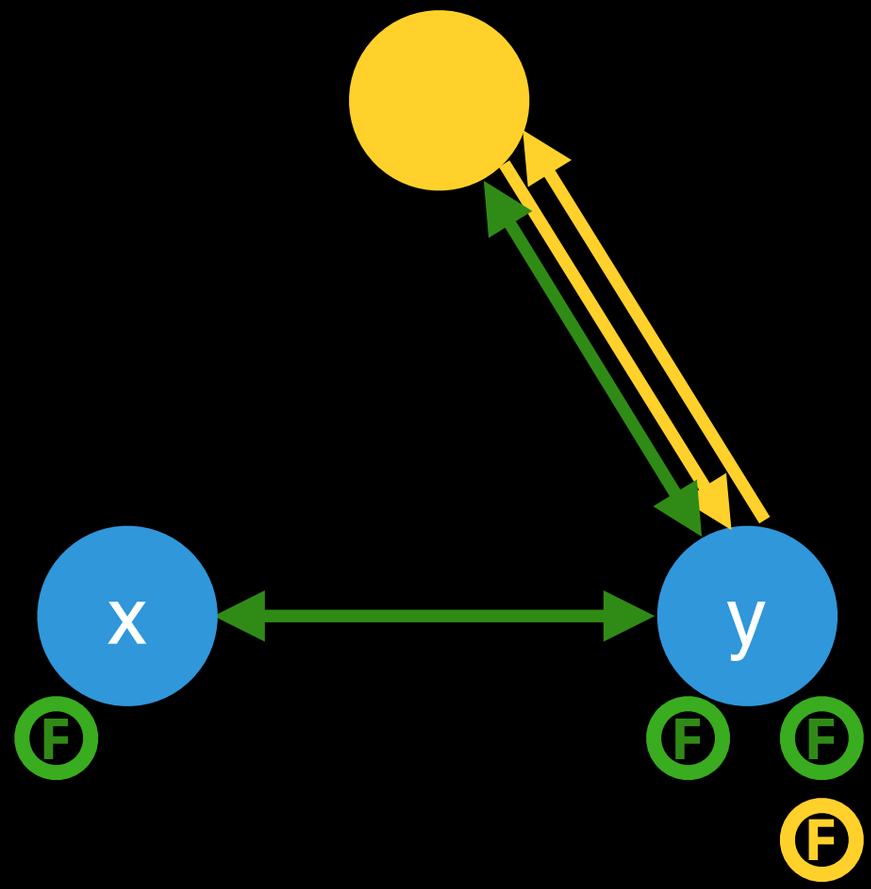

Measurement Host

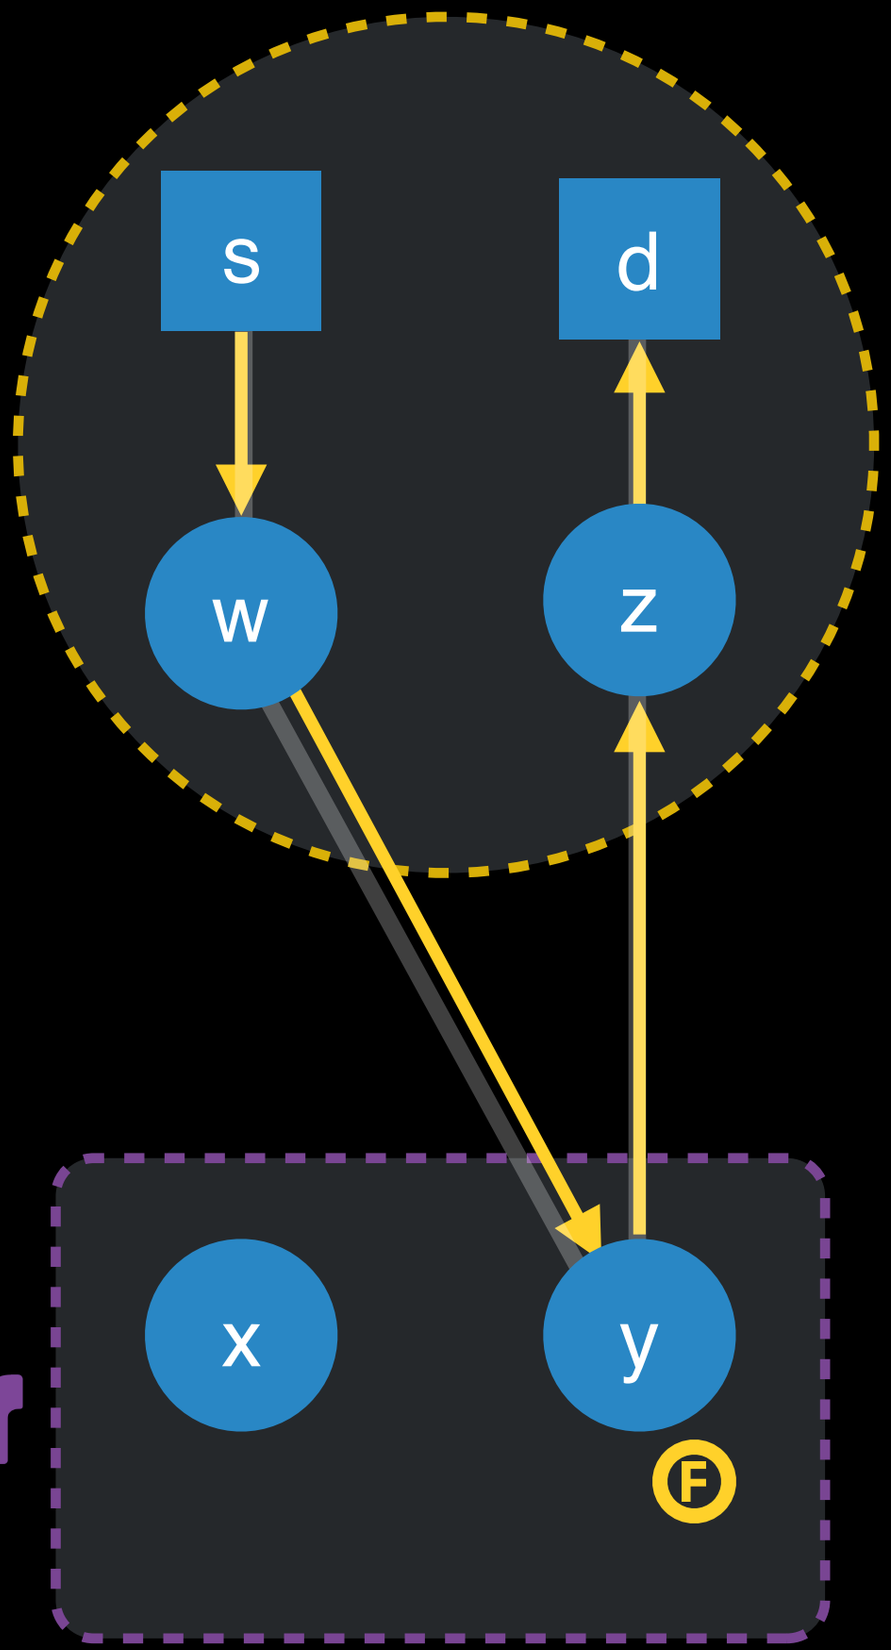

1 2 3

Isolate **RTT** between client and y

Summary

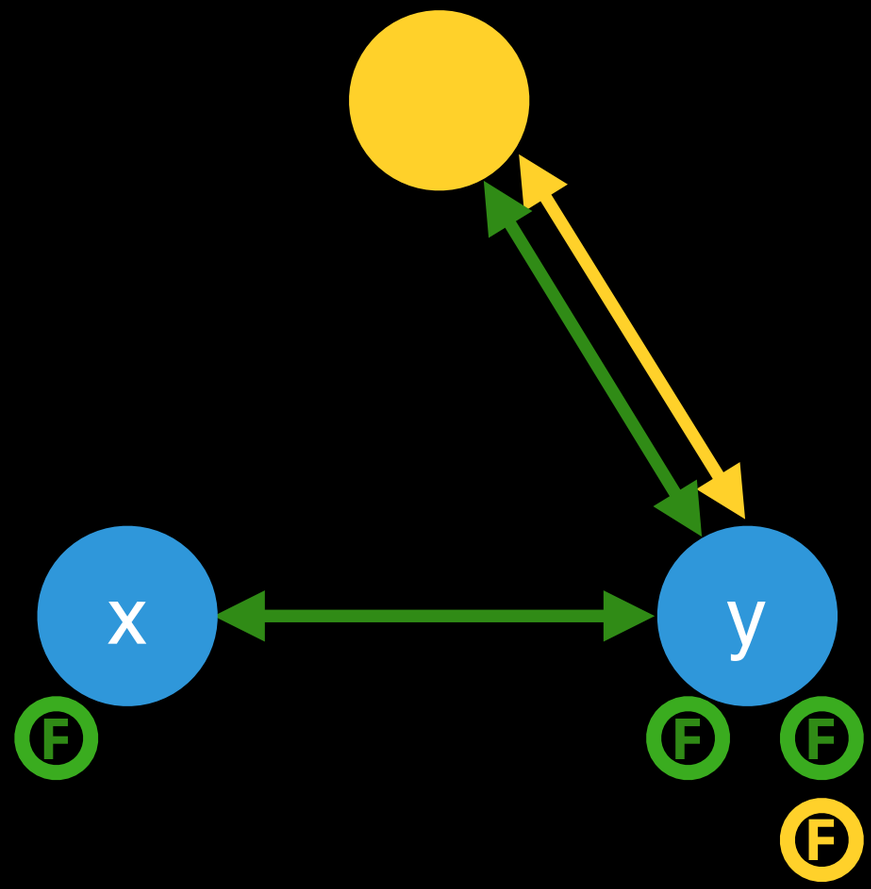

Measurement Host

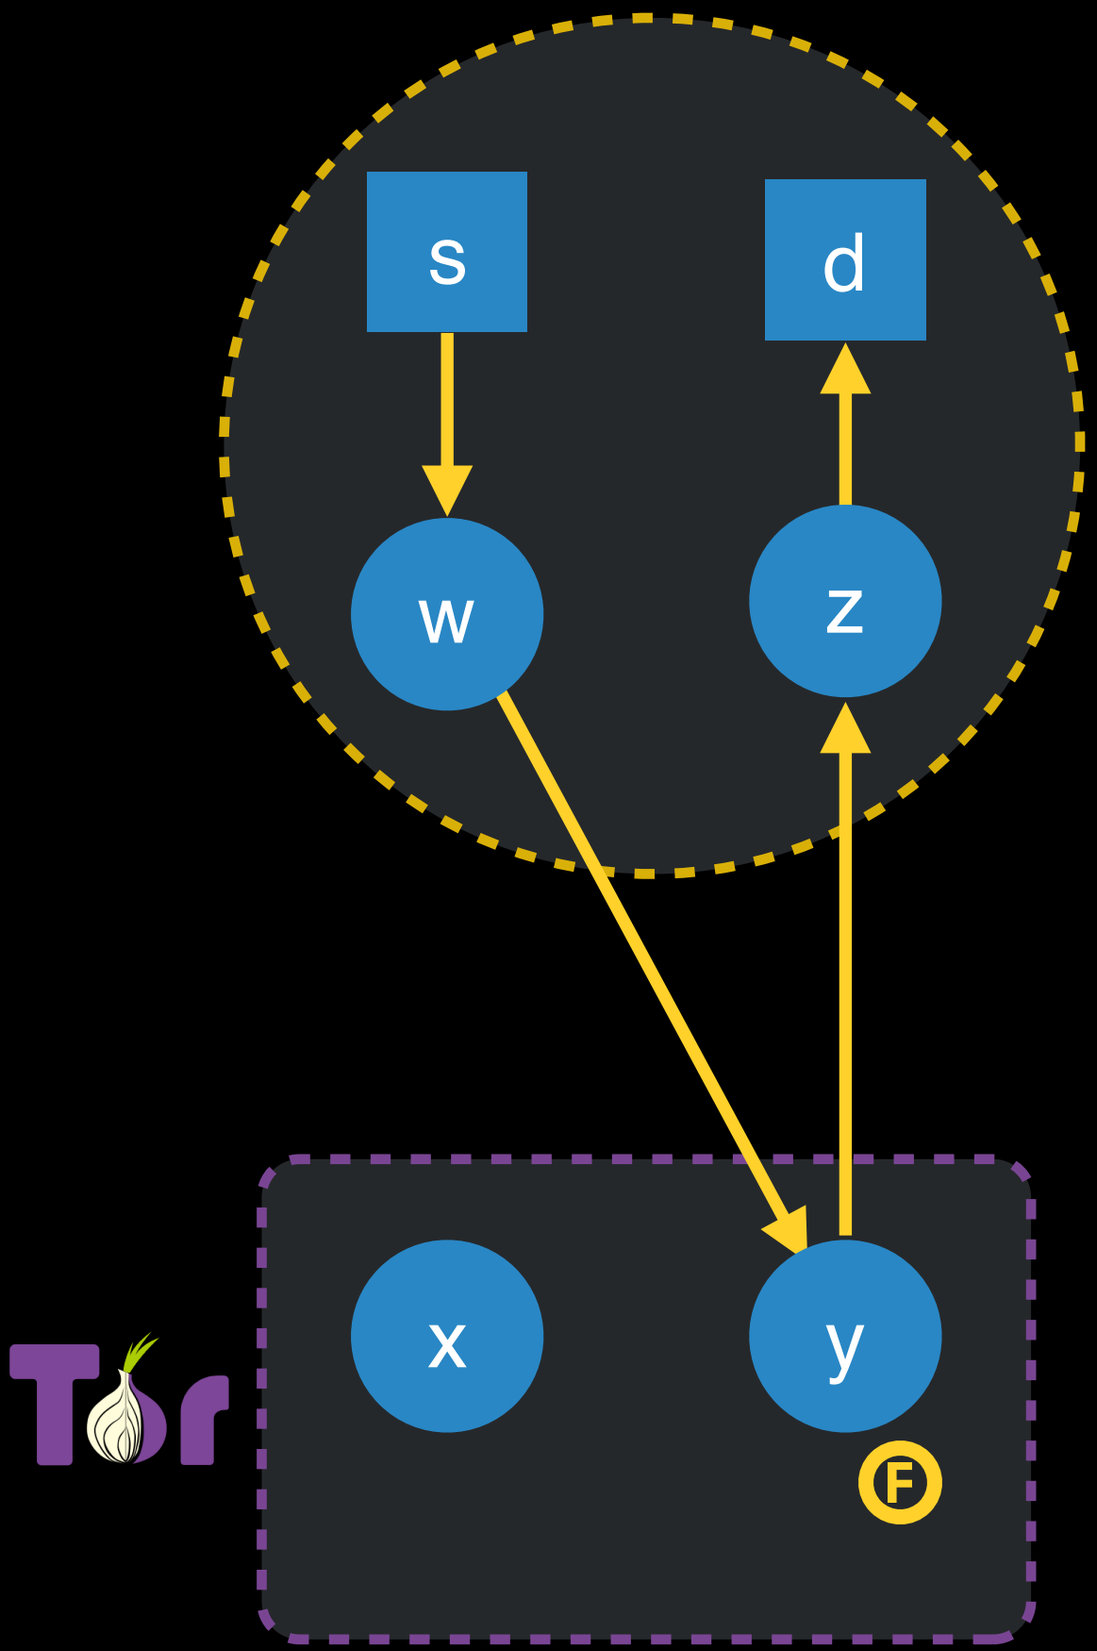

1 2 3

Isolate **RTT** between client and y

Summary

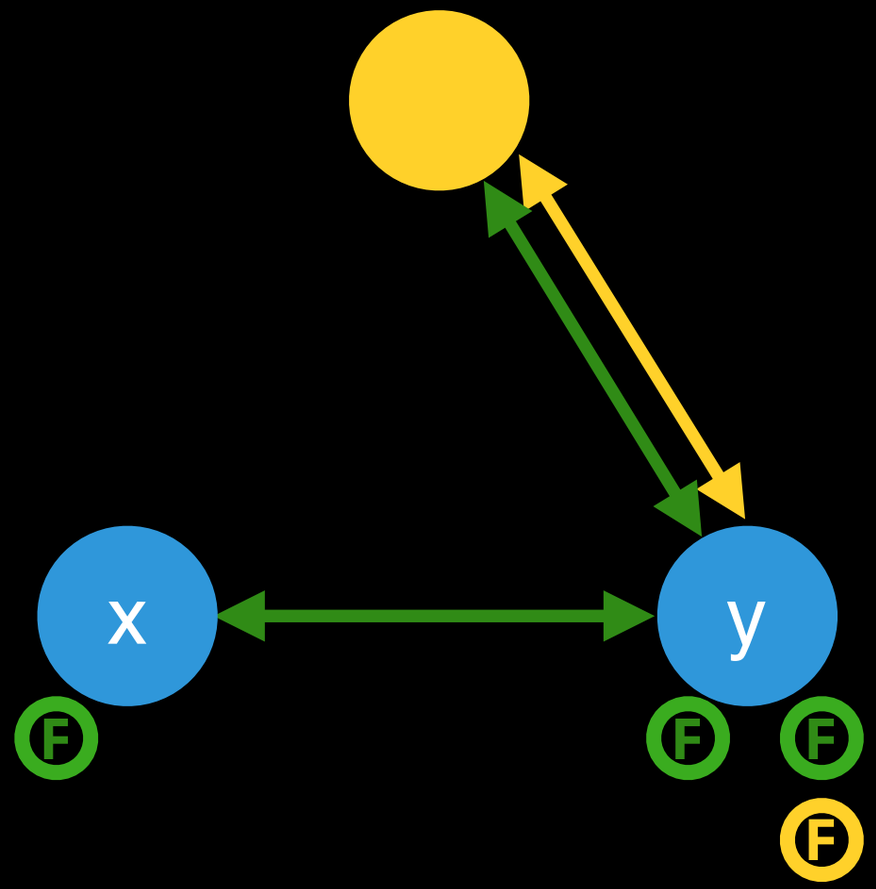

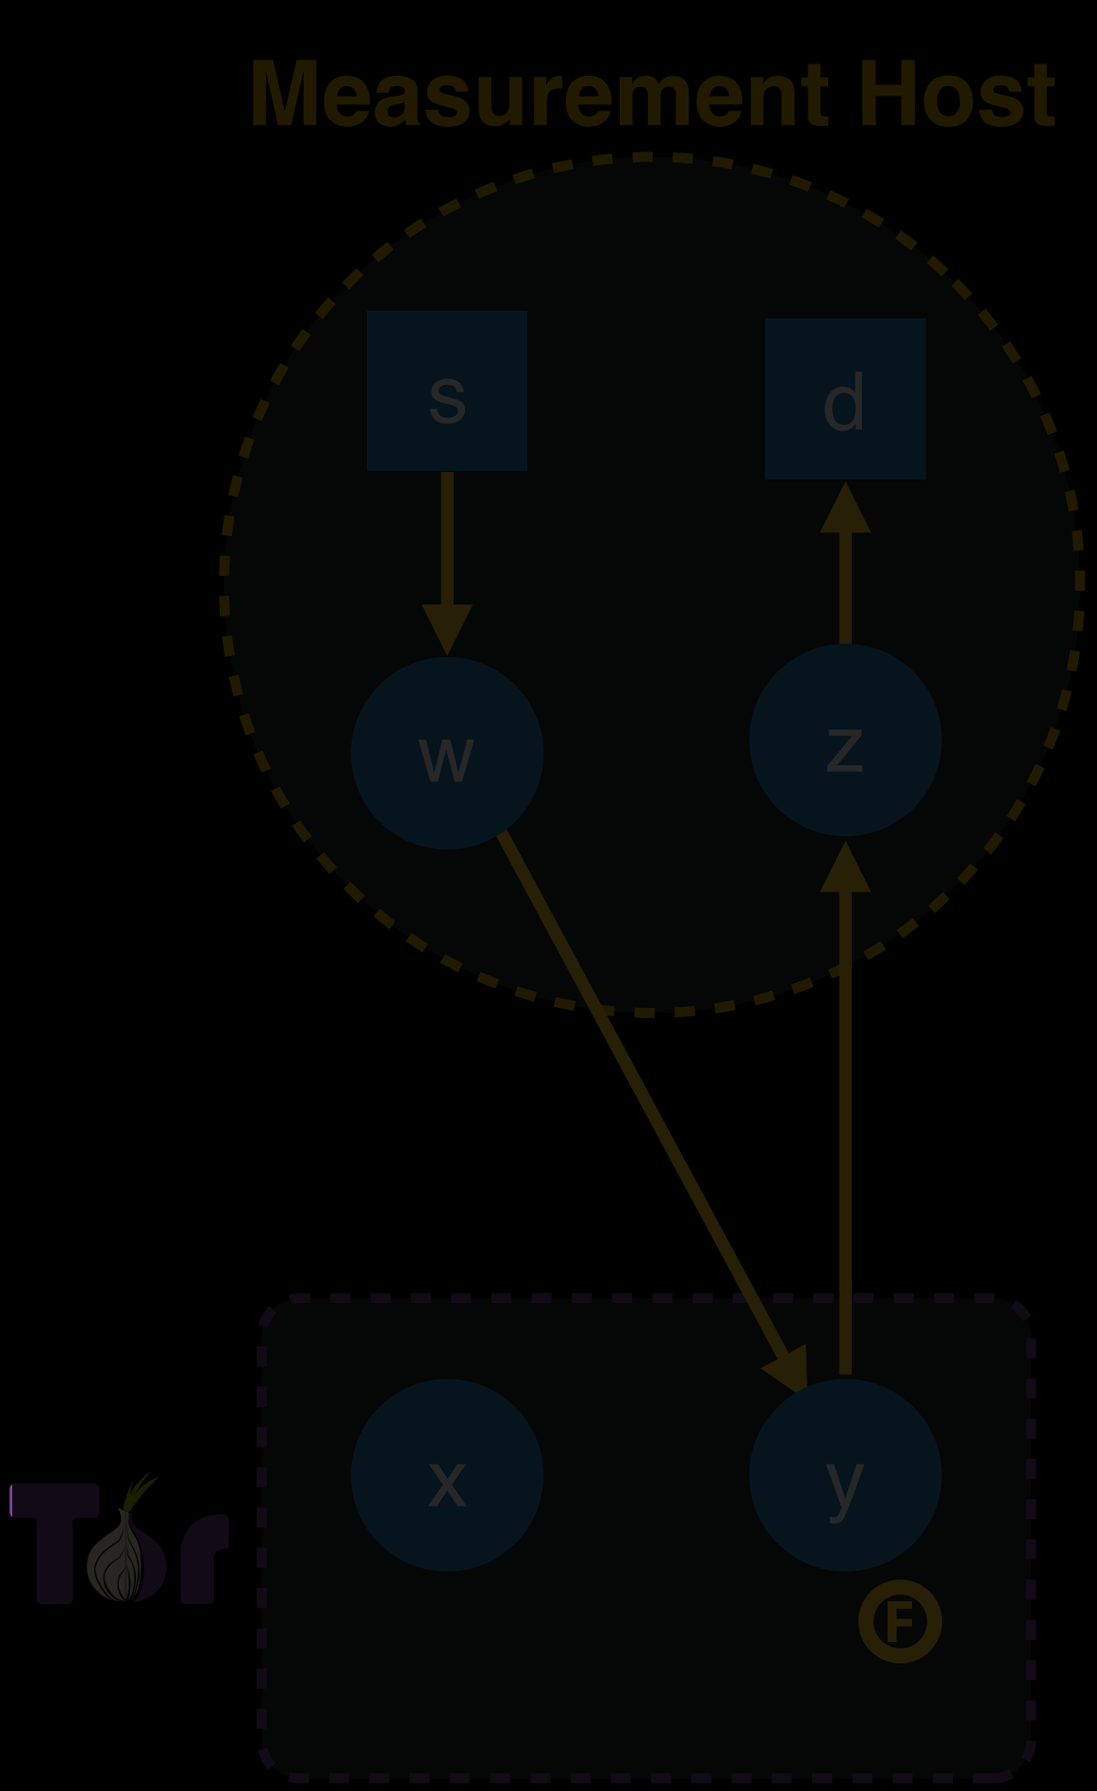

1 – 2 3

Isolate **RTT** between client and y

Summary

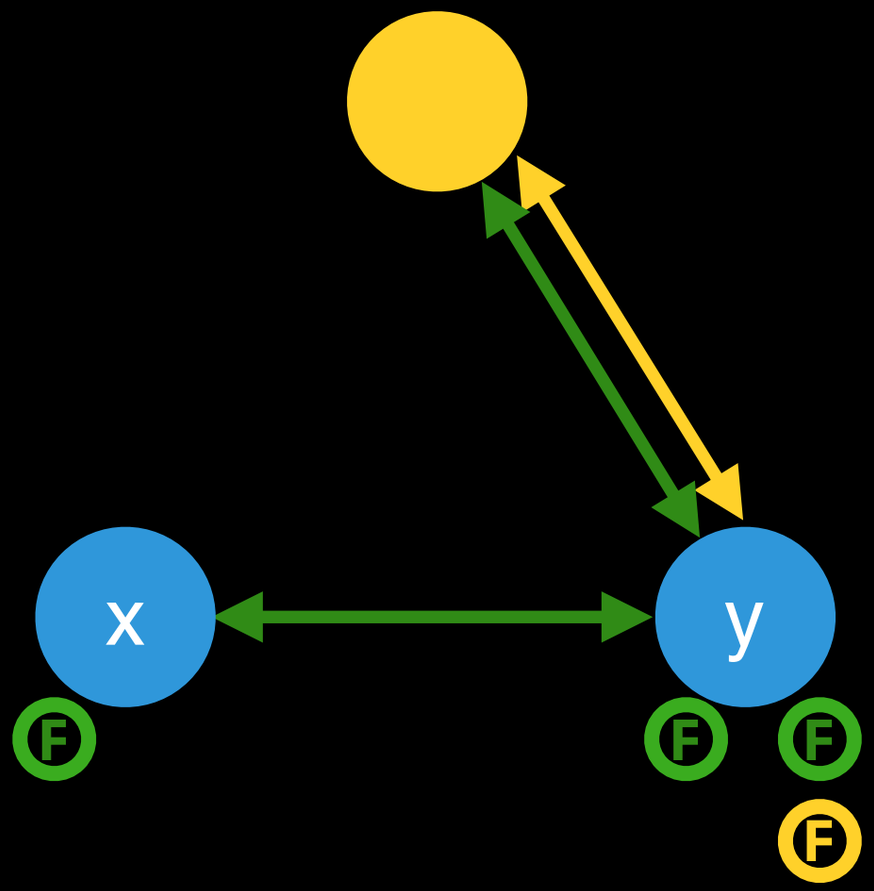

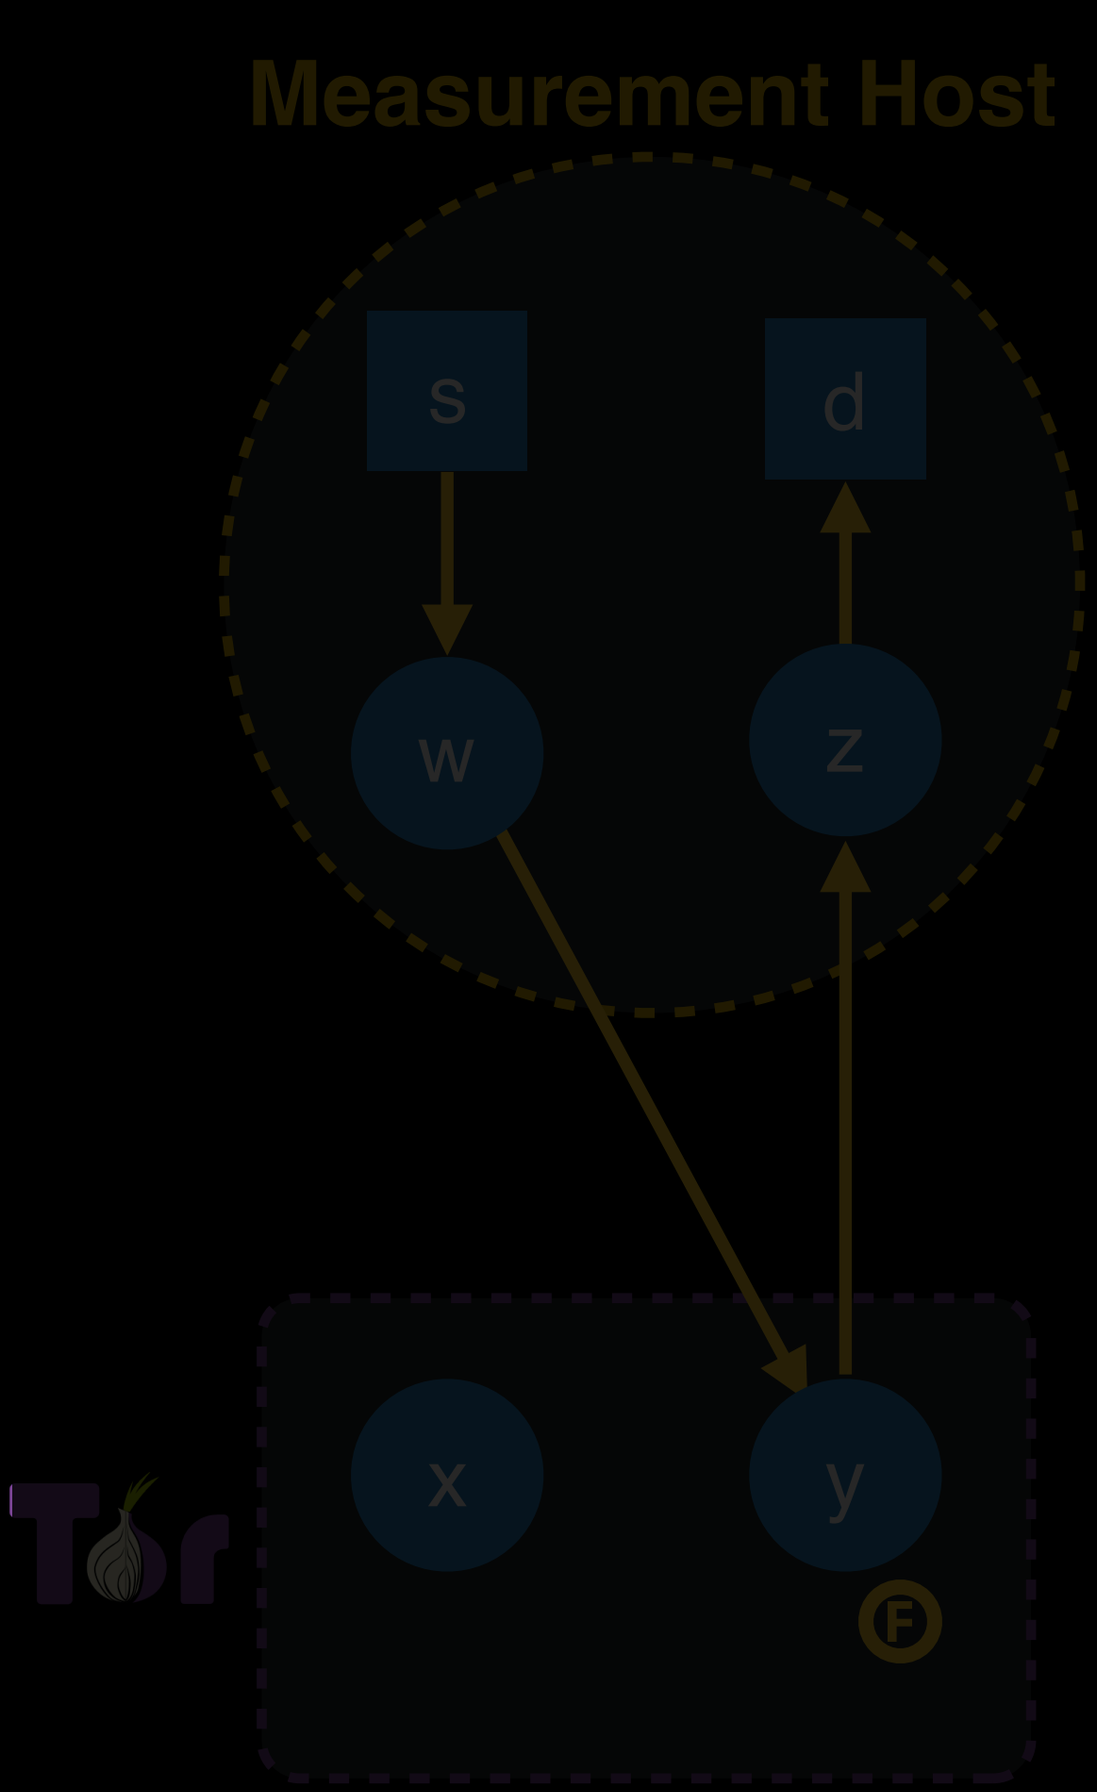

1 – 2 – 3

Isolate **RTT** between client and y

Summary

Subtract latencies

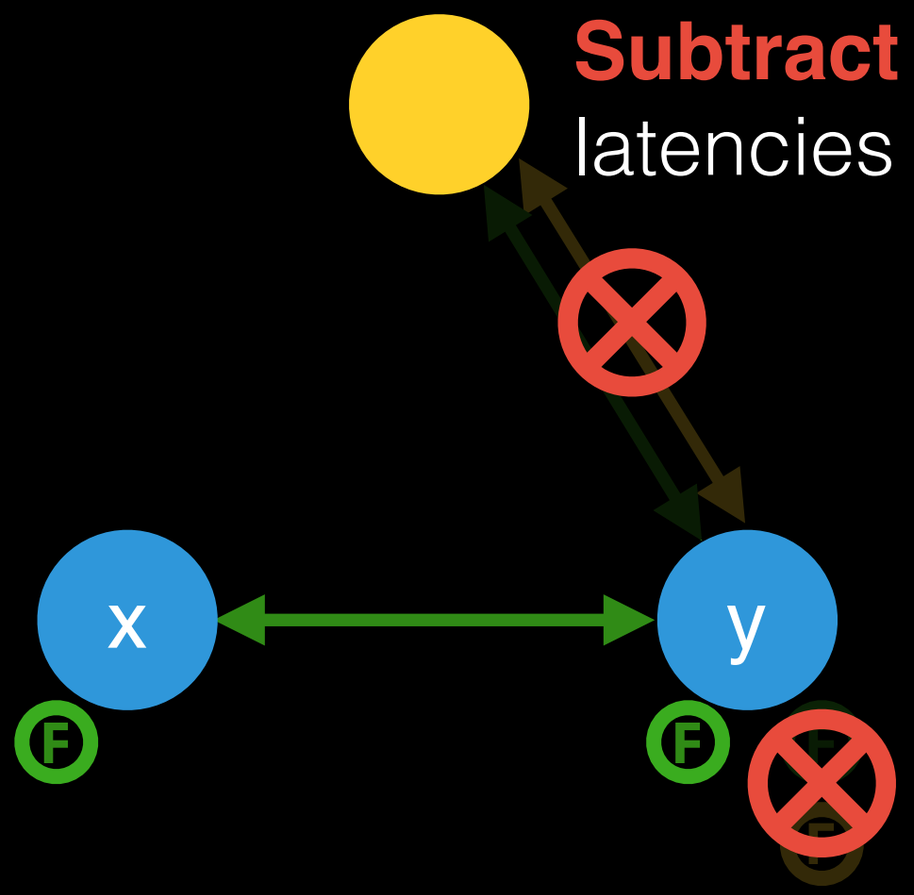

1 – 2 – 3

Isolate **RTT** between client and y

Summary

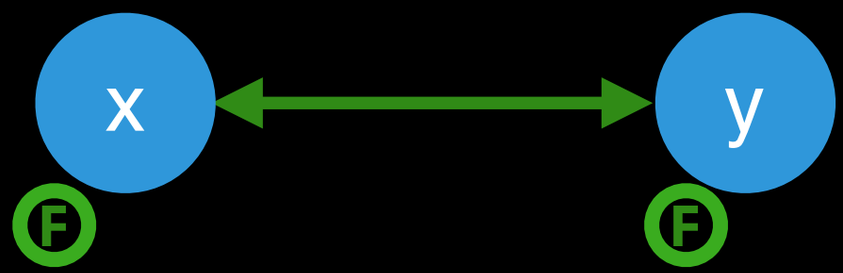

1 – 2 – 3 = RTT(x,y) + Fx + Fy

Summary

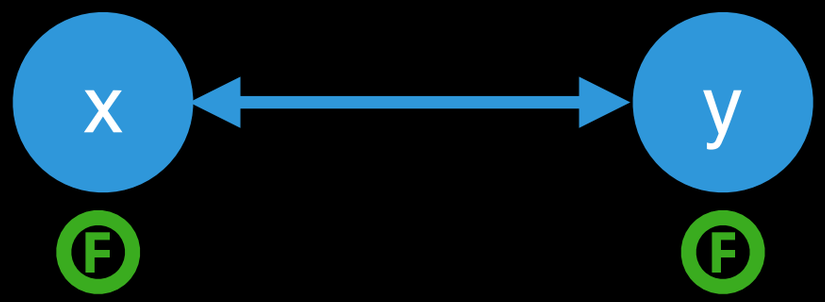

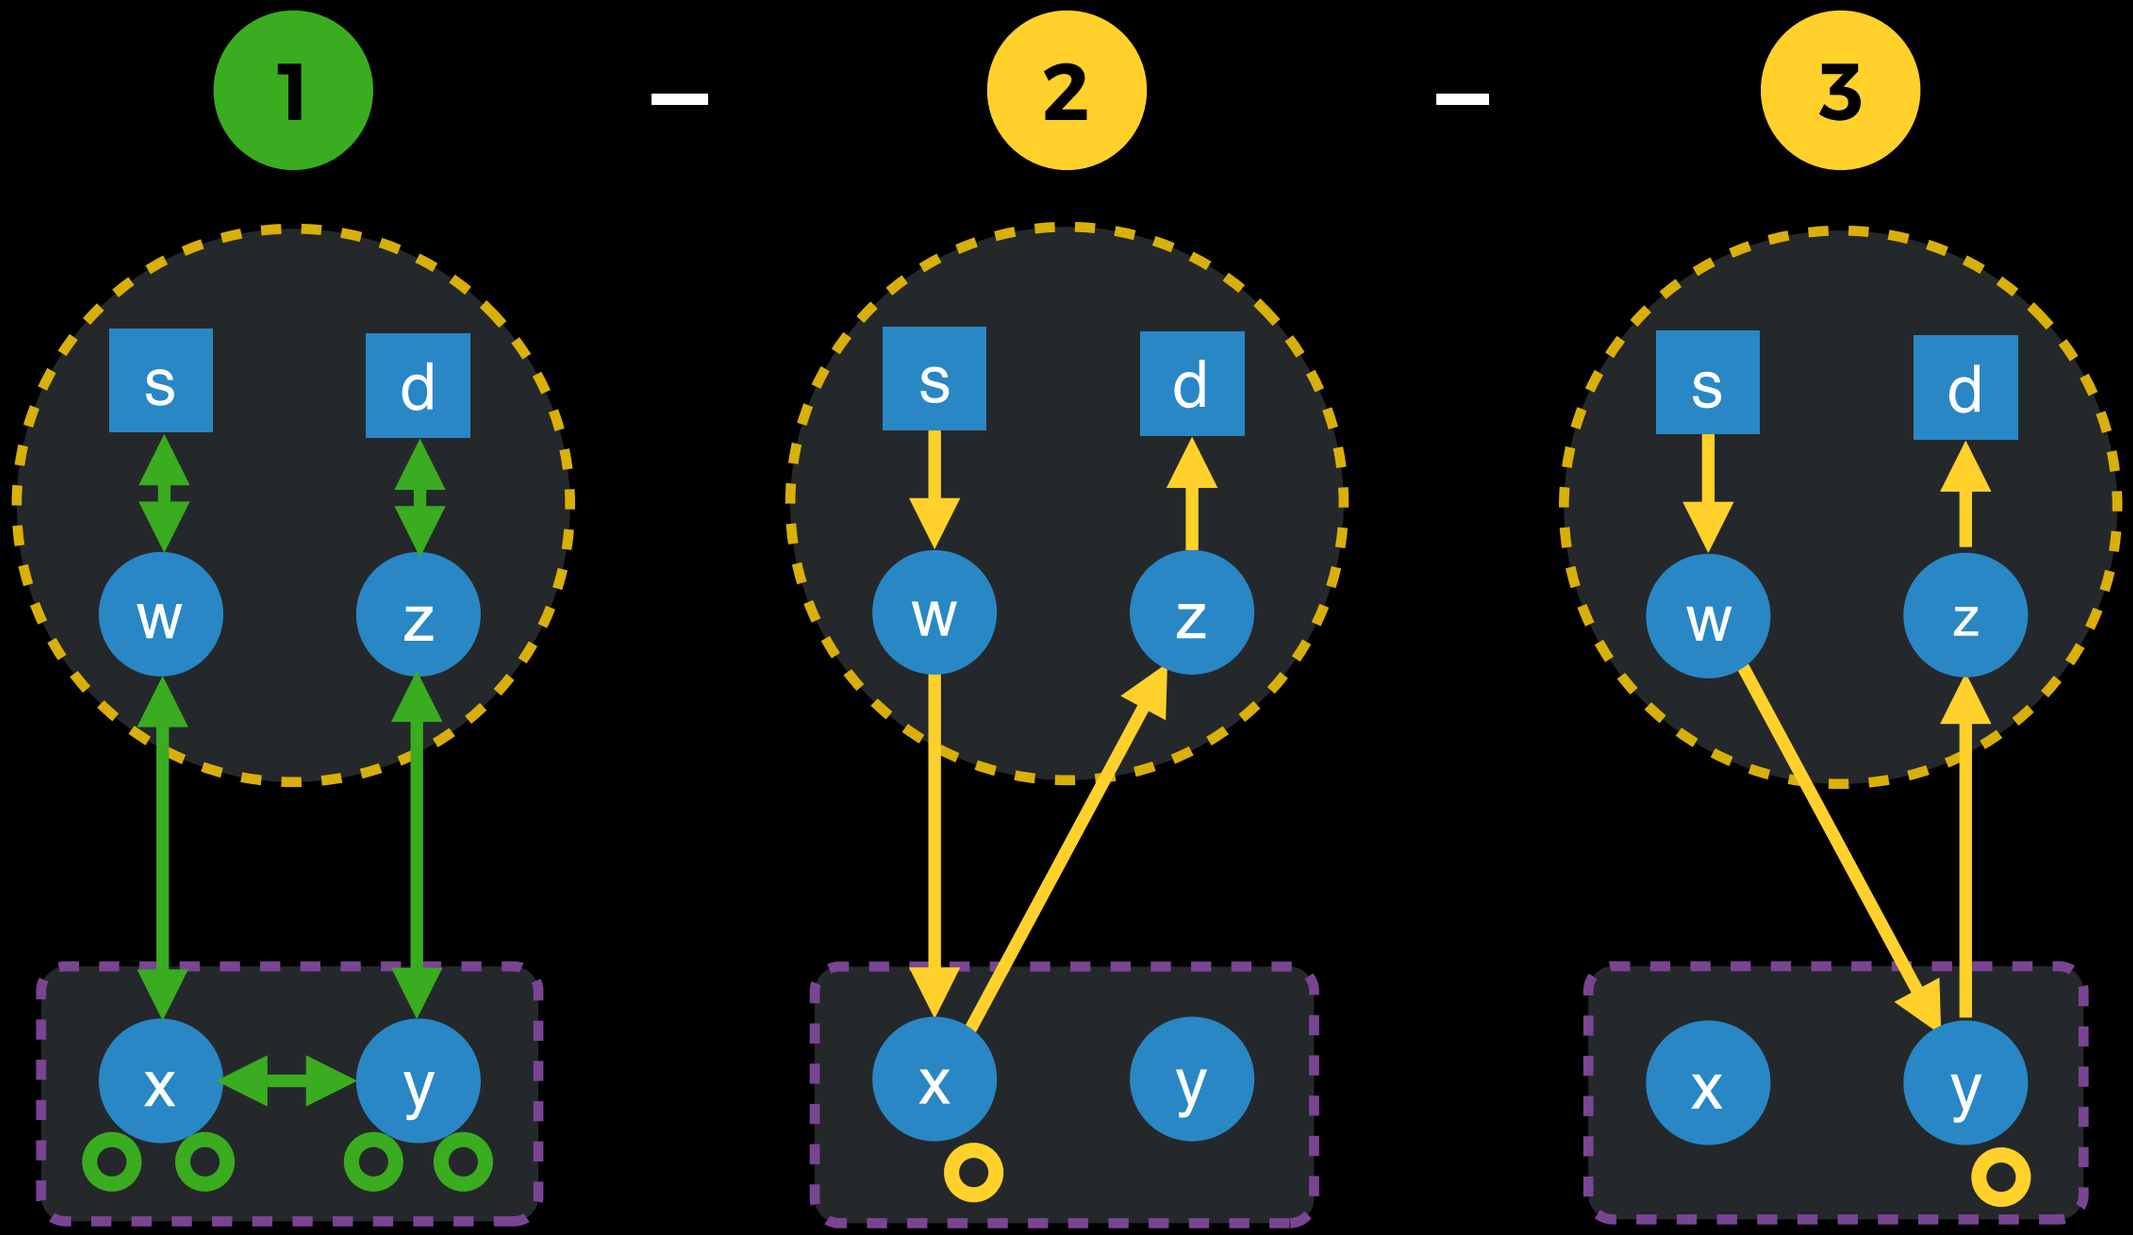
1
2
3
s
d
w
z
x
y
s
d
w
z
x
y
s
d
w
z
x
y

min (X * 1) − min (X * 2) − min (X * 3)

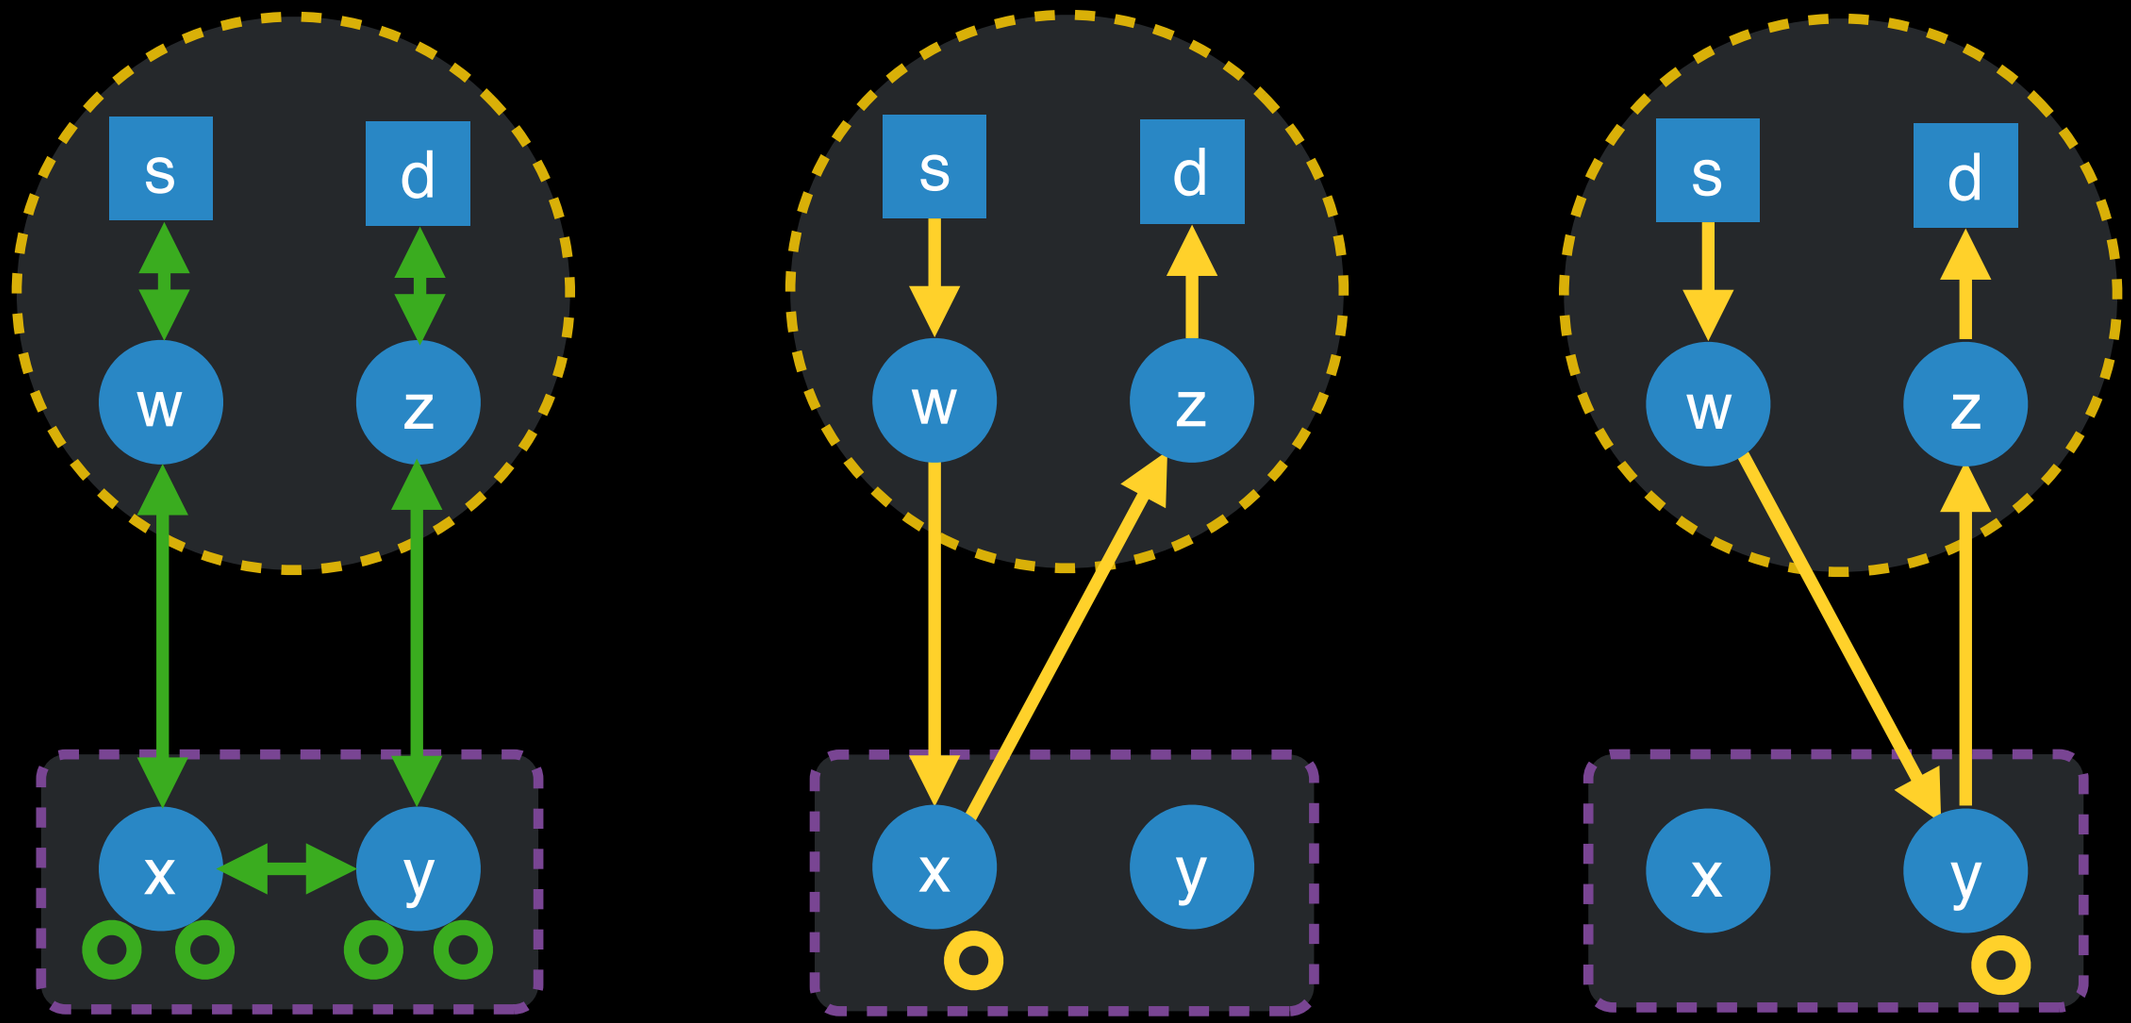

Minimum of multiple, independent samples of each circuit

Ting evaluation

Implemented Ting using the Stem Tor controller

No modifications to the Tor client

How well does Ting work?

- How accurate?
- How many samples?
- How consistent?

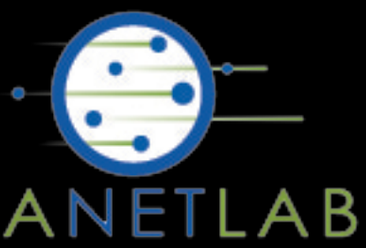

Ting evaluation

How well does Ting work?

- How accurate?
- How many samples?
- How consistent?

PLANETLAB

Tor

31 Tor relays

930 Total pairs

How accurate is Ting?

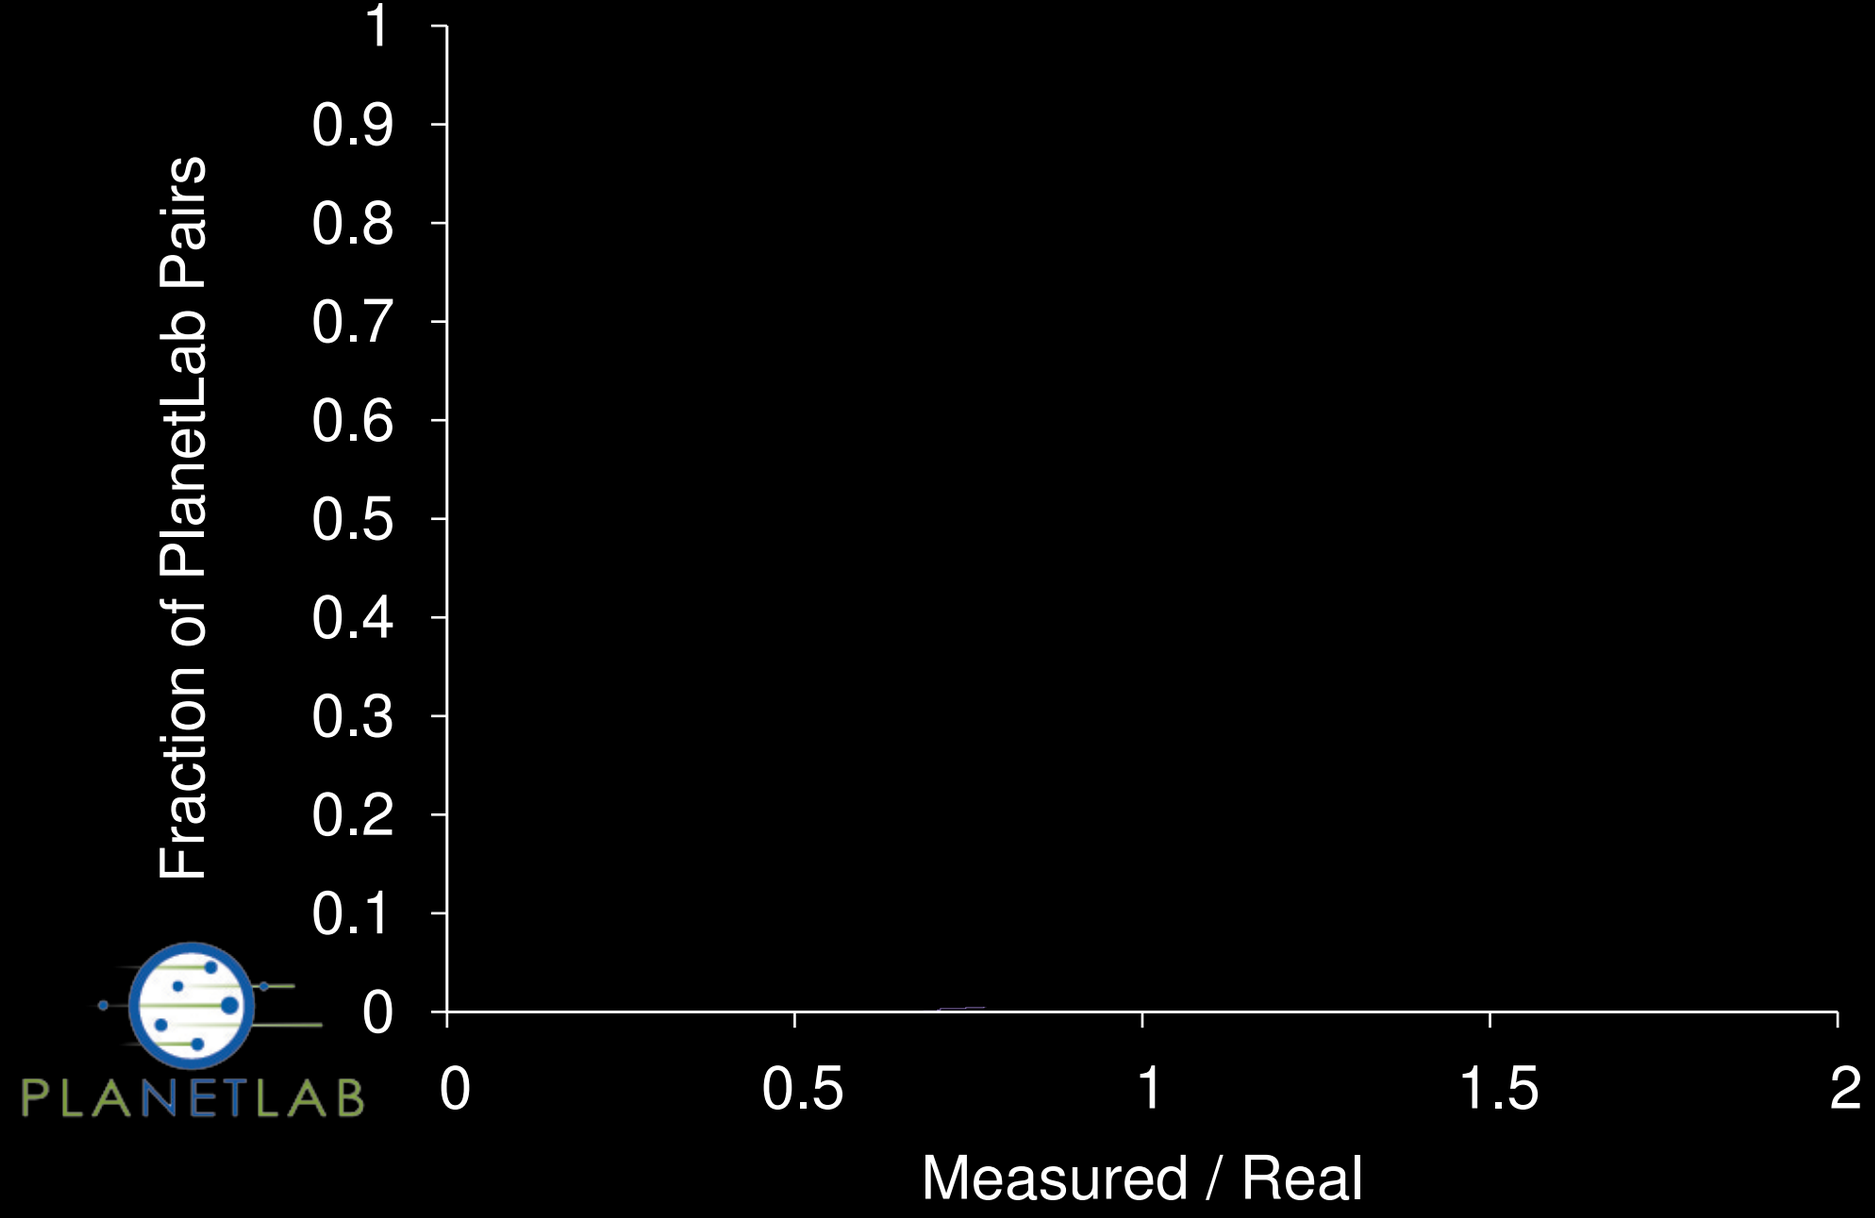

How accurate is Ting?

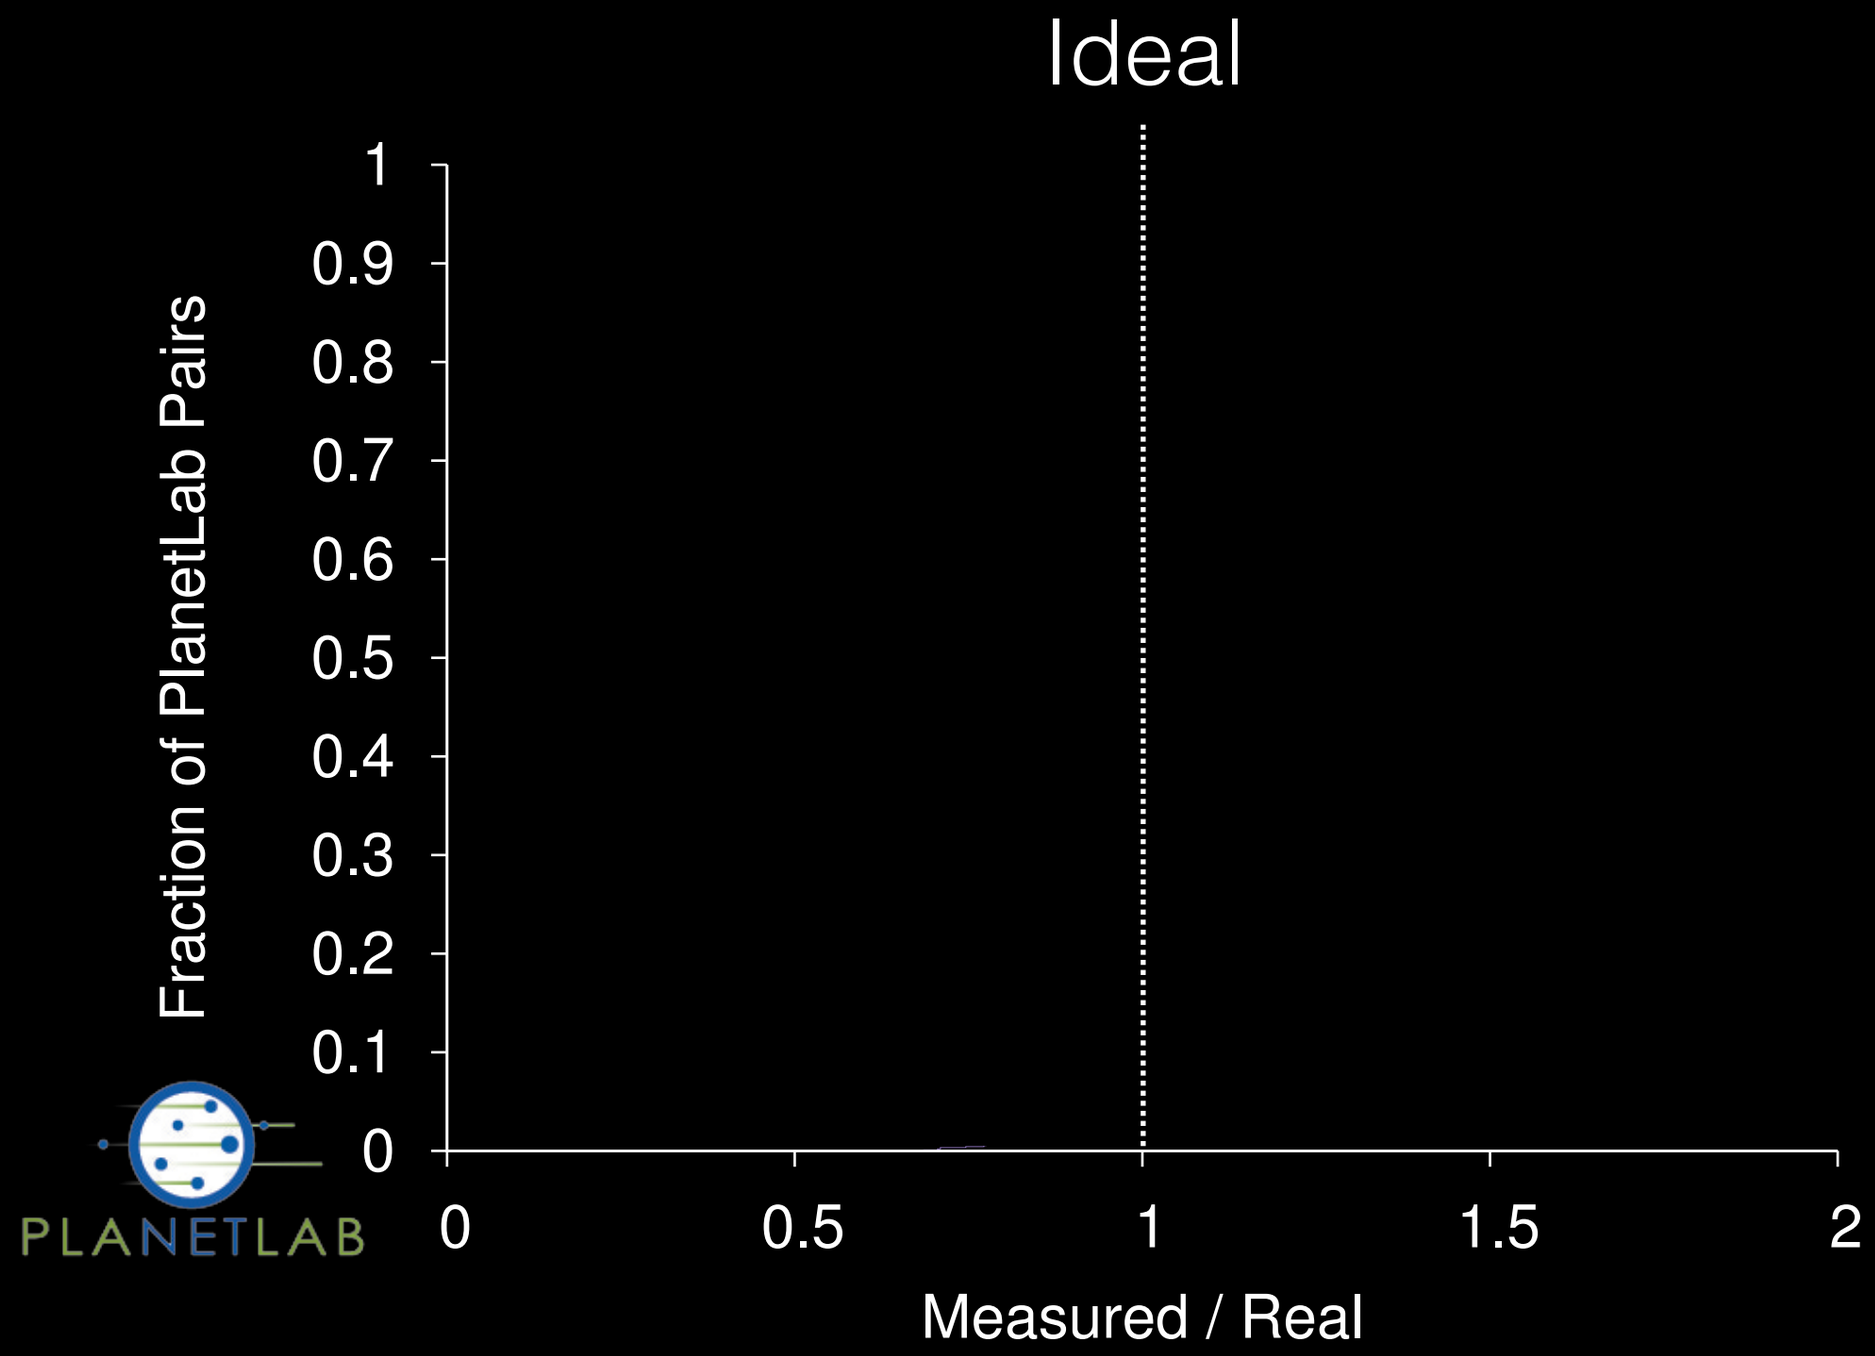

How accurate is Ting?

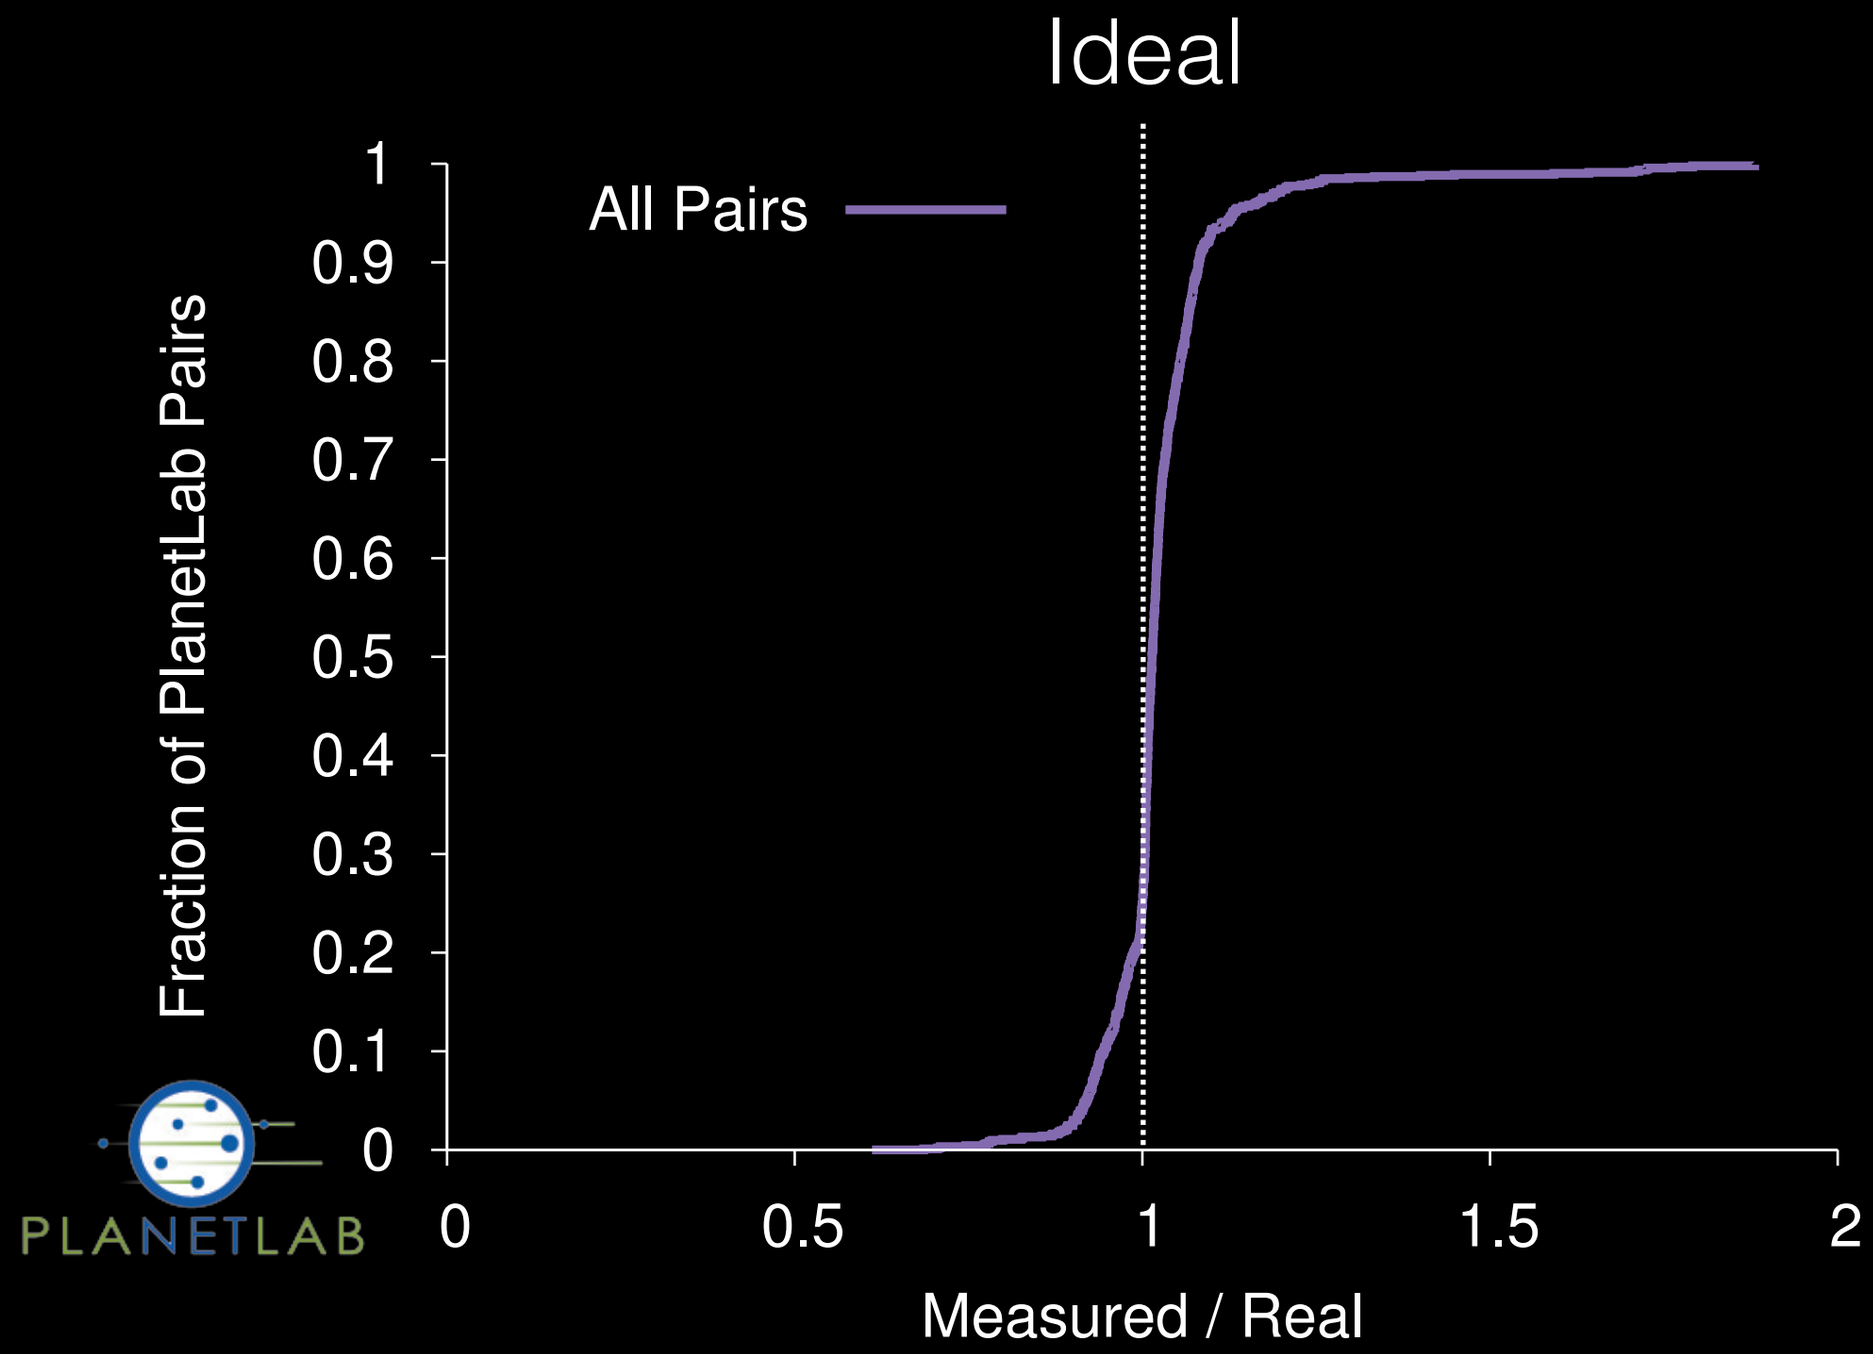

How accurate is Ting?

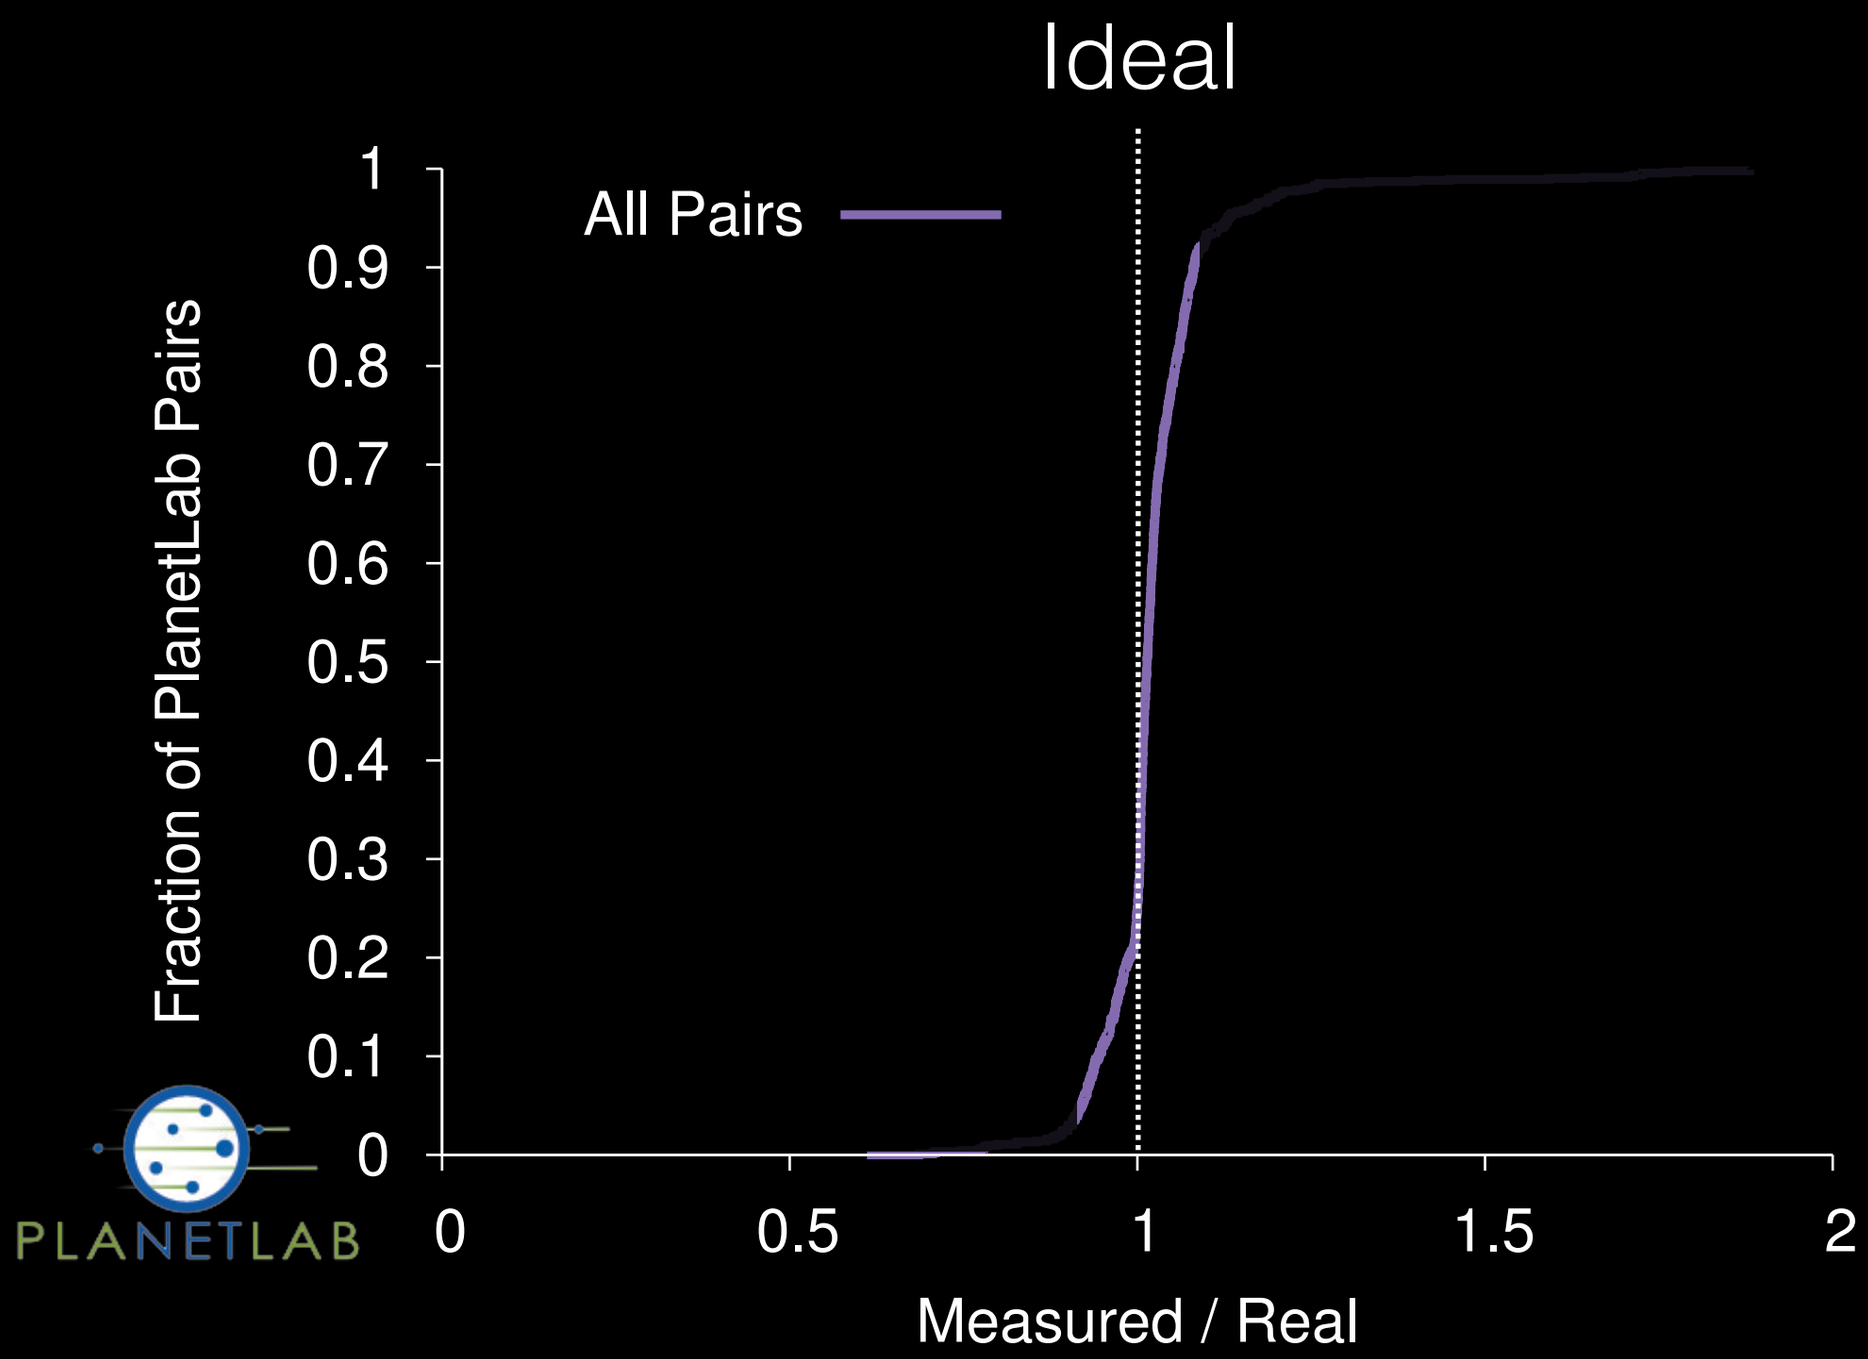

Ting estimates typically within 10% of “real”

How accurate is Ting?

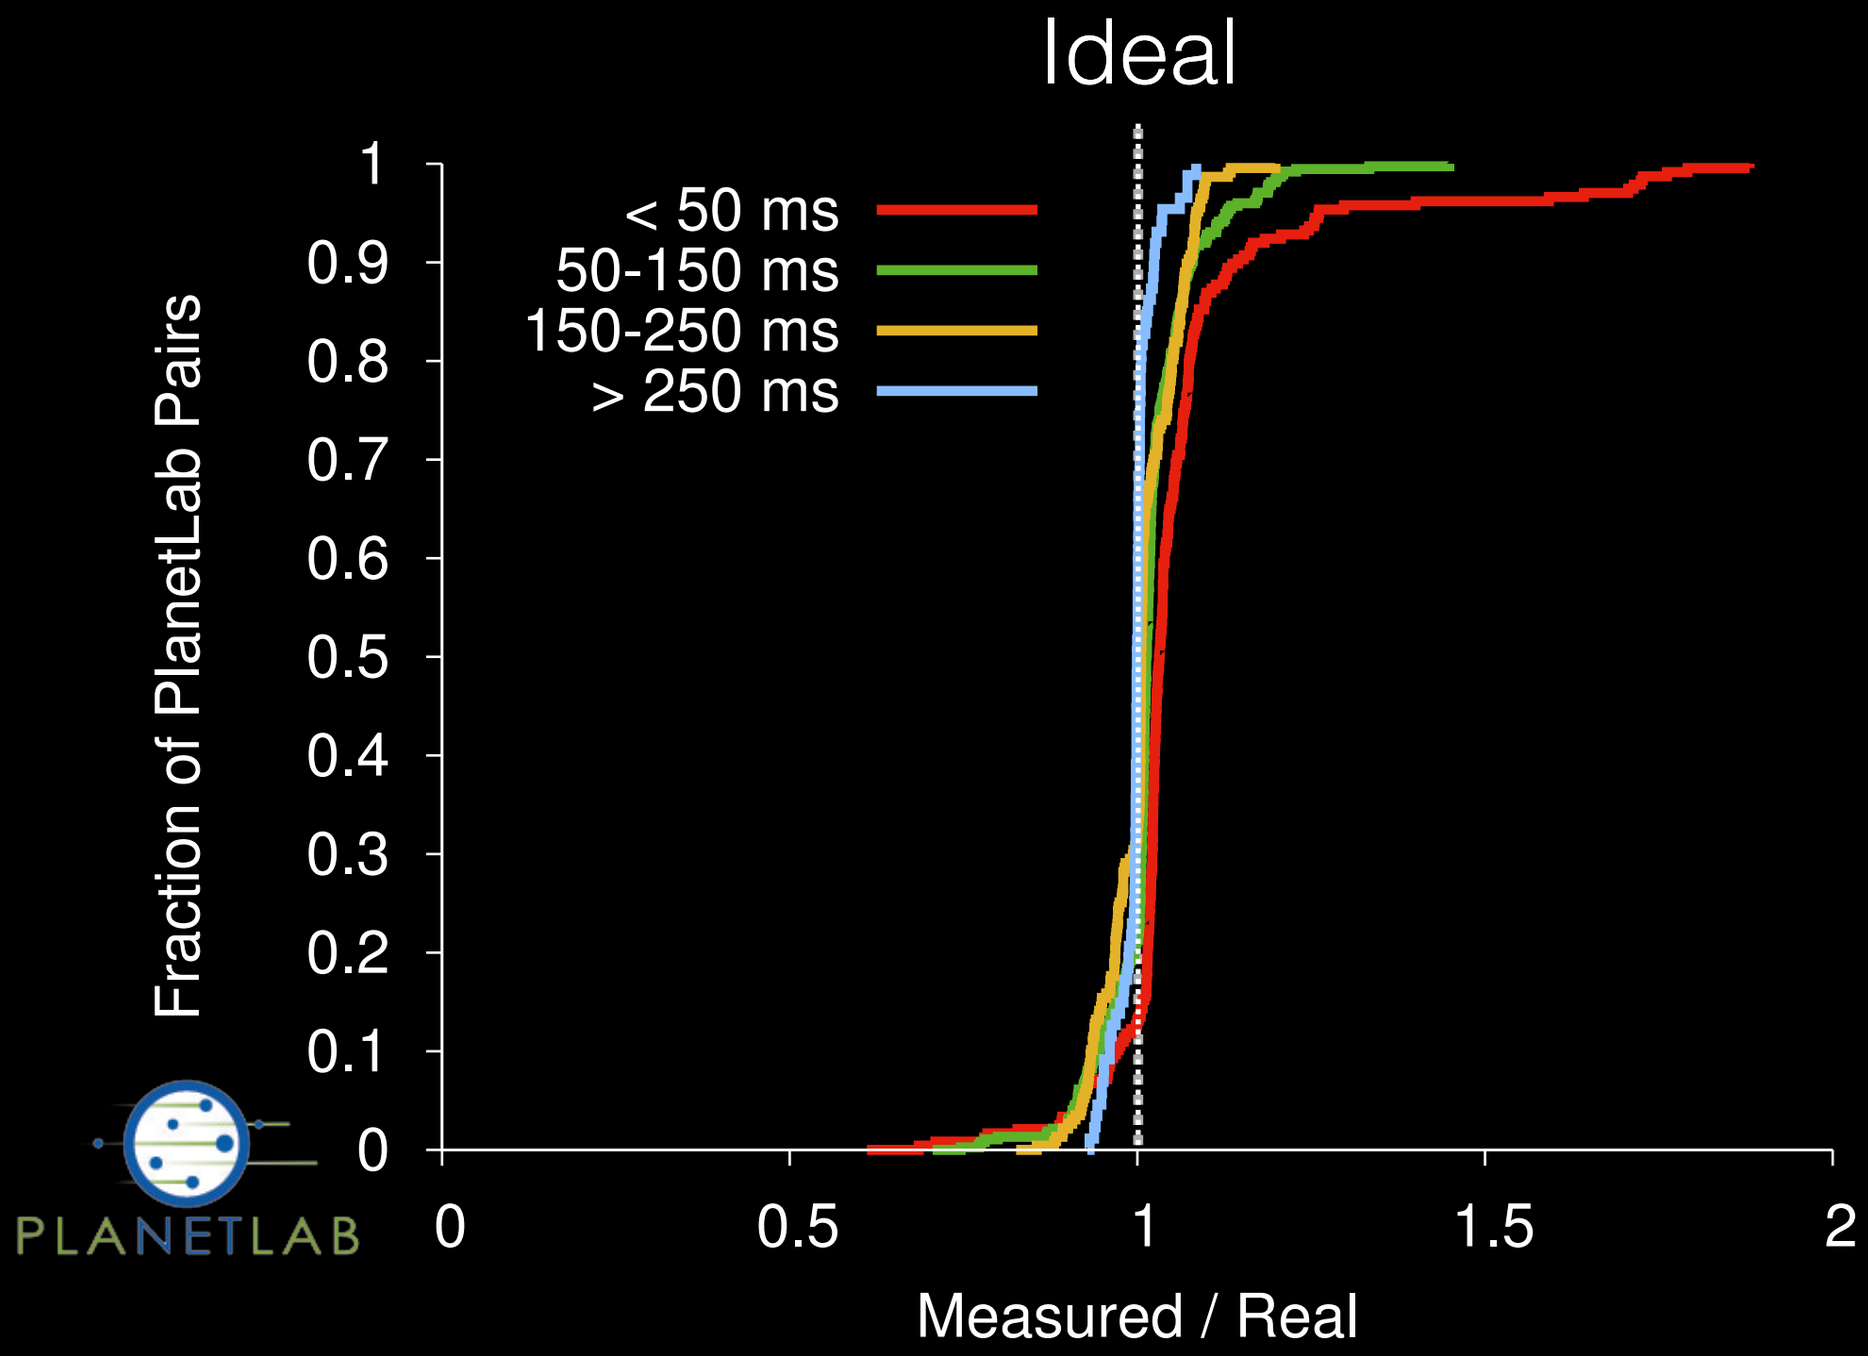

Ting estimates typically within 10% of “real”

How accurate is Ting?

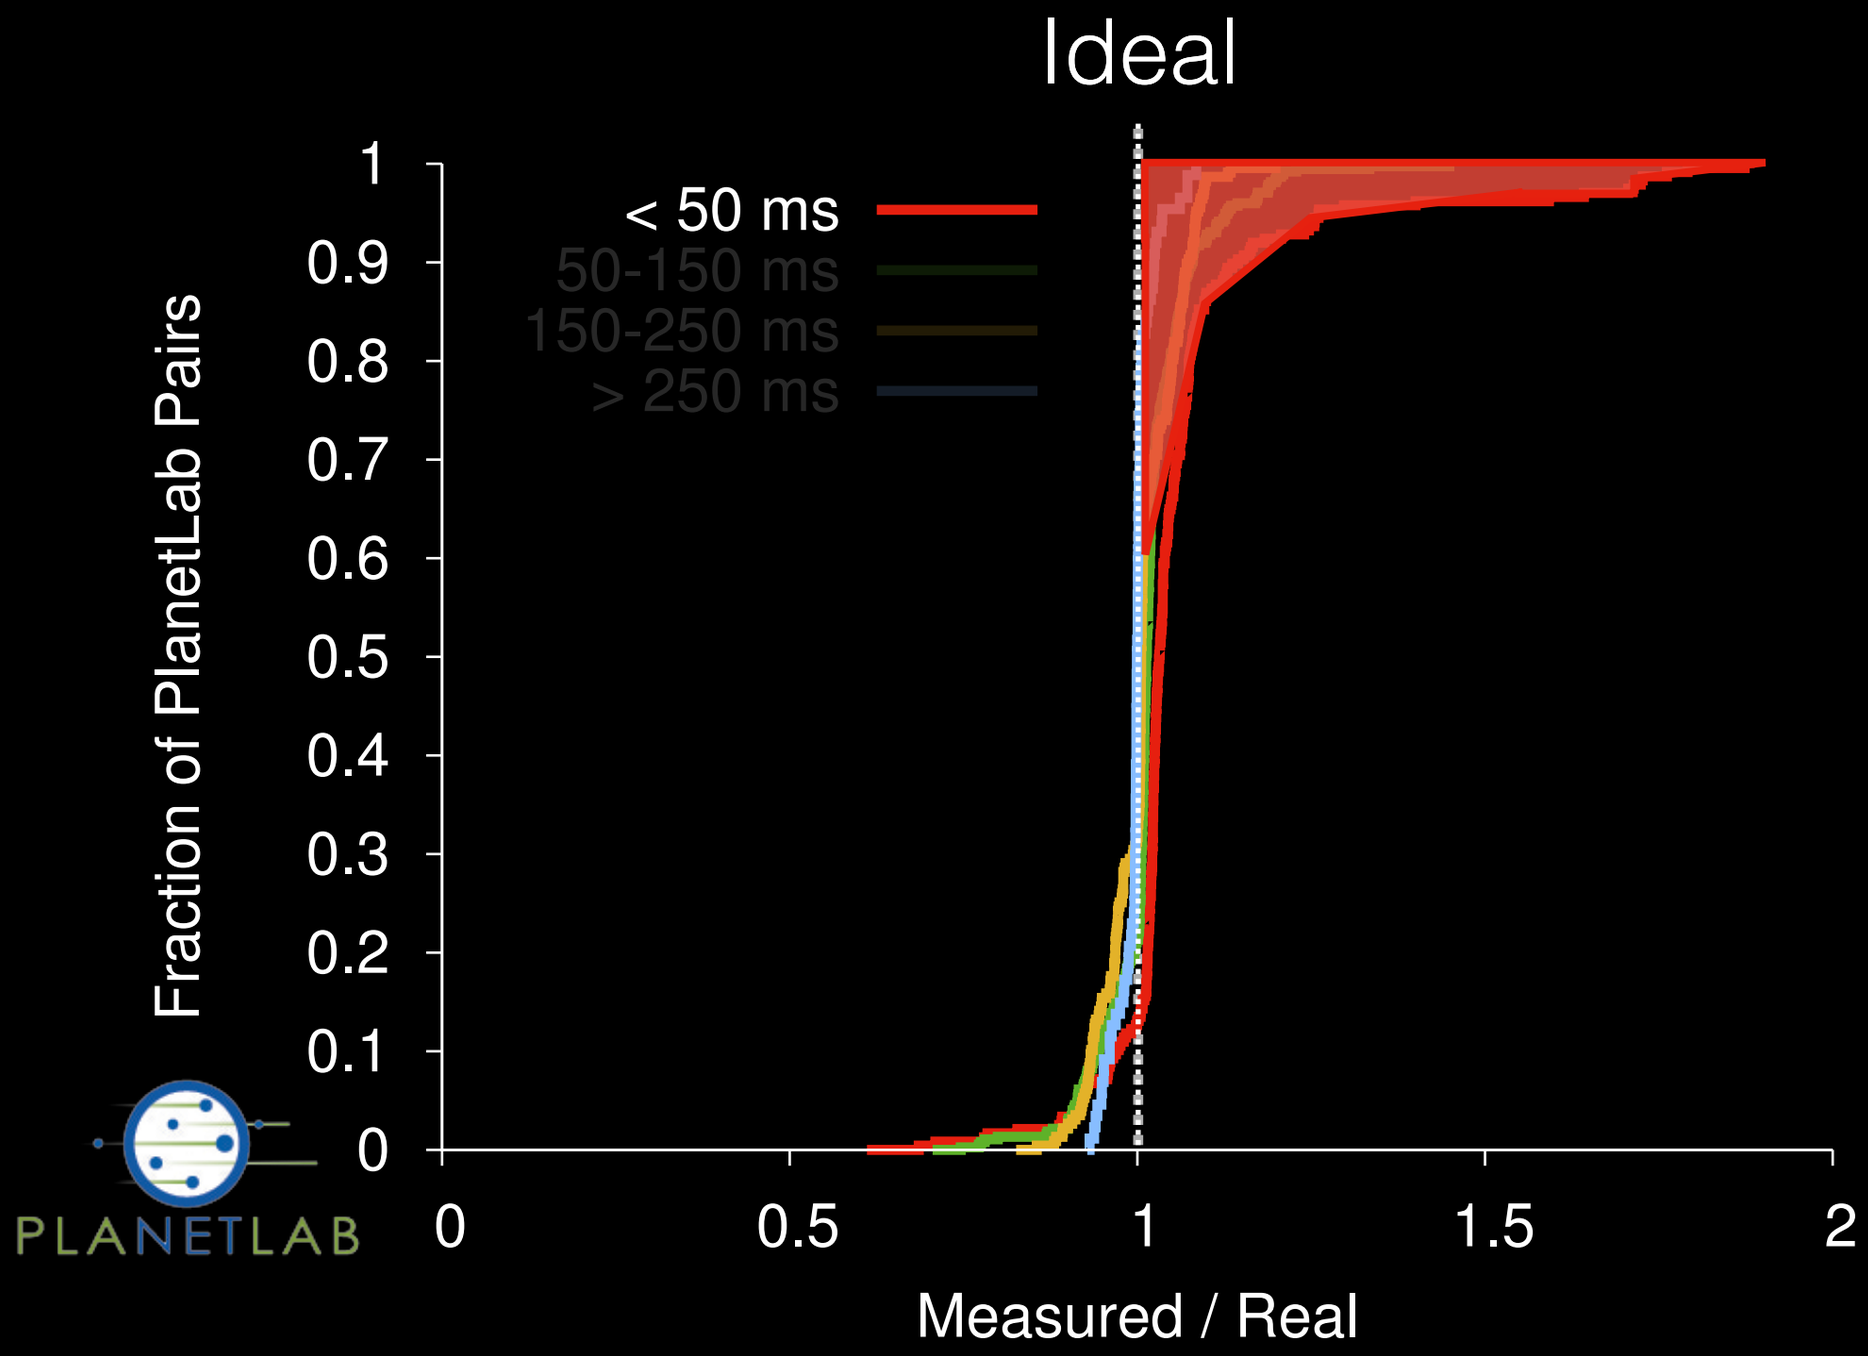

Ting estimates typically within 10% of “real”

How accurate is Ting?

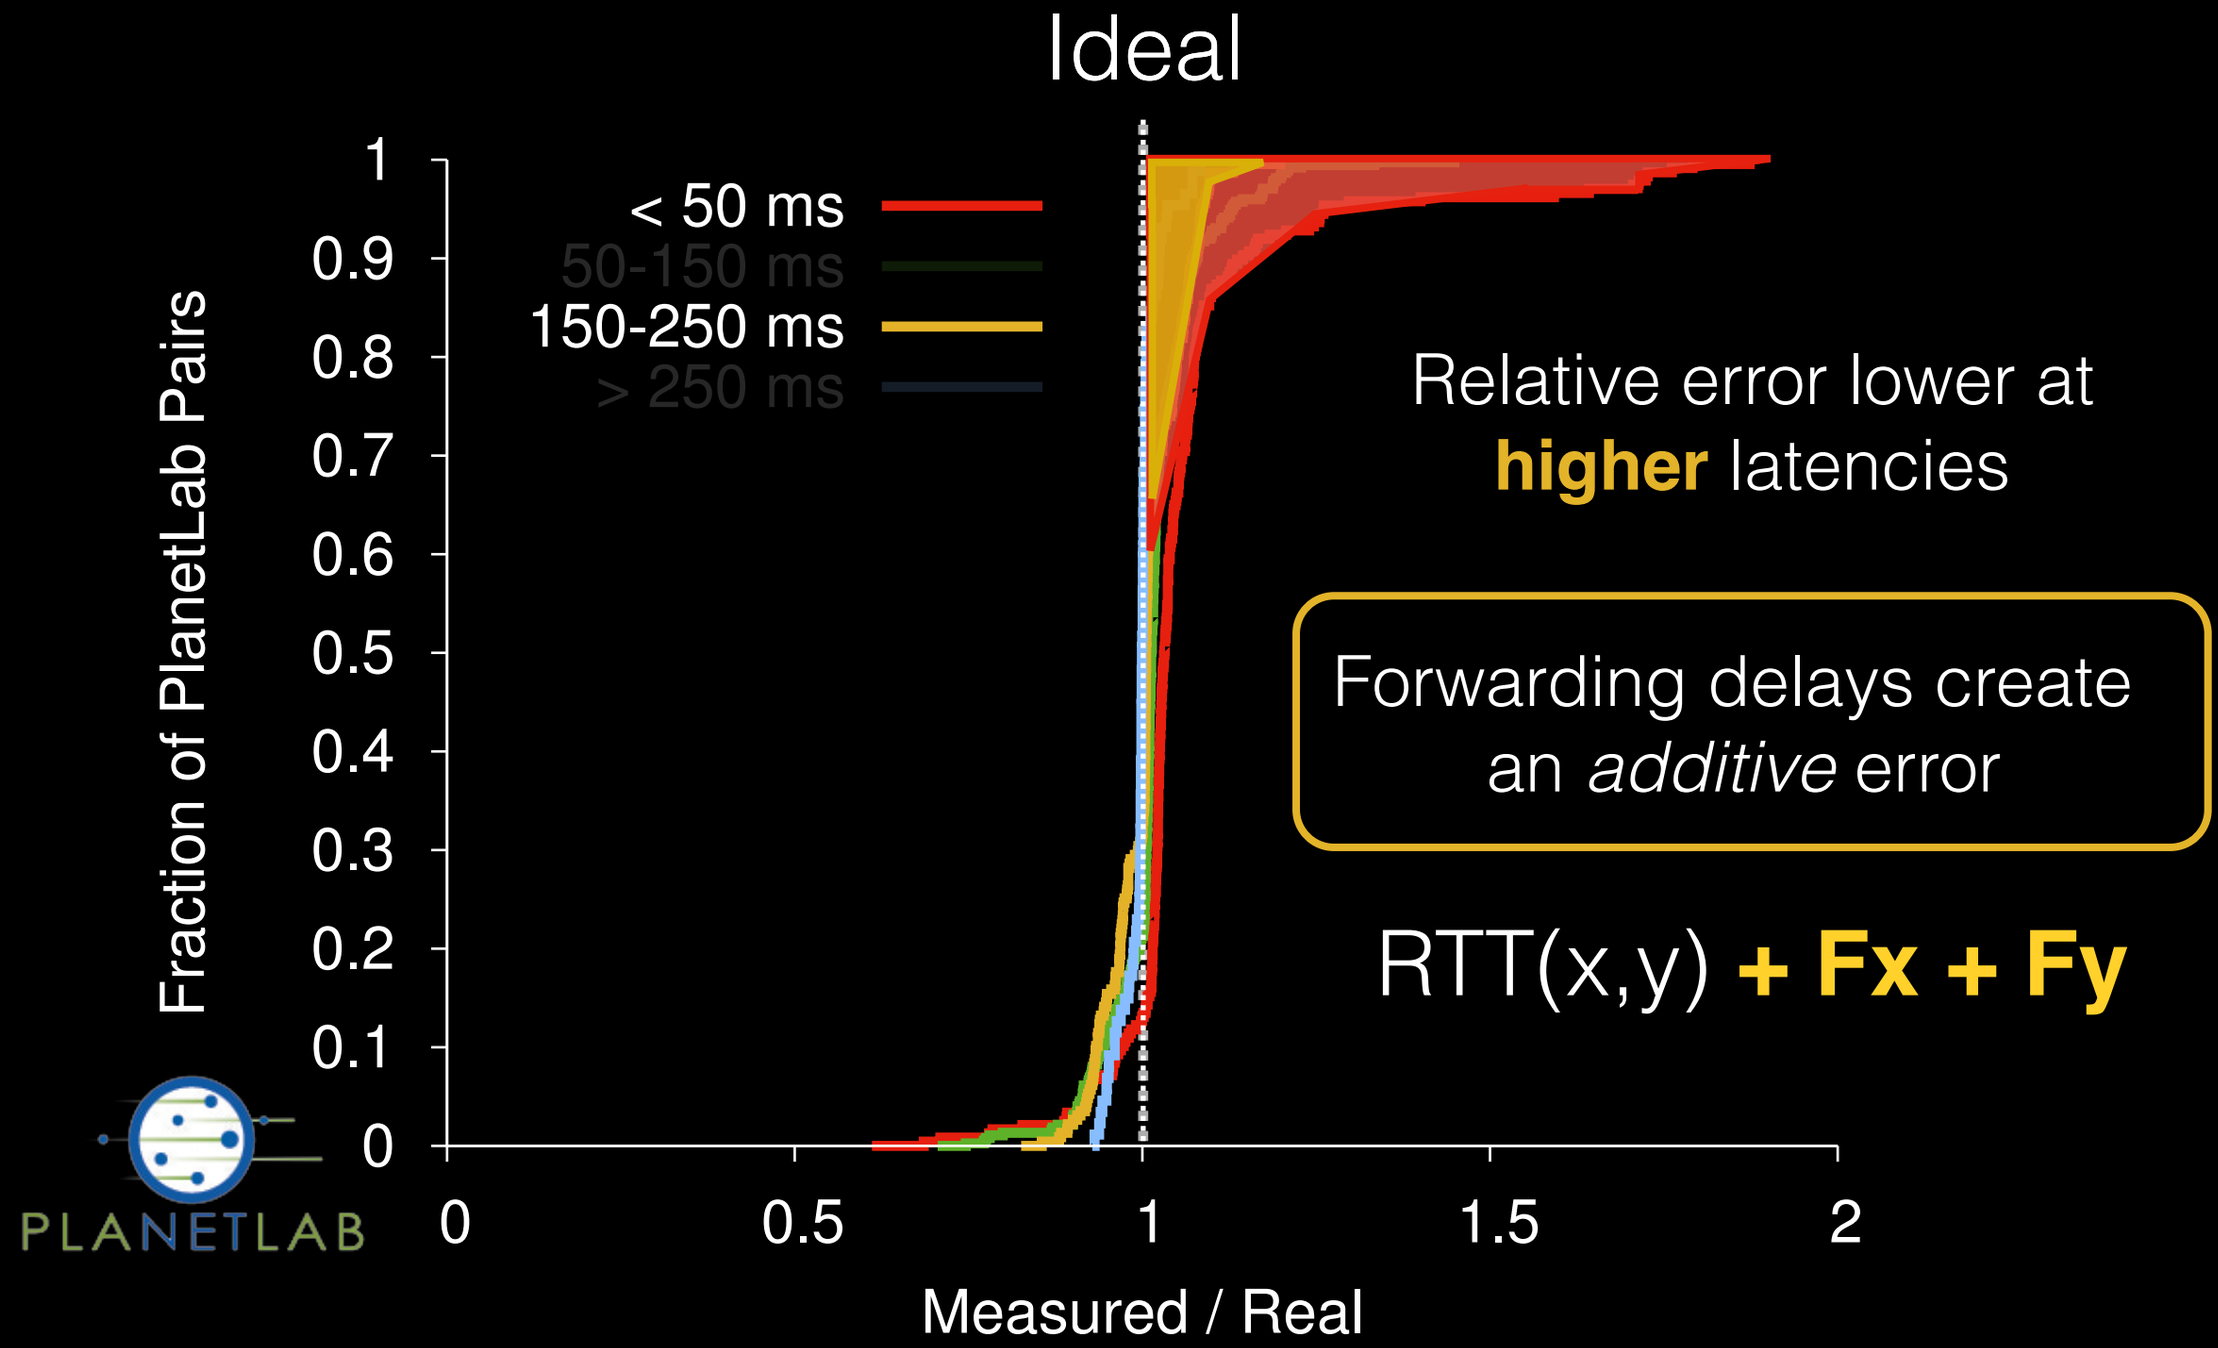

Relative error lower at **higher** latencies

Forwarding delays create an *additive* error

RTT(x,y) **+ Fx + Fy**

Ting estimates typically within 10% of “real”

How many samples does Ting need?

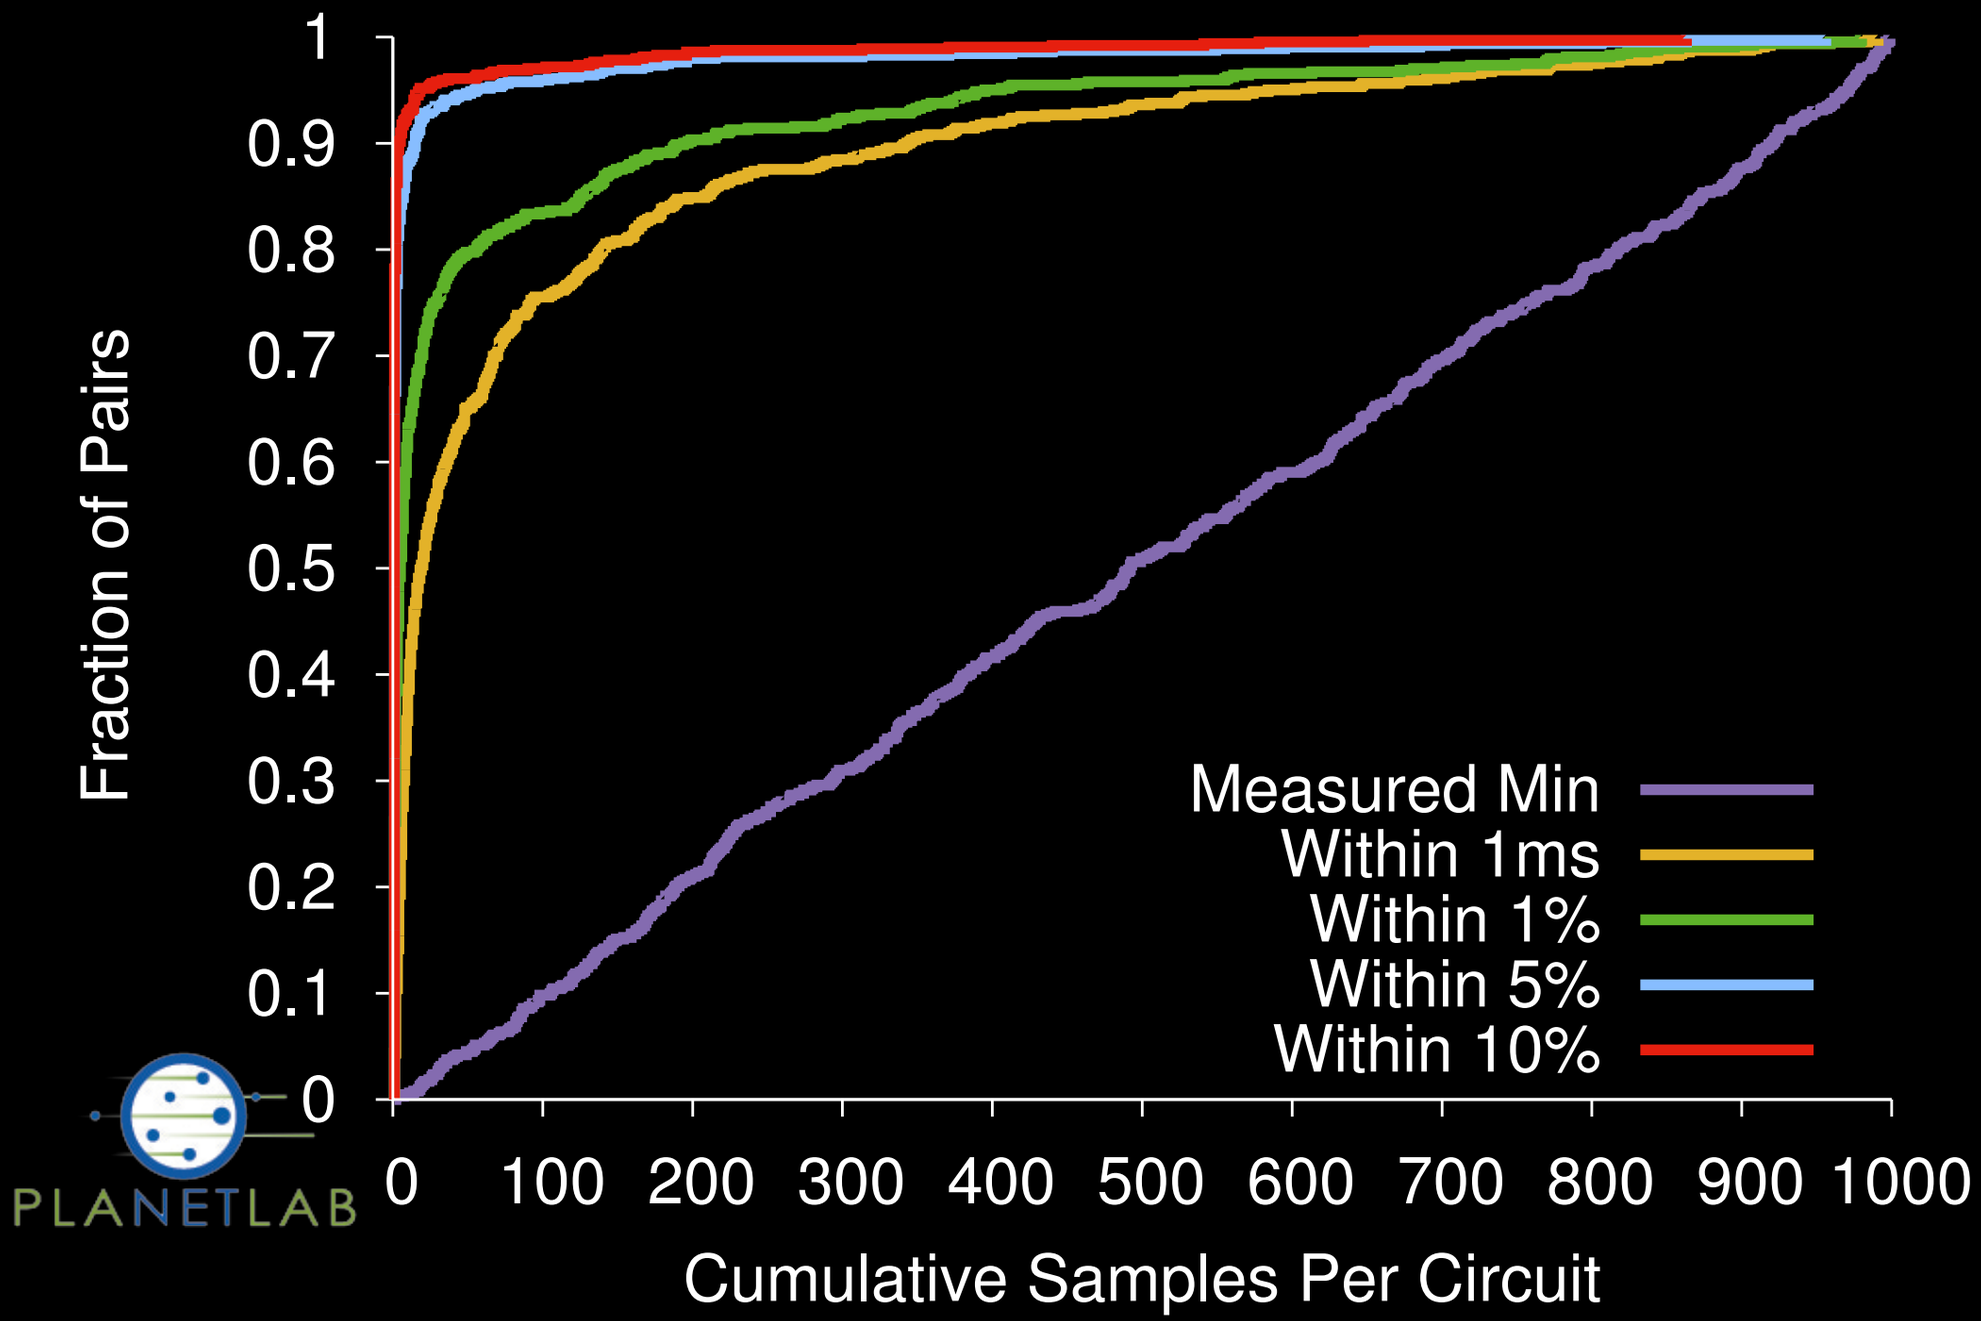

How many samples does Ting need?

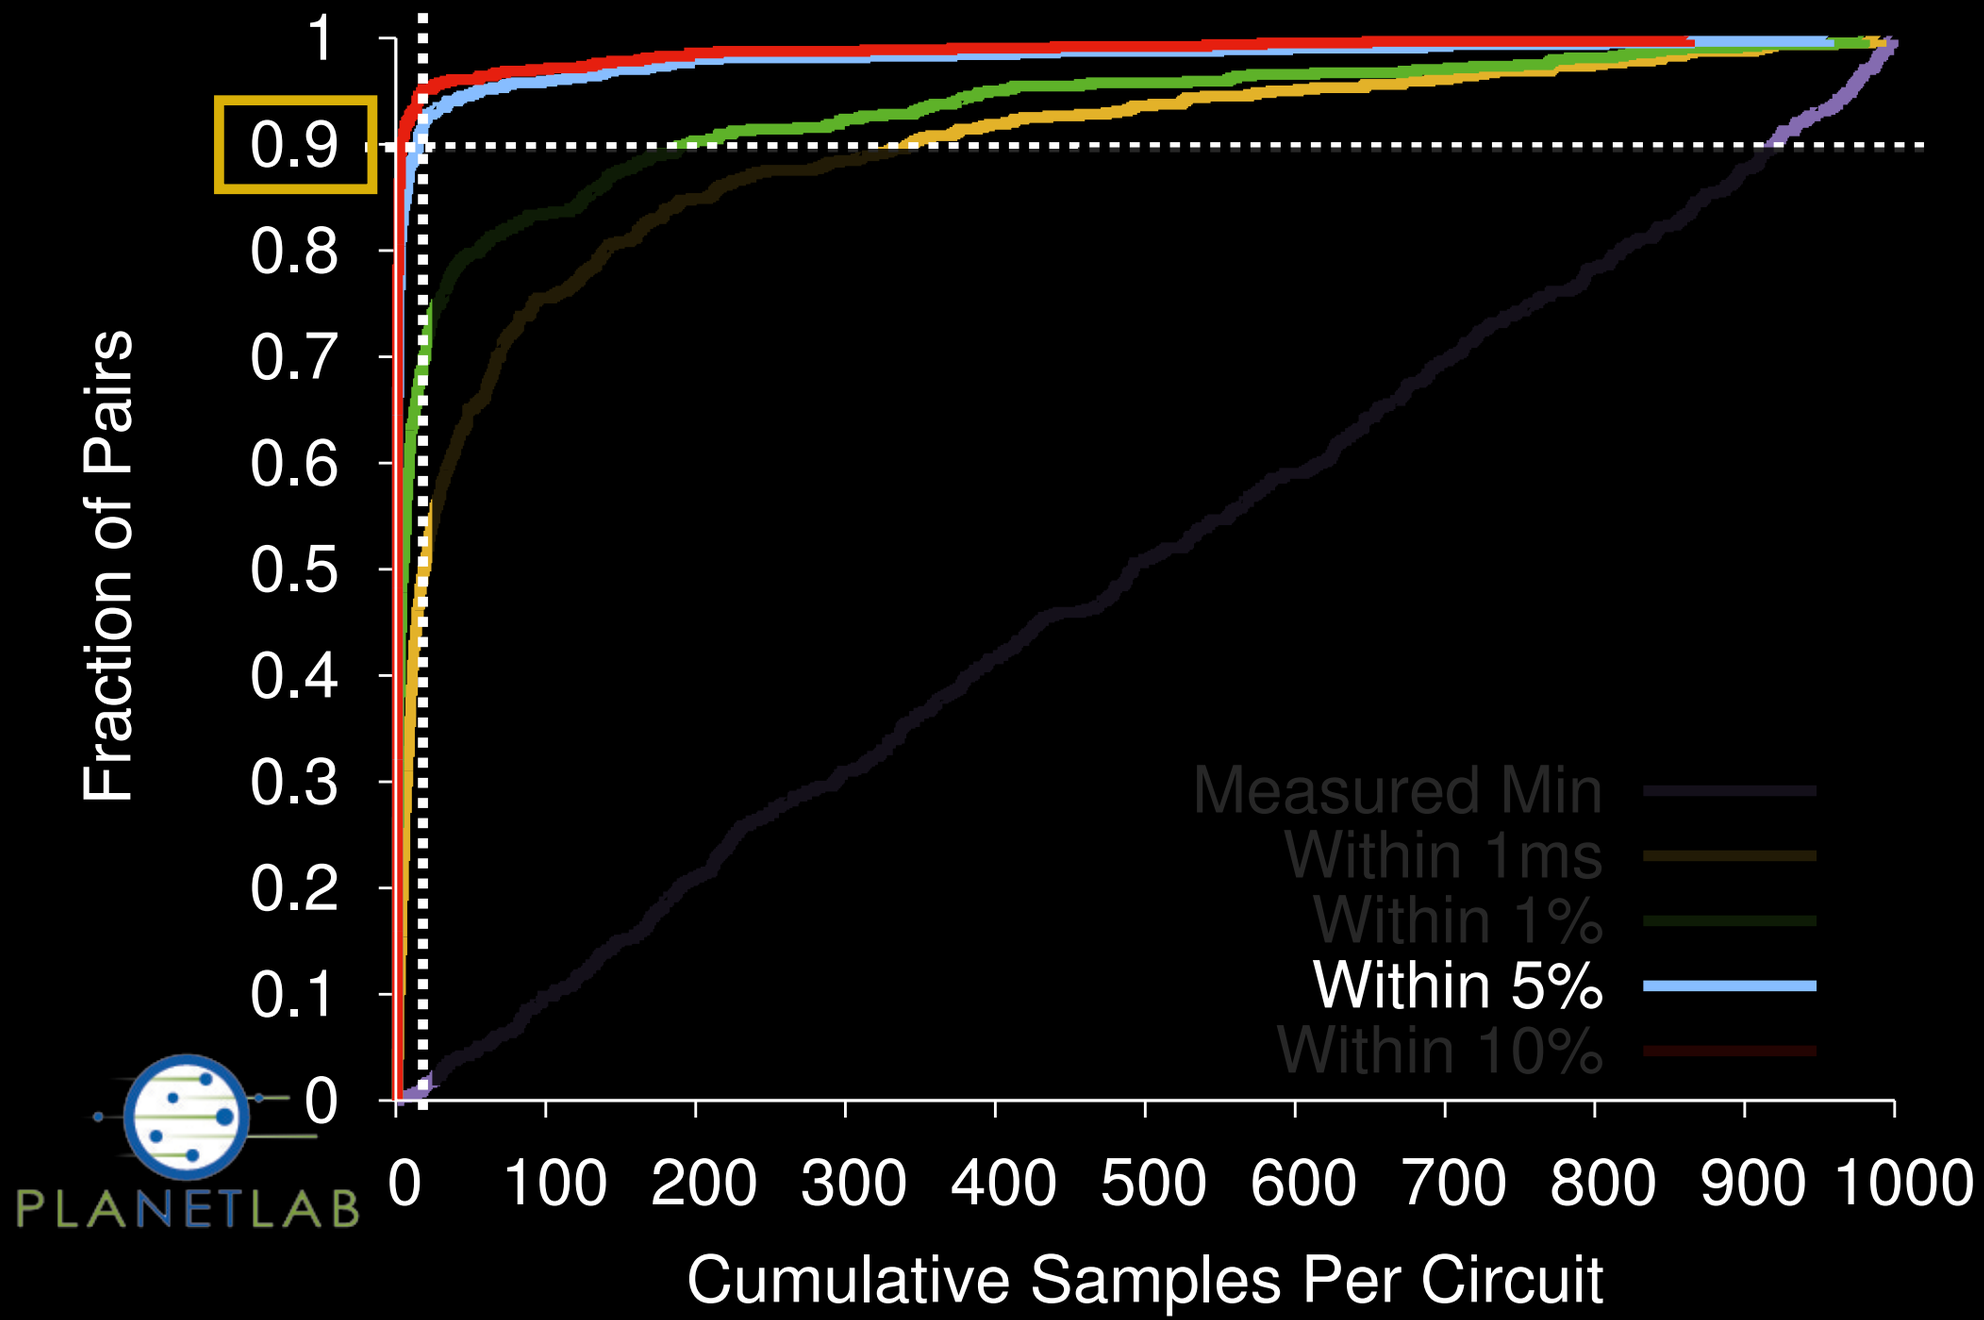

How many samples does Ting need?

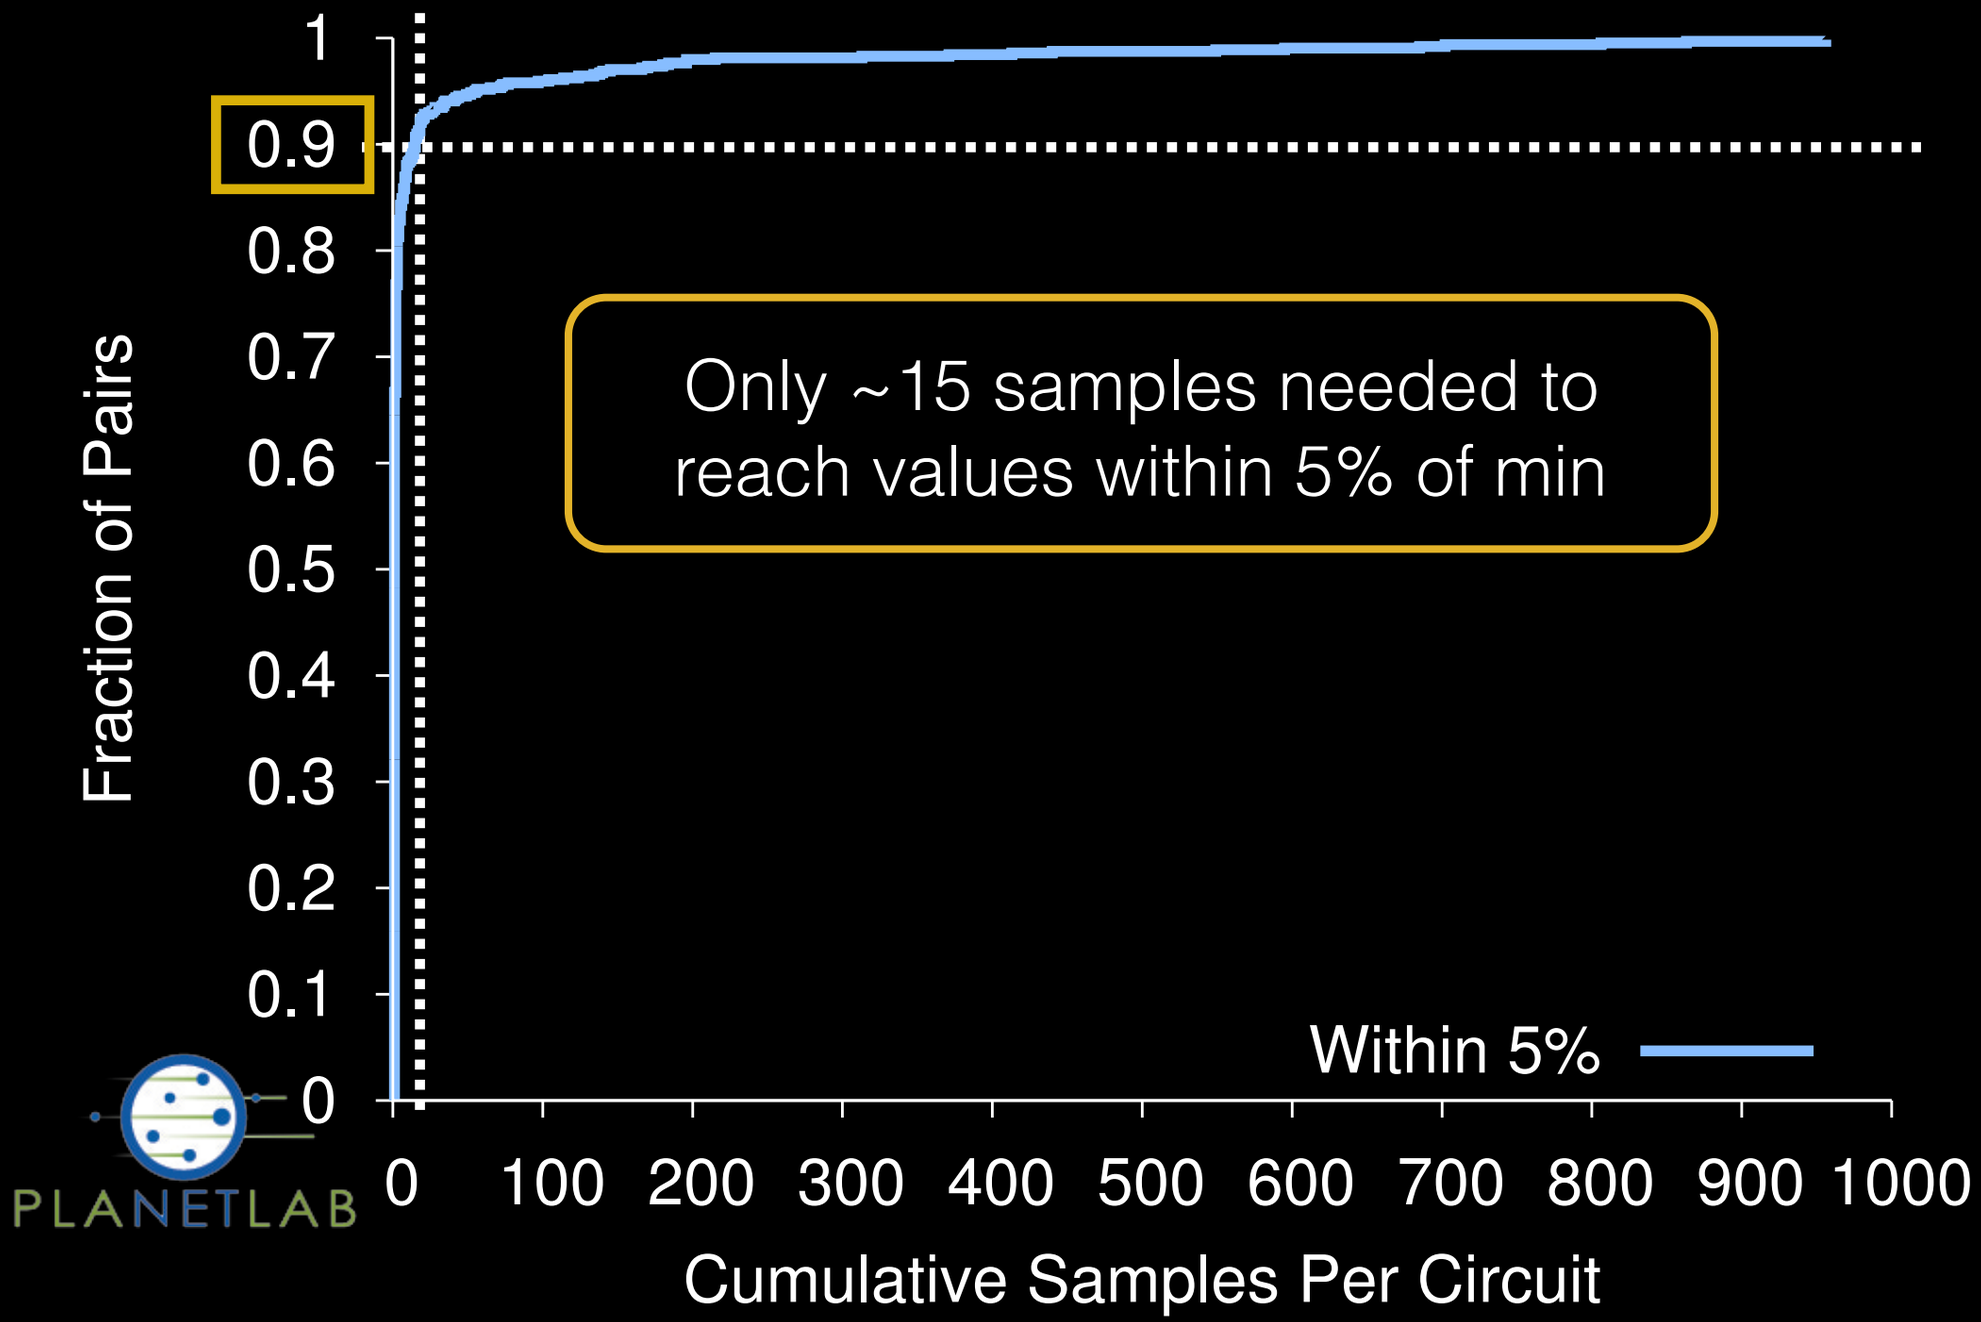

Ting remains accurate with few samples

How consistent are Ting measurements?

30 pairs of real Tor relays, measured **once an hour** over a **week**

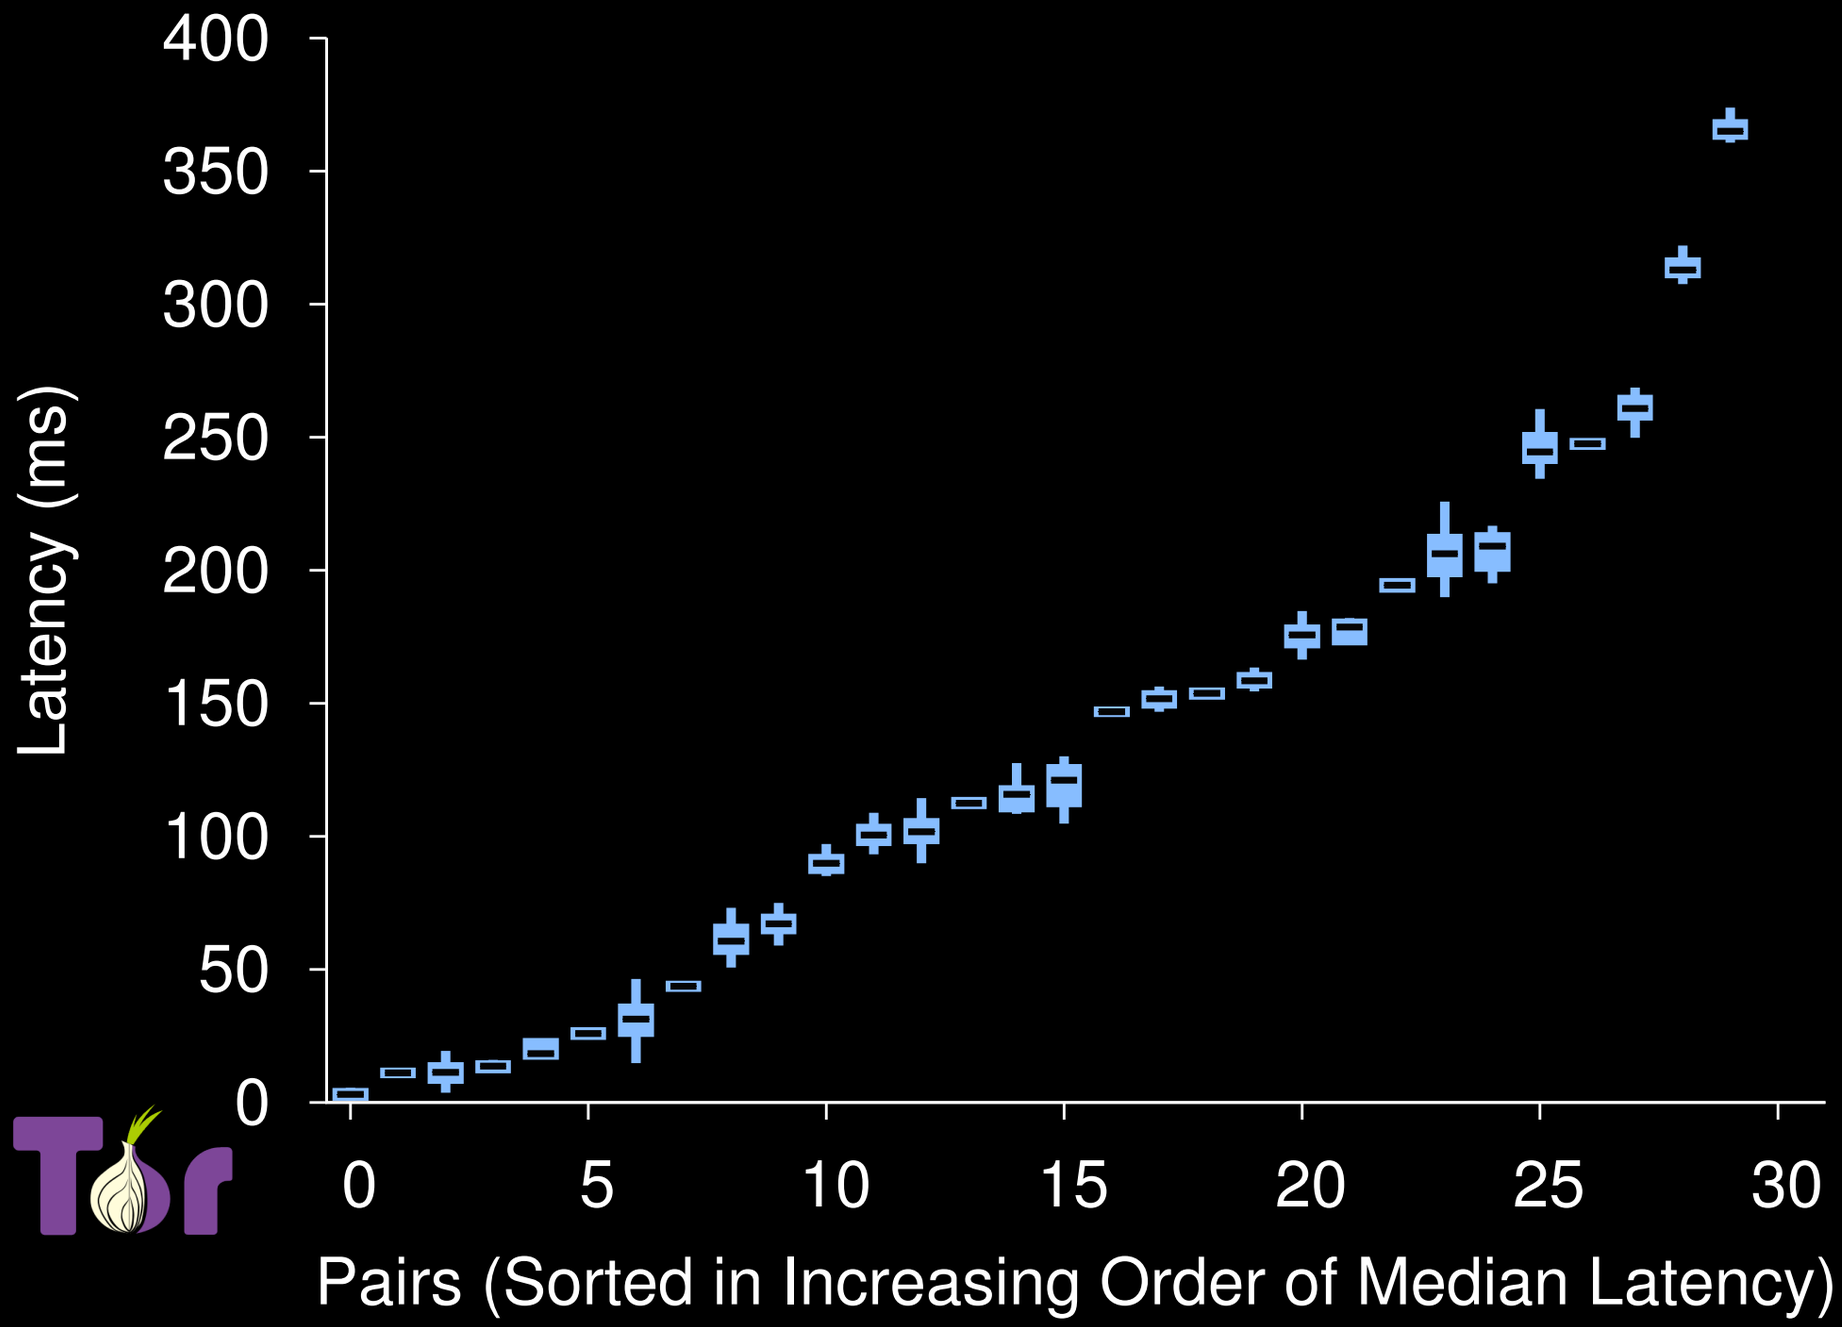

How consistent are Ting measurements?

30 pairs of real Tor relays, measured **once an hour** over a **week**

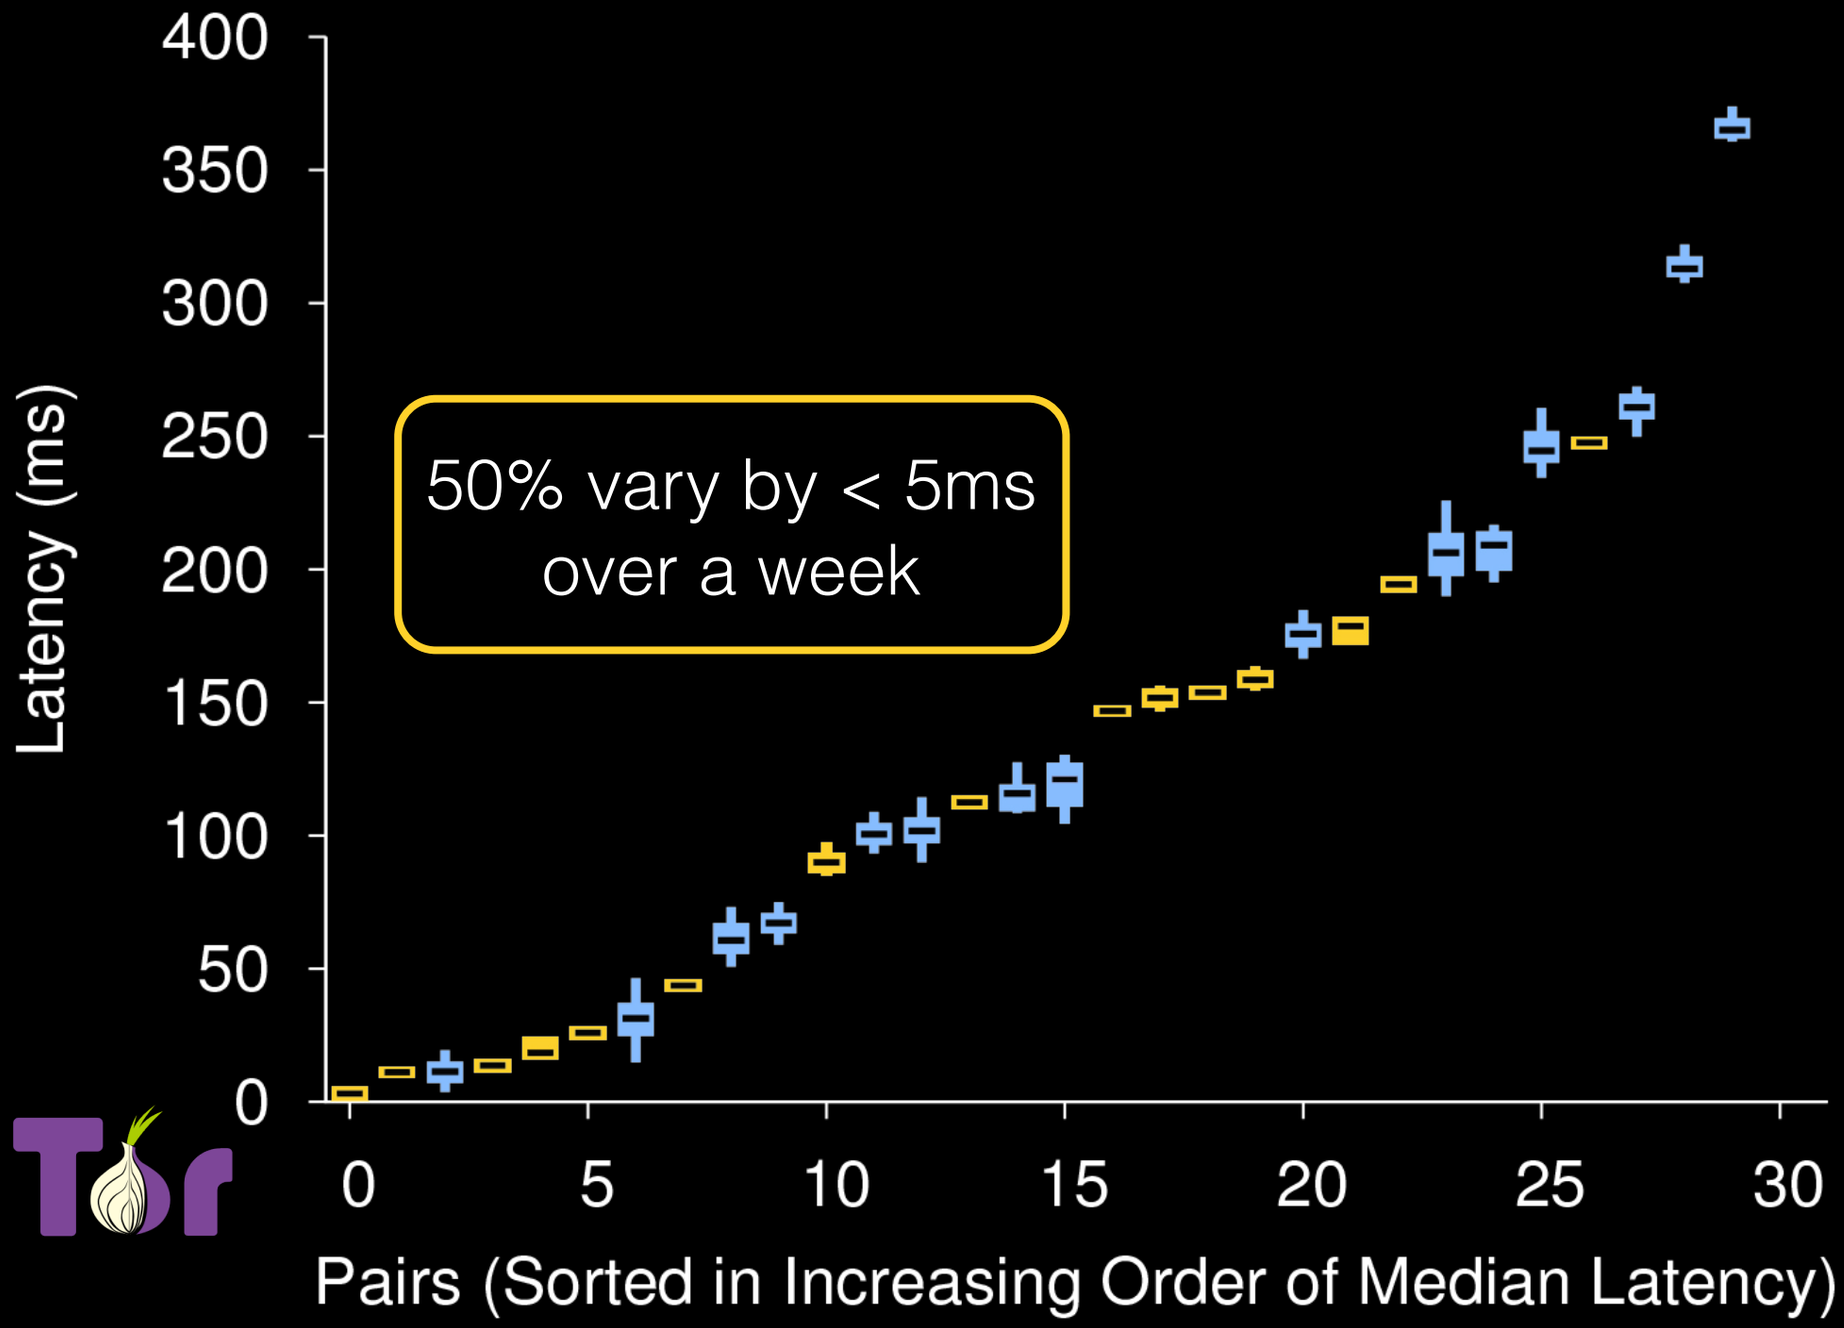

Ting measurements need not be updated often

Evaluation summary

How well does Ting work?

- Typically within 10% of real latency
- Remains accurate with few samples
- Vary by only a few ms over time

Evaluation summary

How well does Ting work?

- Typically within 10% of real latency
- Remains accurate with few samples
- Vary by only a few ms over time

Ting's **stability and accuracy** permit collection of an **all-pairs RTT** dataset

All-pairs RTT dataset

50 Tor relays outside of our control

1,225 pairs in all-pairs RTT dataset

Geographic distribution

Latency distribution

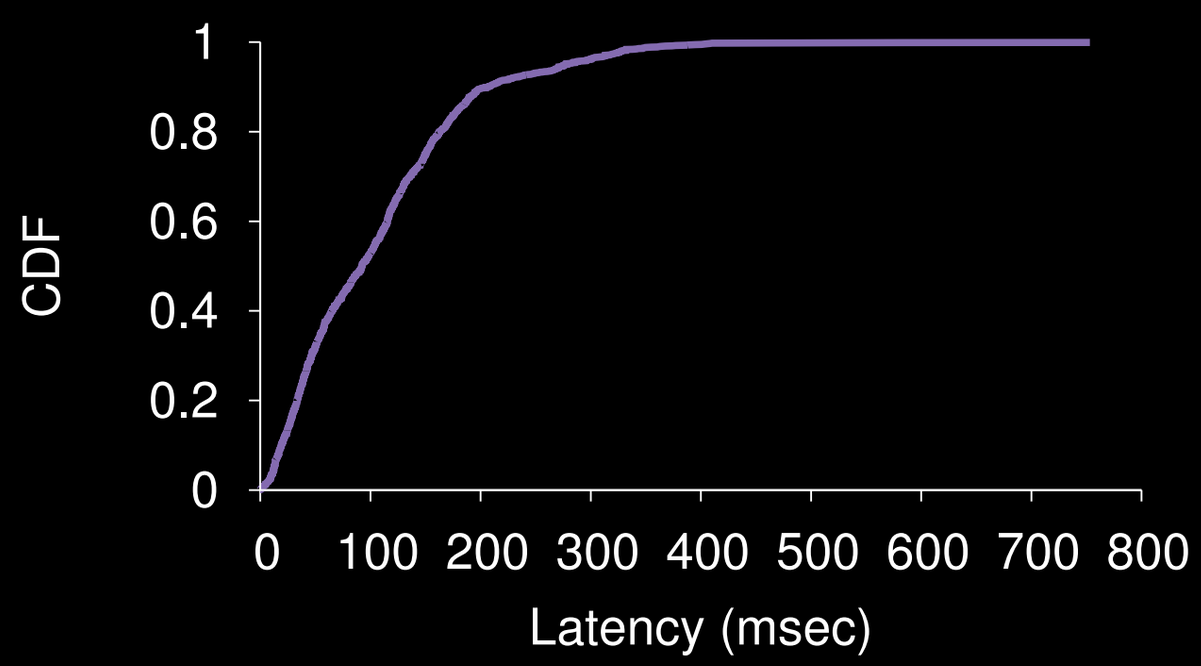

Applications

1. Speeding up deanyonmization of Tor circuits
2. Improving Tor's path selection algorithm
3. Gain insight into non-Tor nodes

Applications

1. Speeding up deanyonmization of Tor circuits

Classic traffic-analysis attack

[Murdoch and Danezis, 2005]

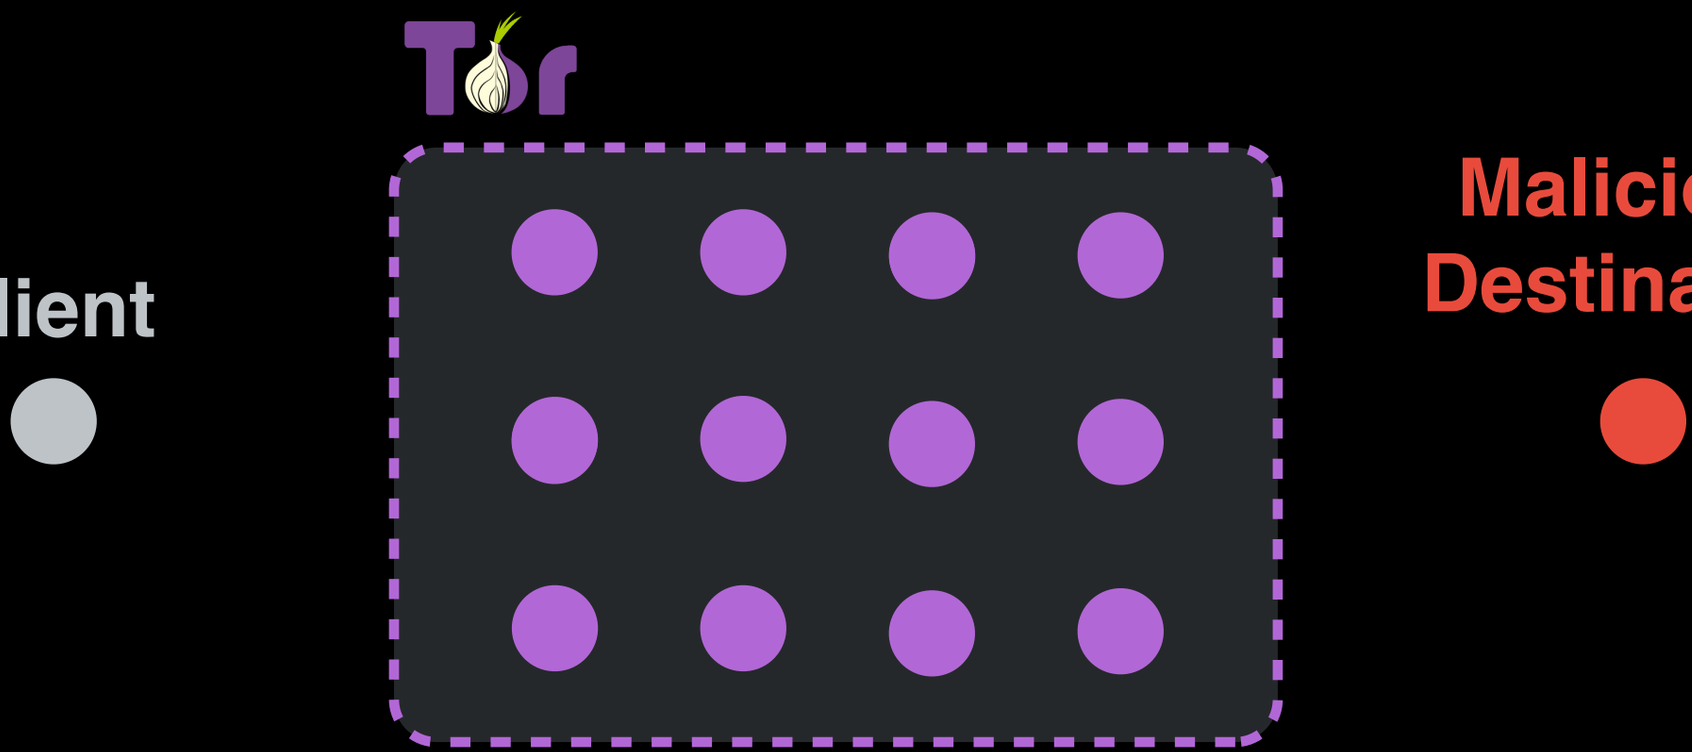

Classic traffic-analysis attack

[Murdoch and Danezis, 2005]

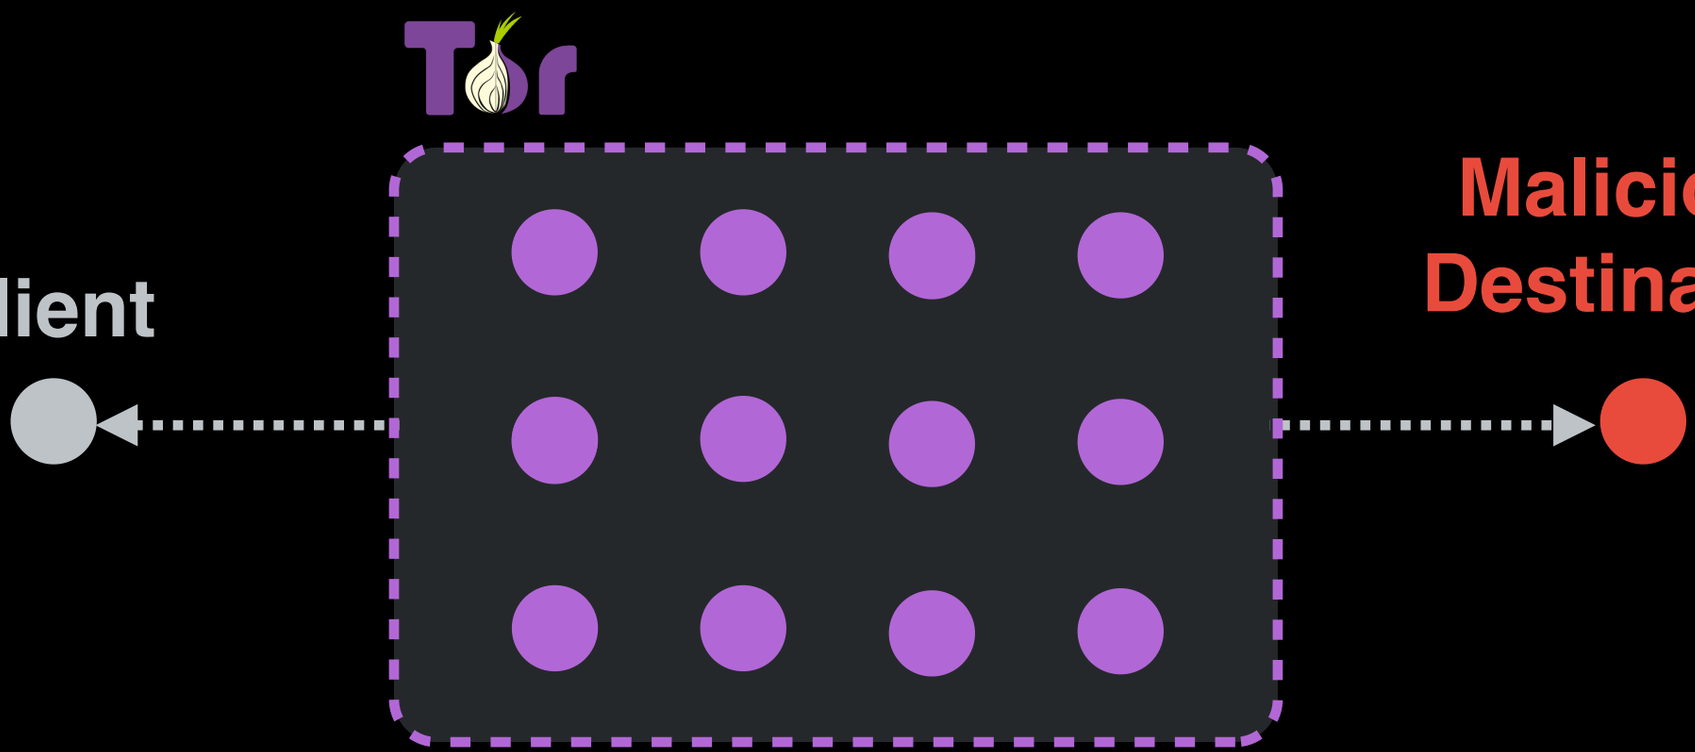

Classic traffic-analysis attack

[Murdoch and Danezis, 2005]

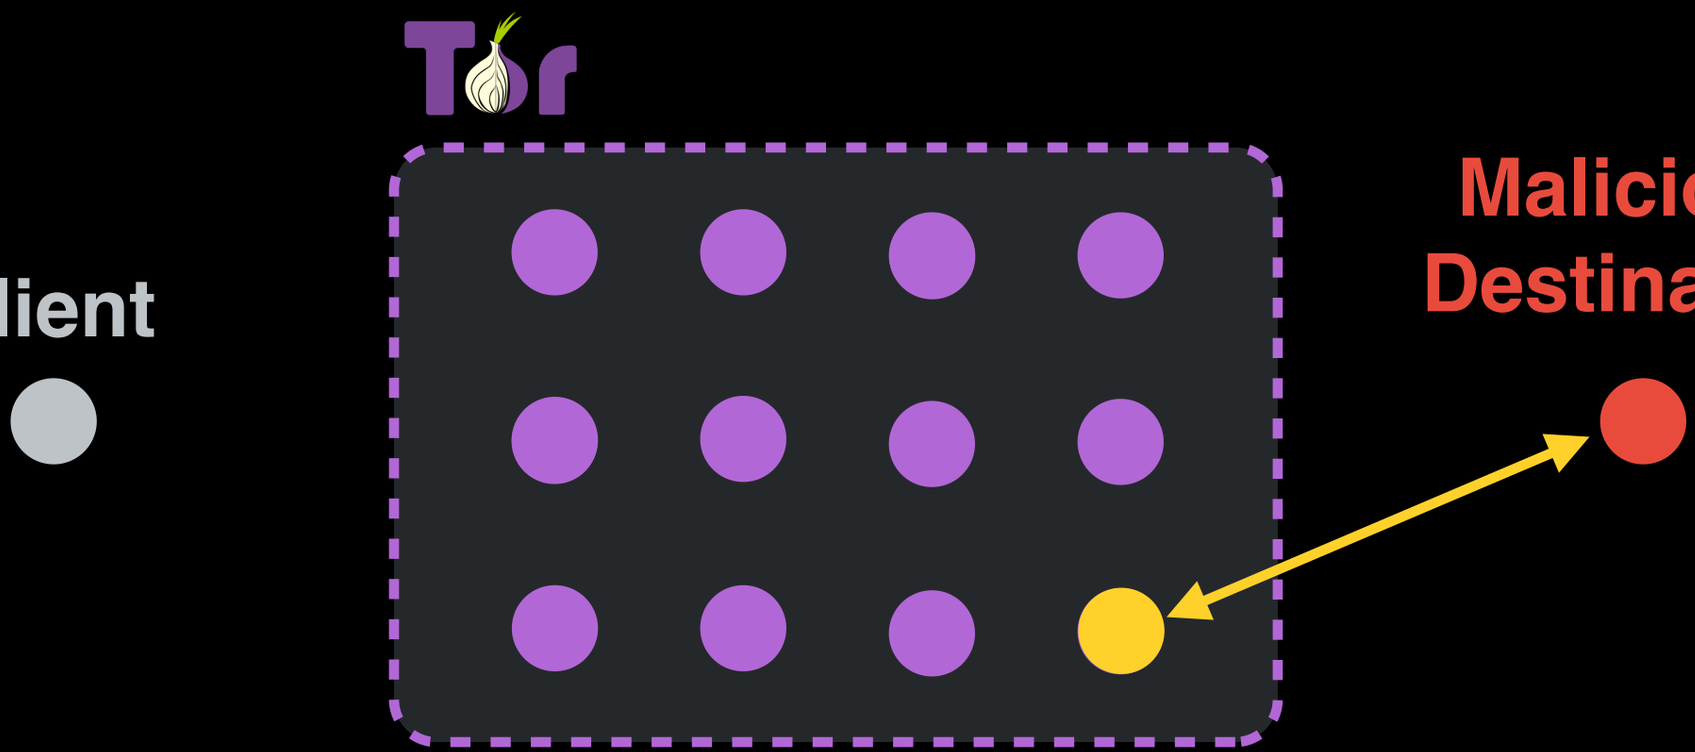

Classic traffic-analysis attack

[Murdoch and Danezis, 2005]

Attacker's Goal: find all nodes in the **circuit**

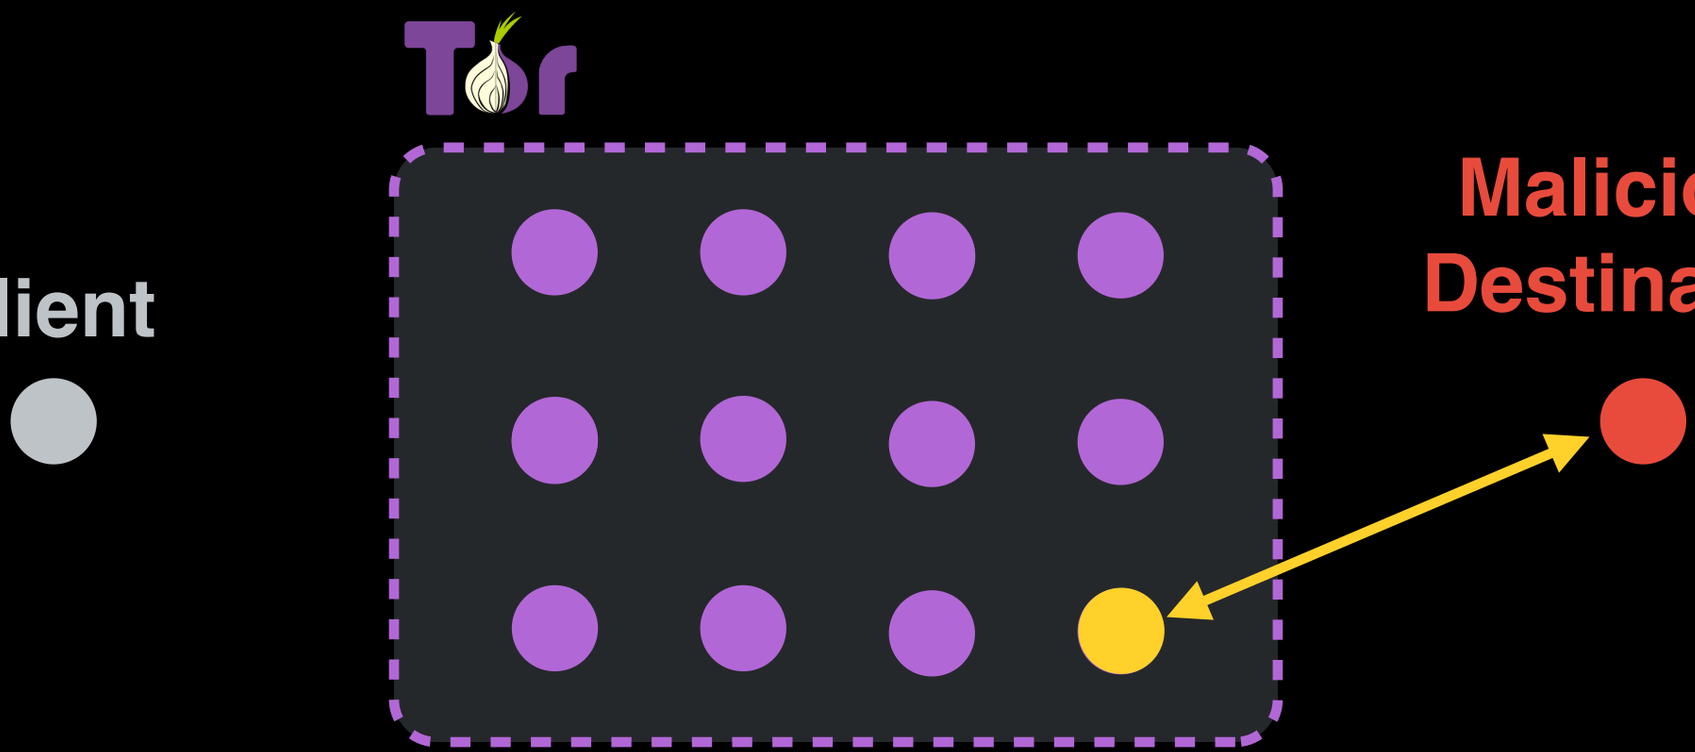

Classic traffic-analysis attack

[Murdoch and Danezis, 2005]

Attacker's Goal: find all nodes in the **circuit**

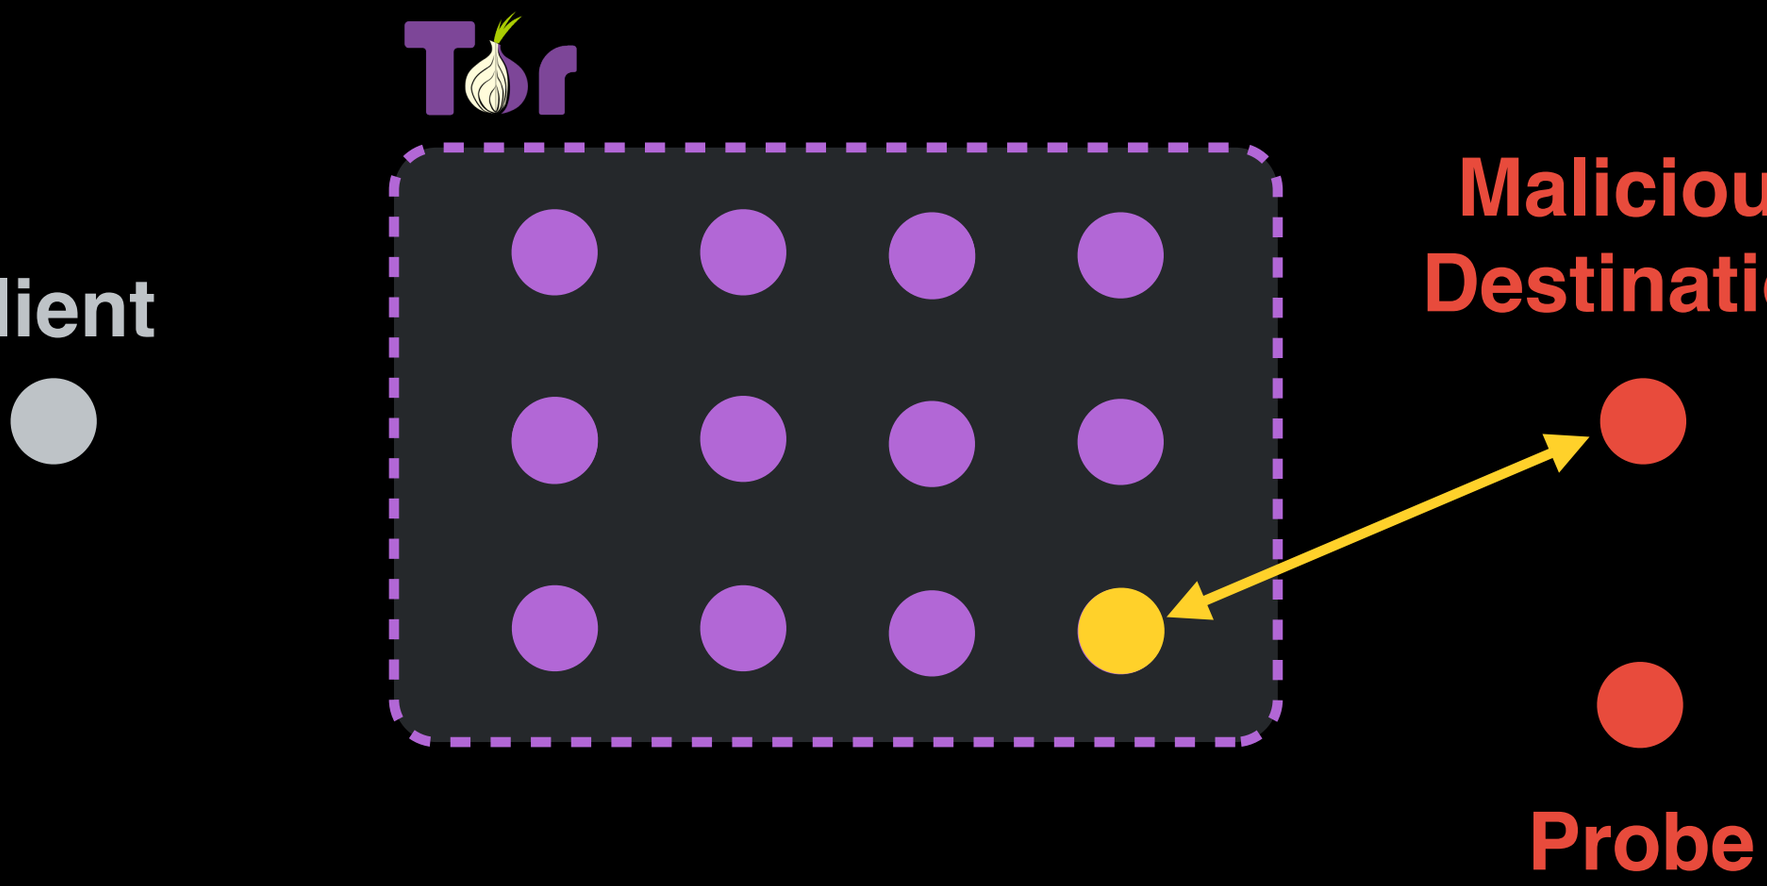

Classic traffic-analysis attack

[Murdoch and Danezis, 2005]

Attacker's Goal: find all nodes in the **circuit**

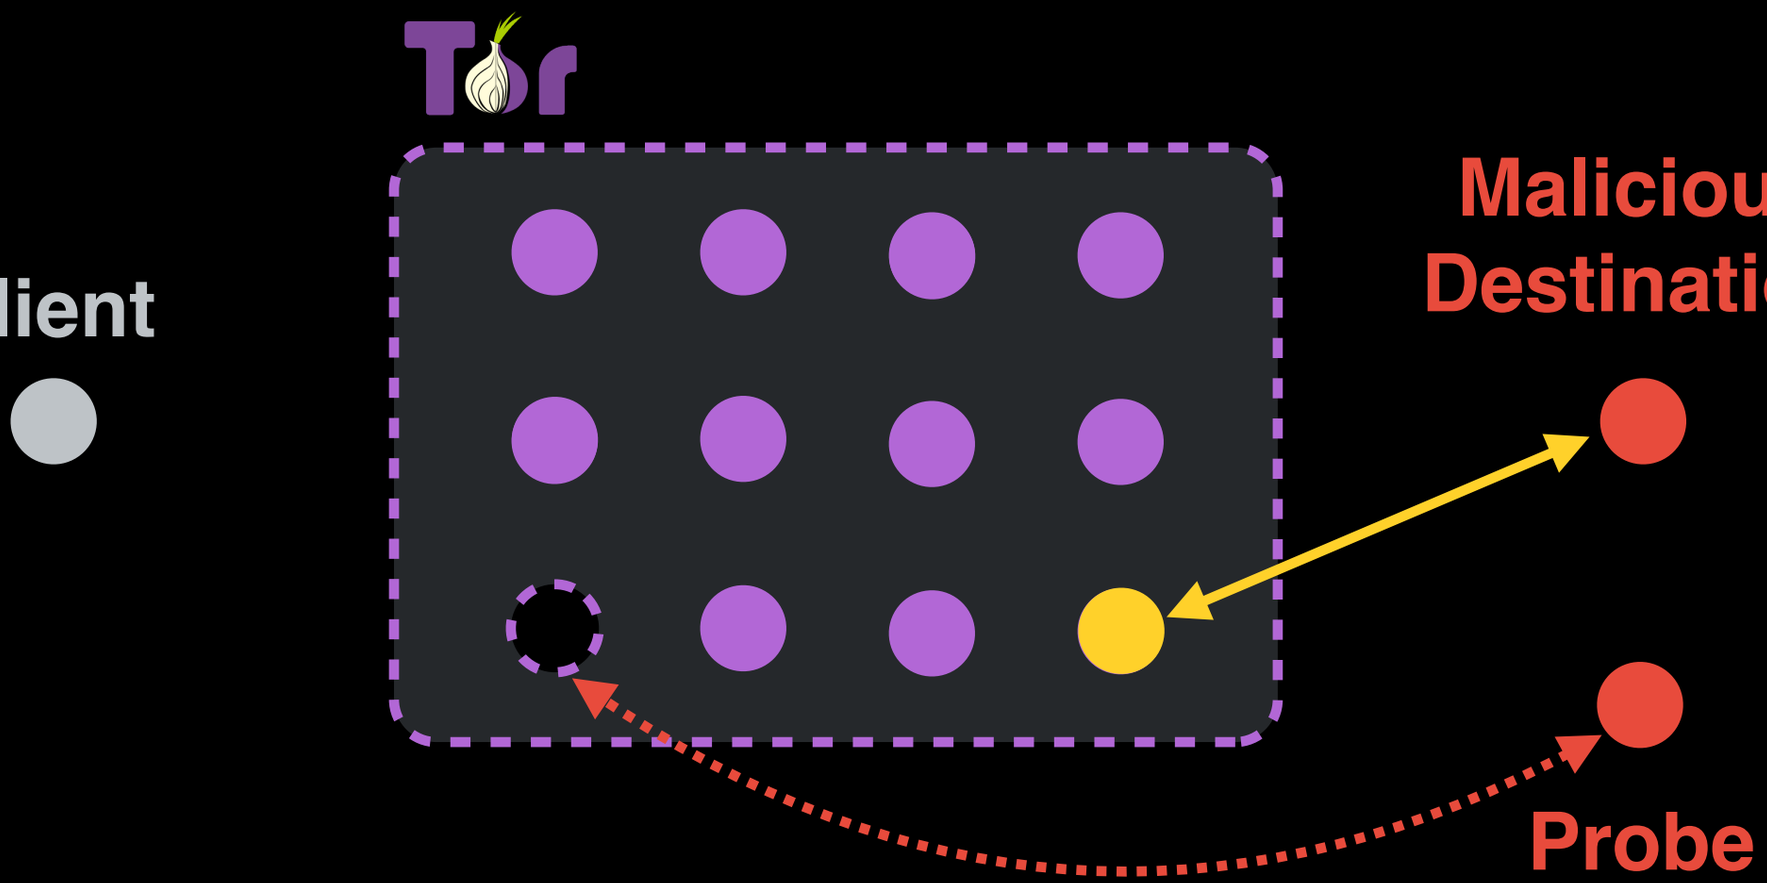

Classic traffic-analysis attack

[Murdoch and Danezis, 2005]

Attacker's Goal: find all nodes in the **circuit**

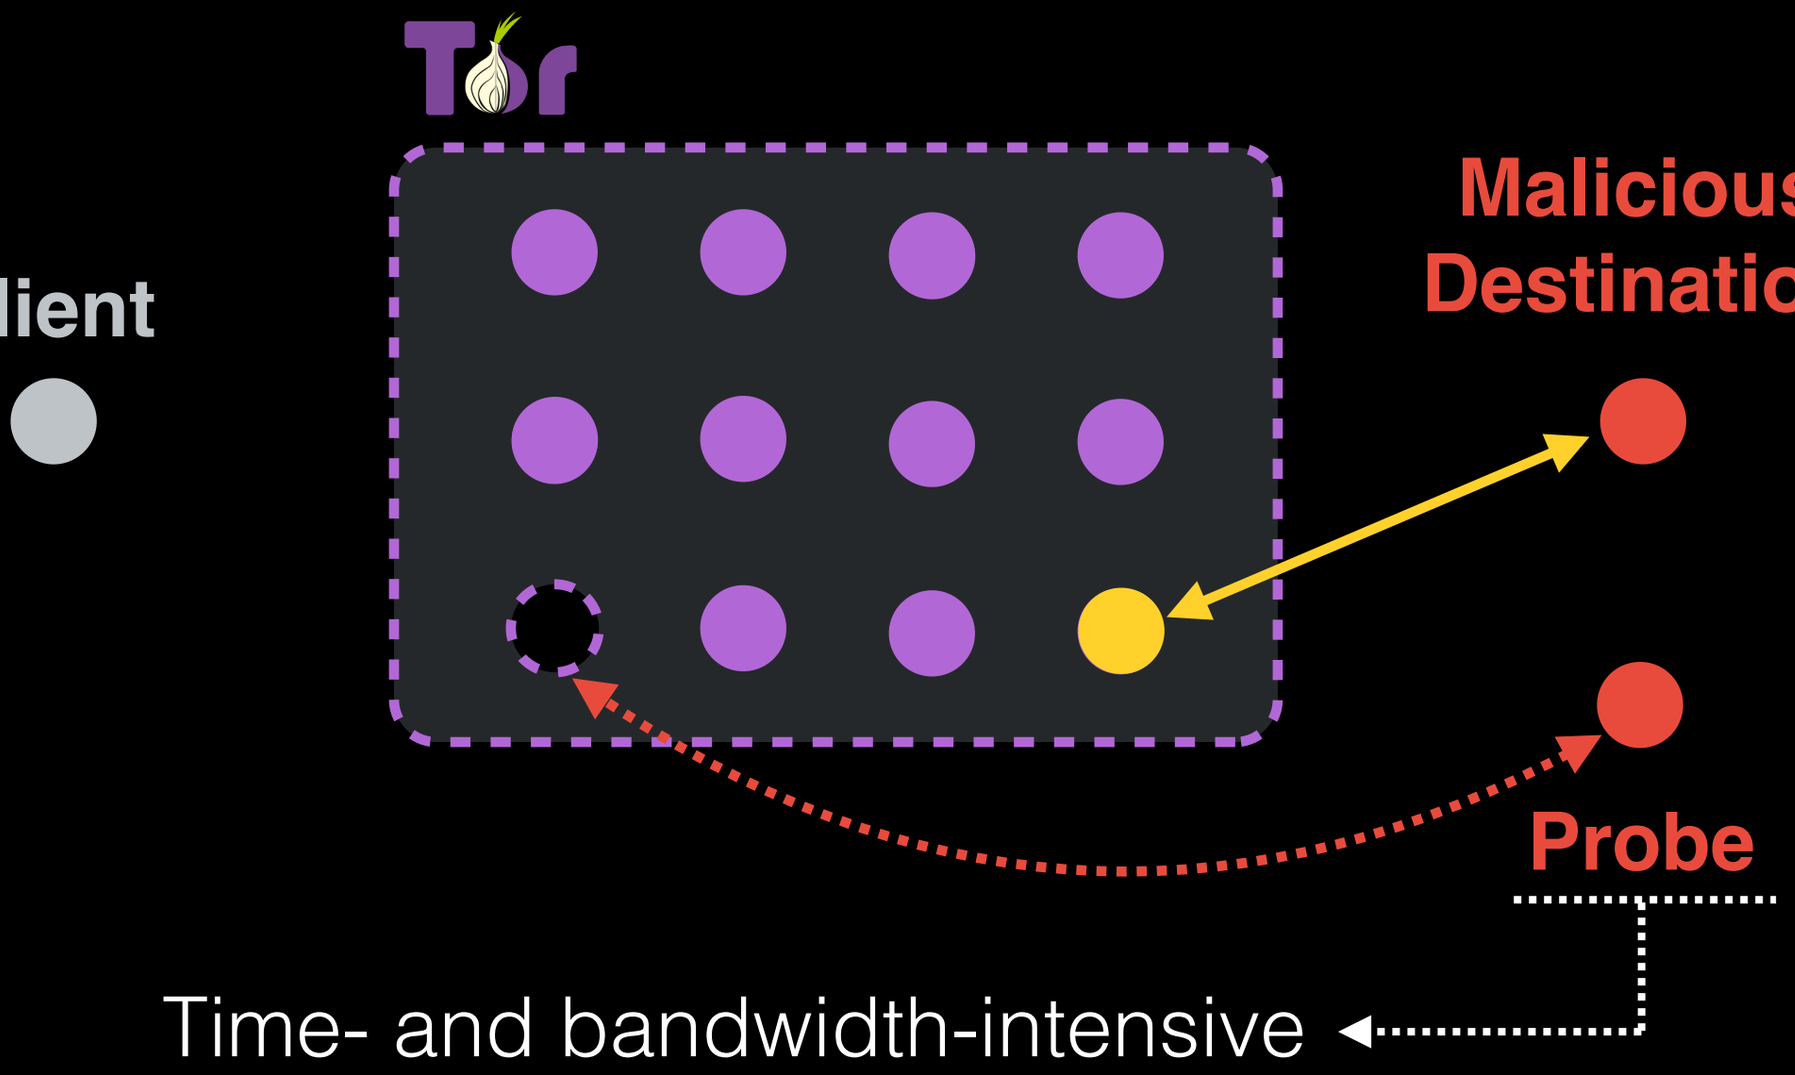

Time- and bandwidth-intensive

Classic traffic-analysis attack

[Murdoch and Danezis, 2005]

Attacker's Goal: find all nodes in the **circuit**

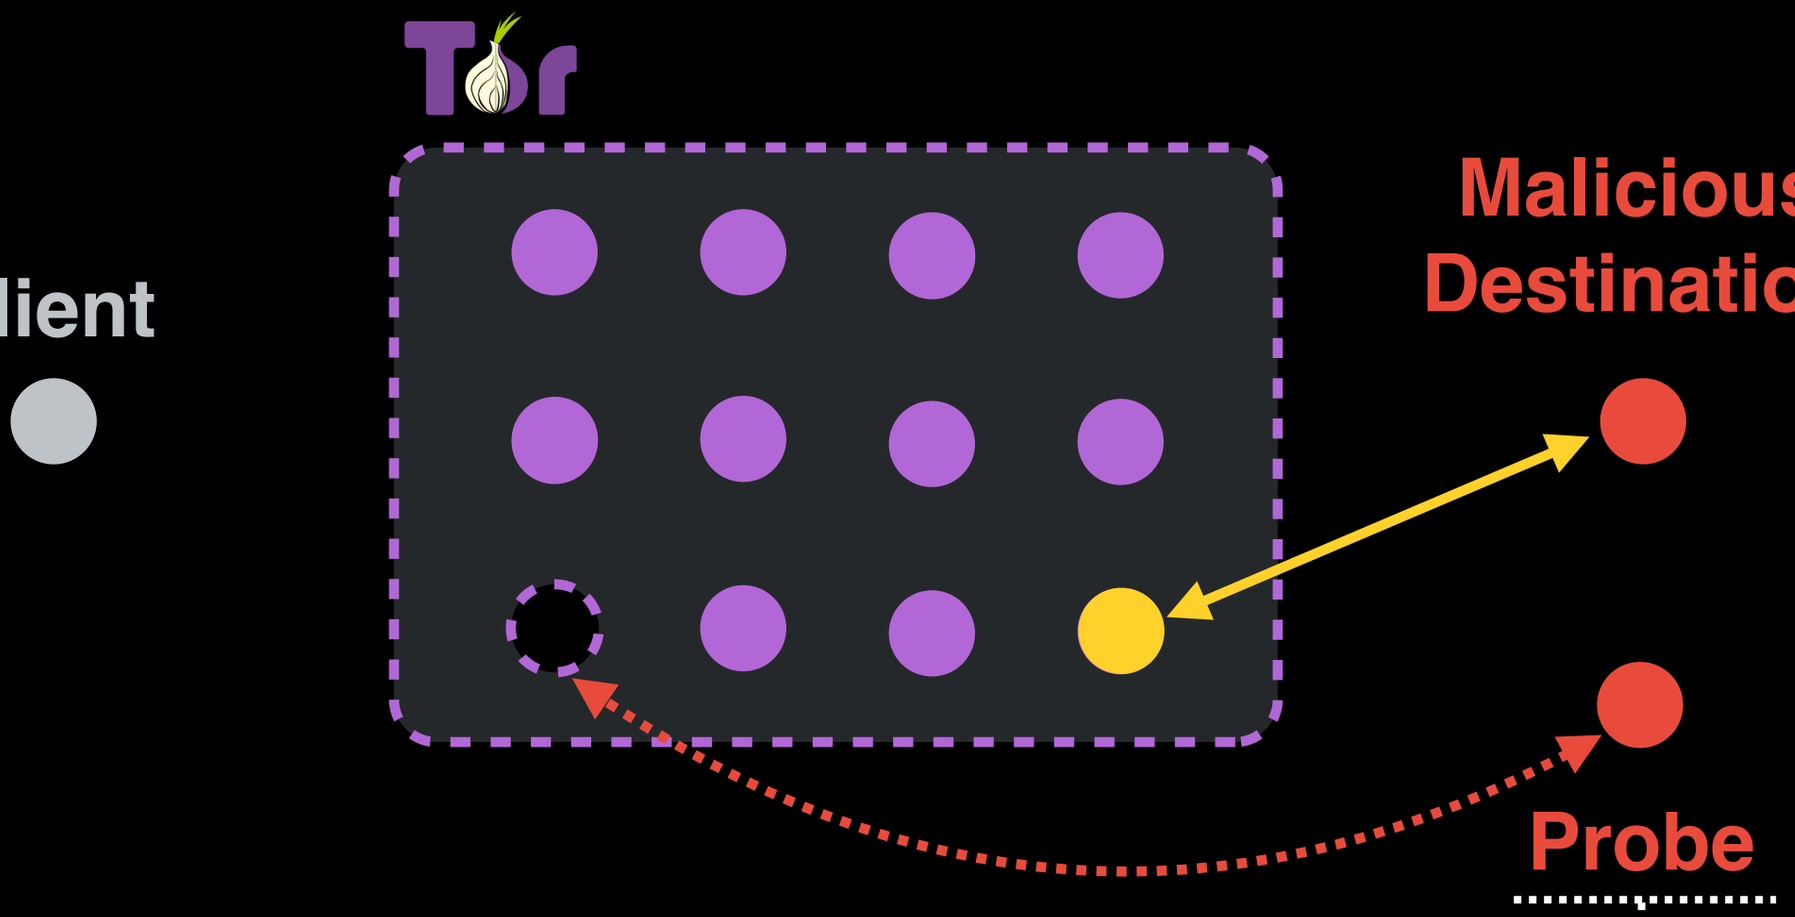

Time- and bandwidth-intensive

Done in an unguided fashion

Classic traffic-analysis attack

[Murdoch and Danezis, 2005]

Attacker's Goal: find all nodes in the **circuit**

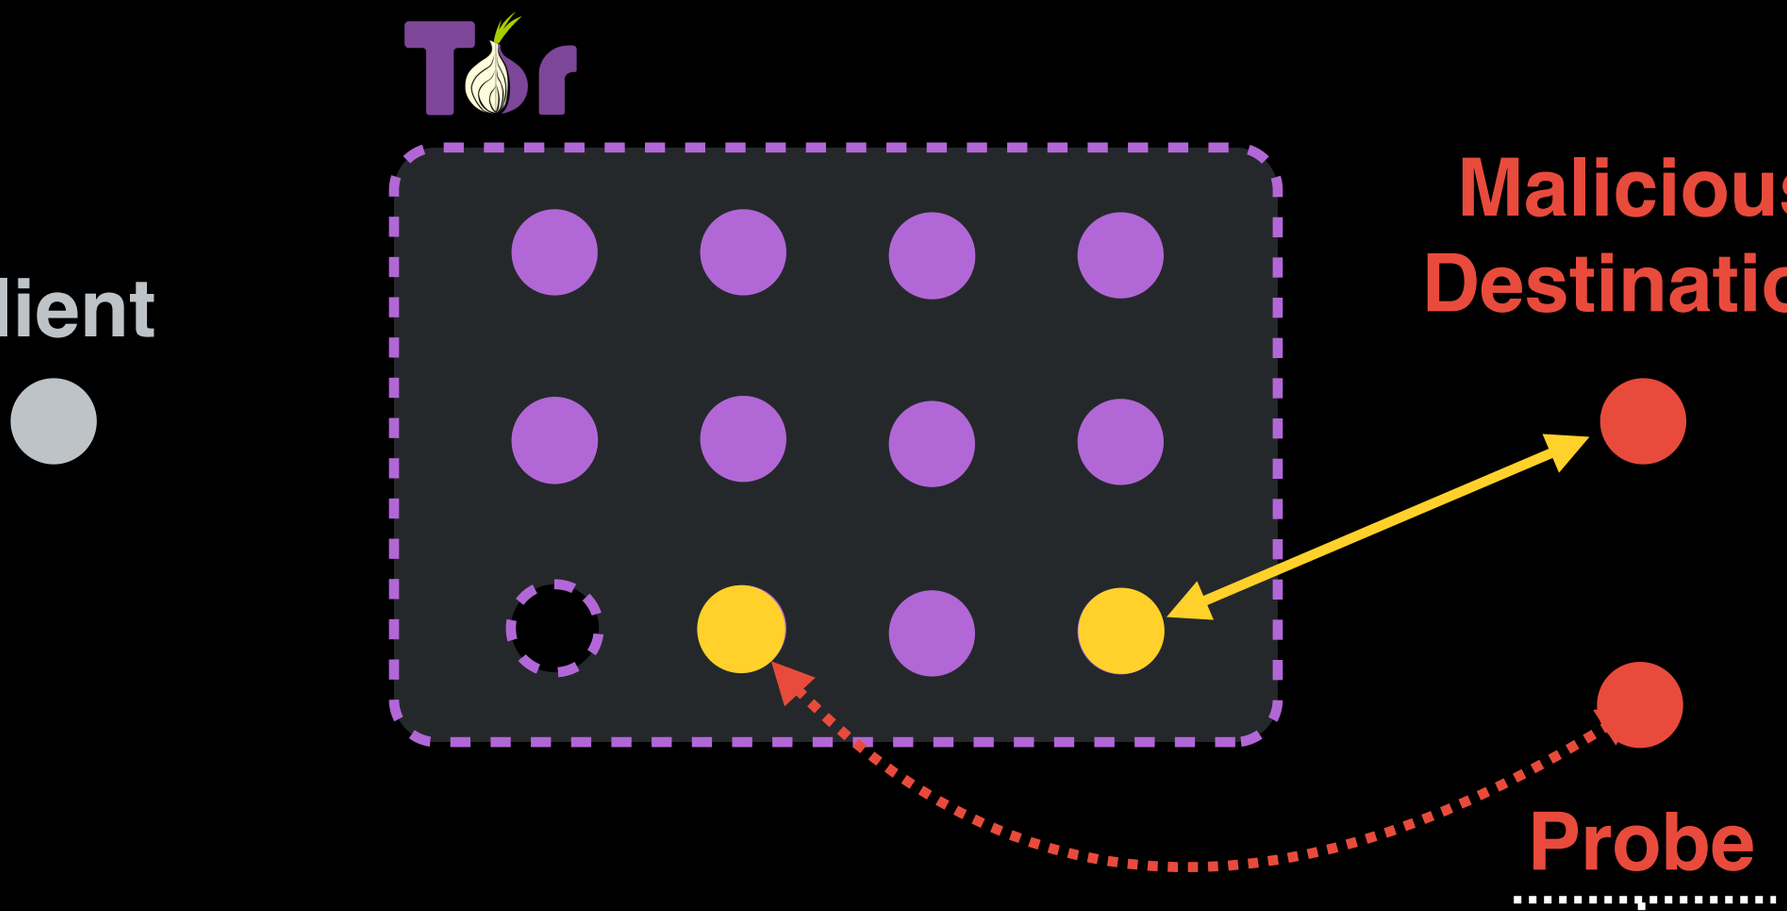

Time- and bandwidth-intensive

Done in an unguided fashion

Classic traffic-analysis attack

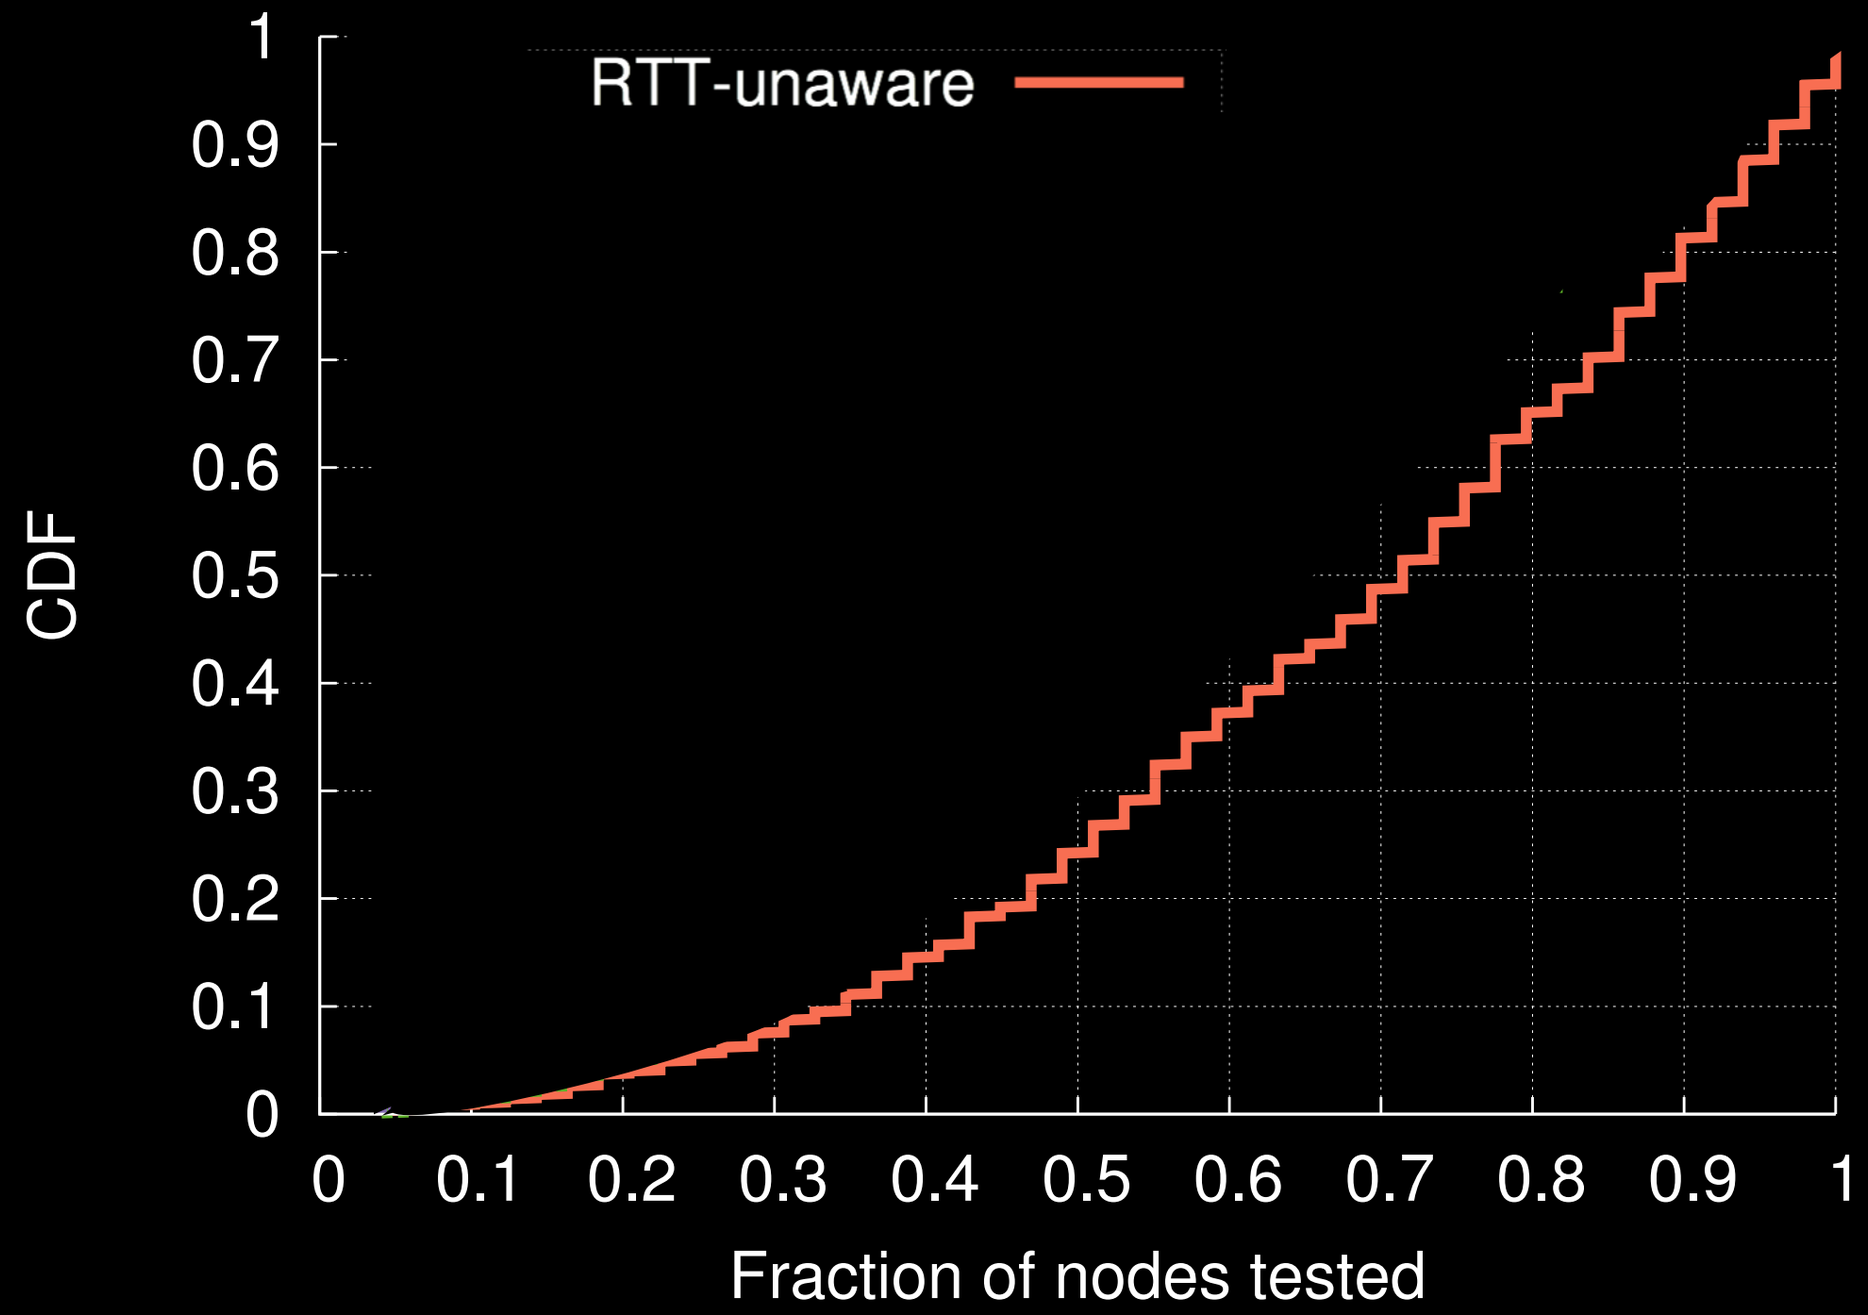

Classic traffic-analysis attack

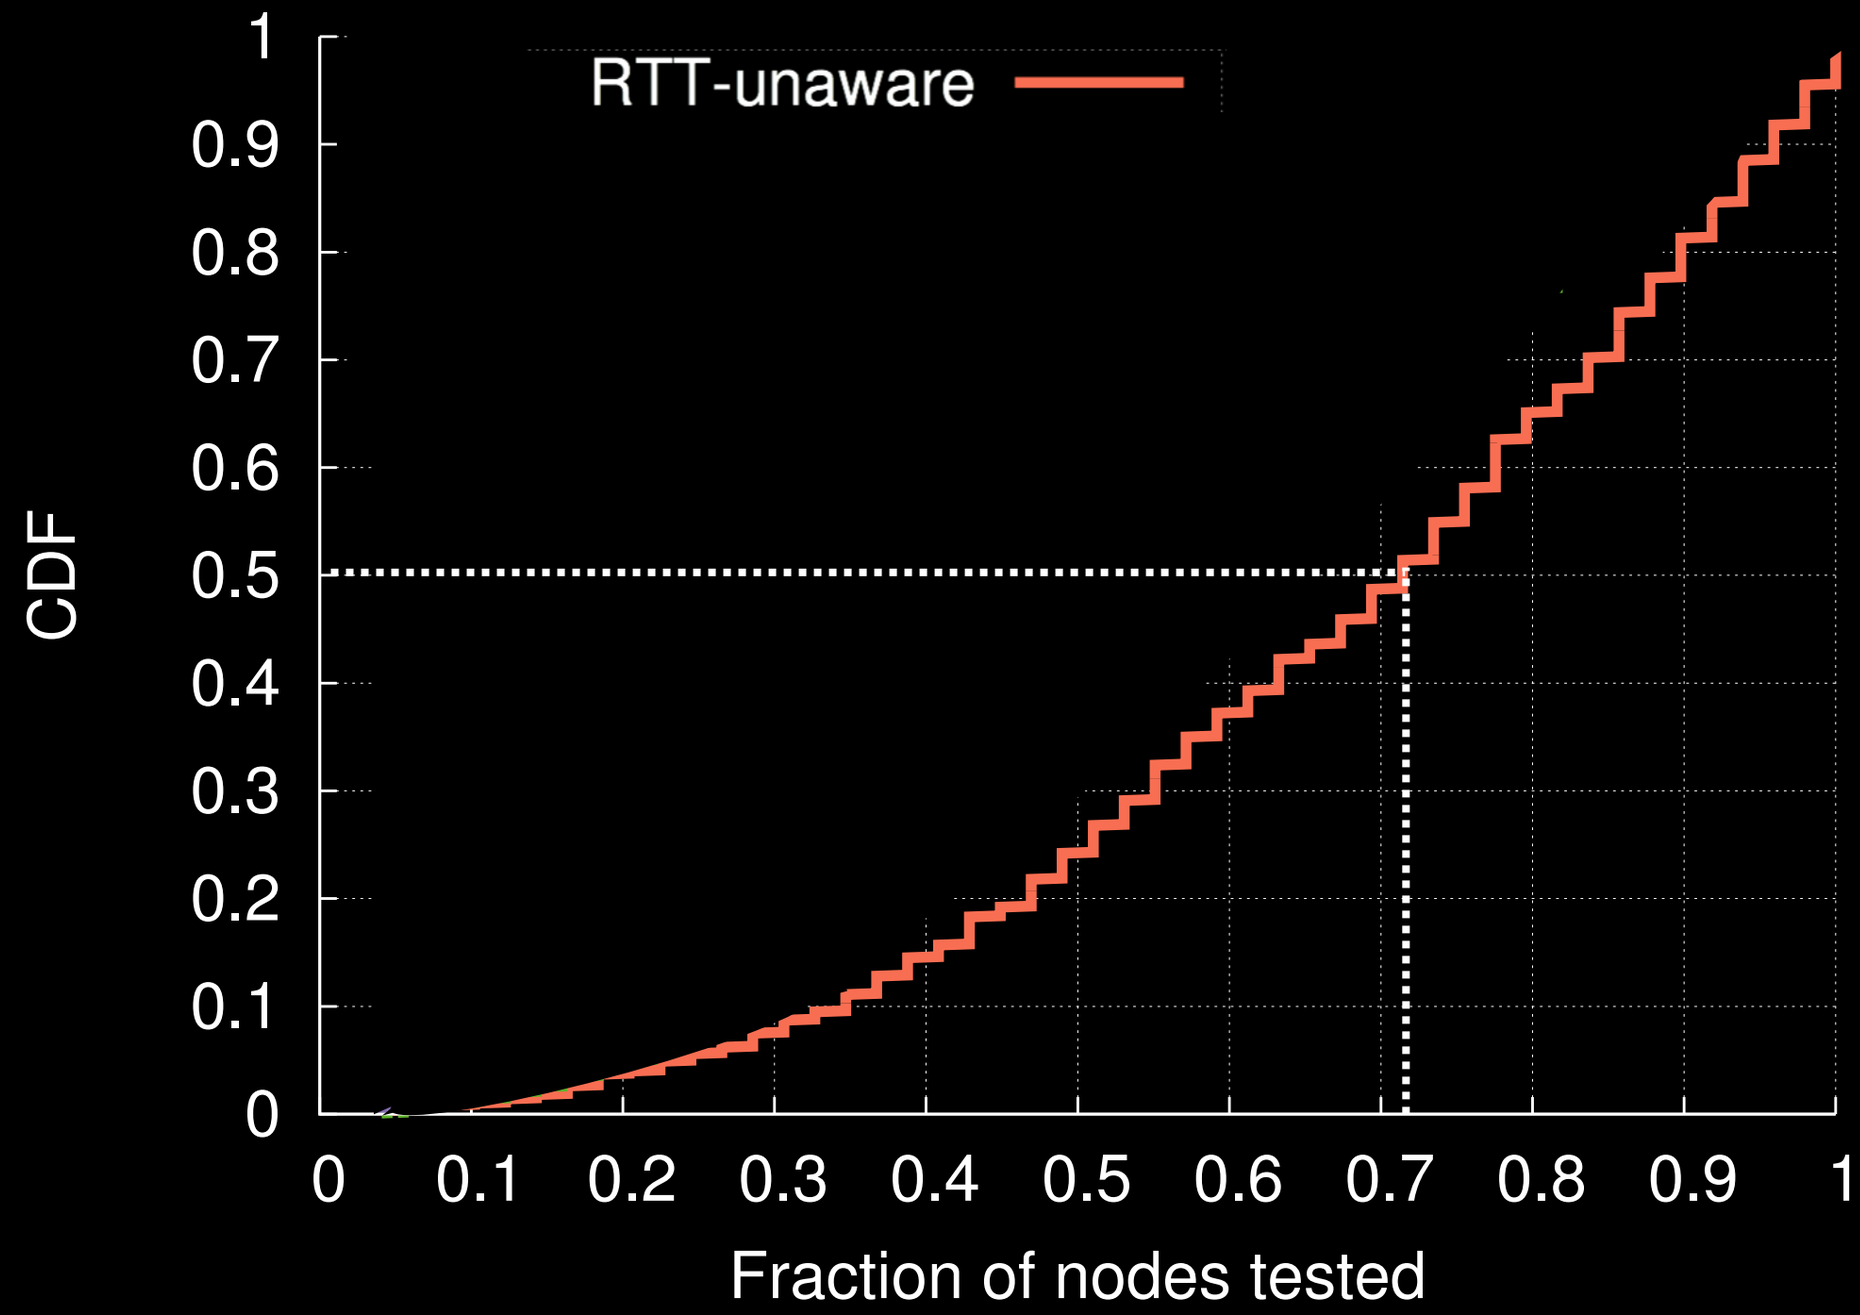

Faster deanonymization with Ting

Apply what the attacker knows about latencies

Faster deanonymization with Ting

Apply what the attacker knows about latencies

Faster deanonymization with Ting

Apply what the attacker knows about latencies

e2e RTT

Directly measured

Faster deanonymization with Ting

Apply what the attacker knows about latencies

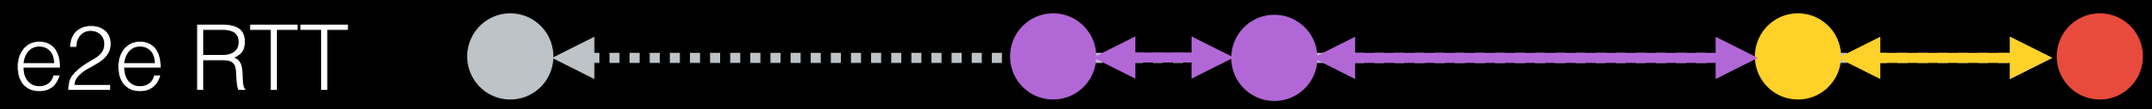

Pre-measured with Ting

Directly measured

Faster deanonymization with Ting

Apply what the attacker knows about latencies

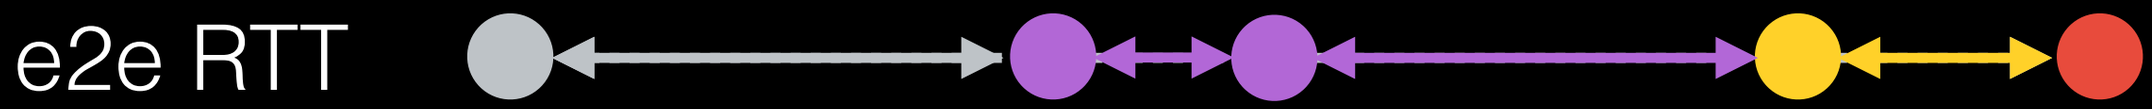

Client's RTT to the entry node is unknown

Pre-measured with Ting

Directly measured

Faster deanonymization with Ting

Apply what the attacker knows about latencies

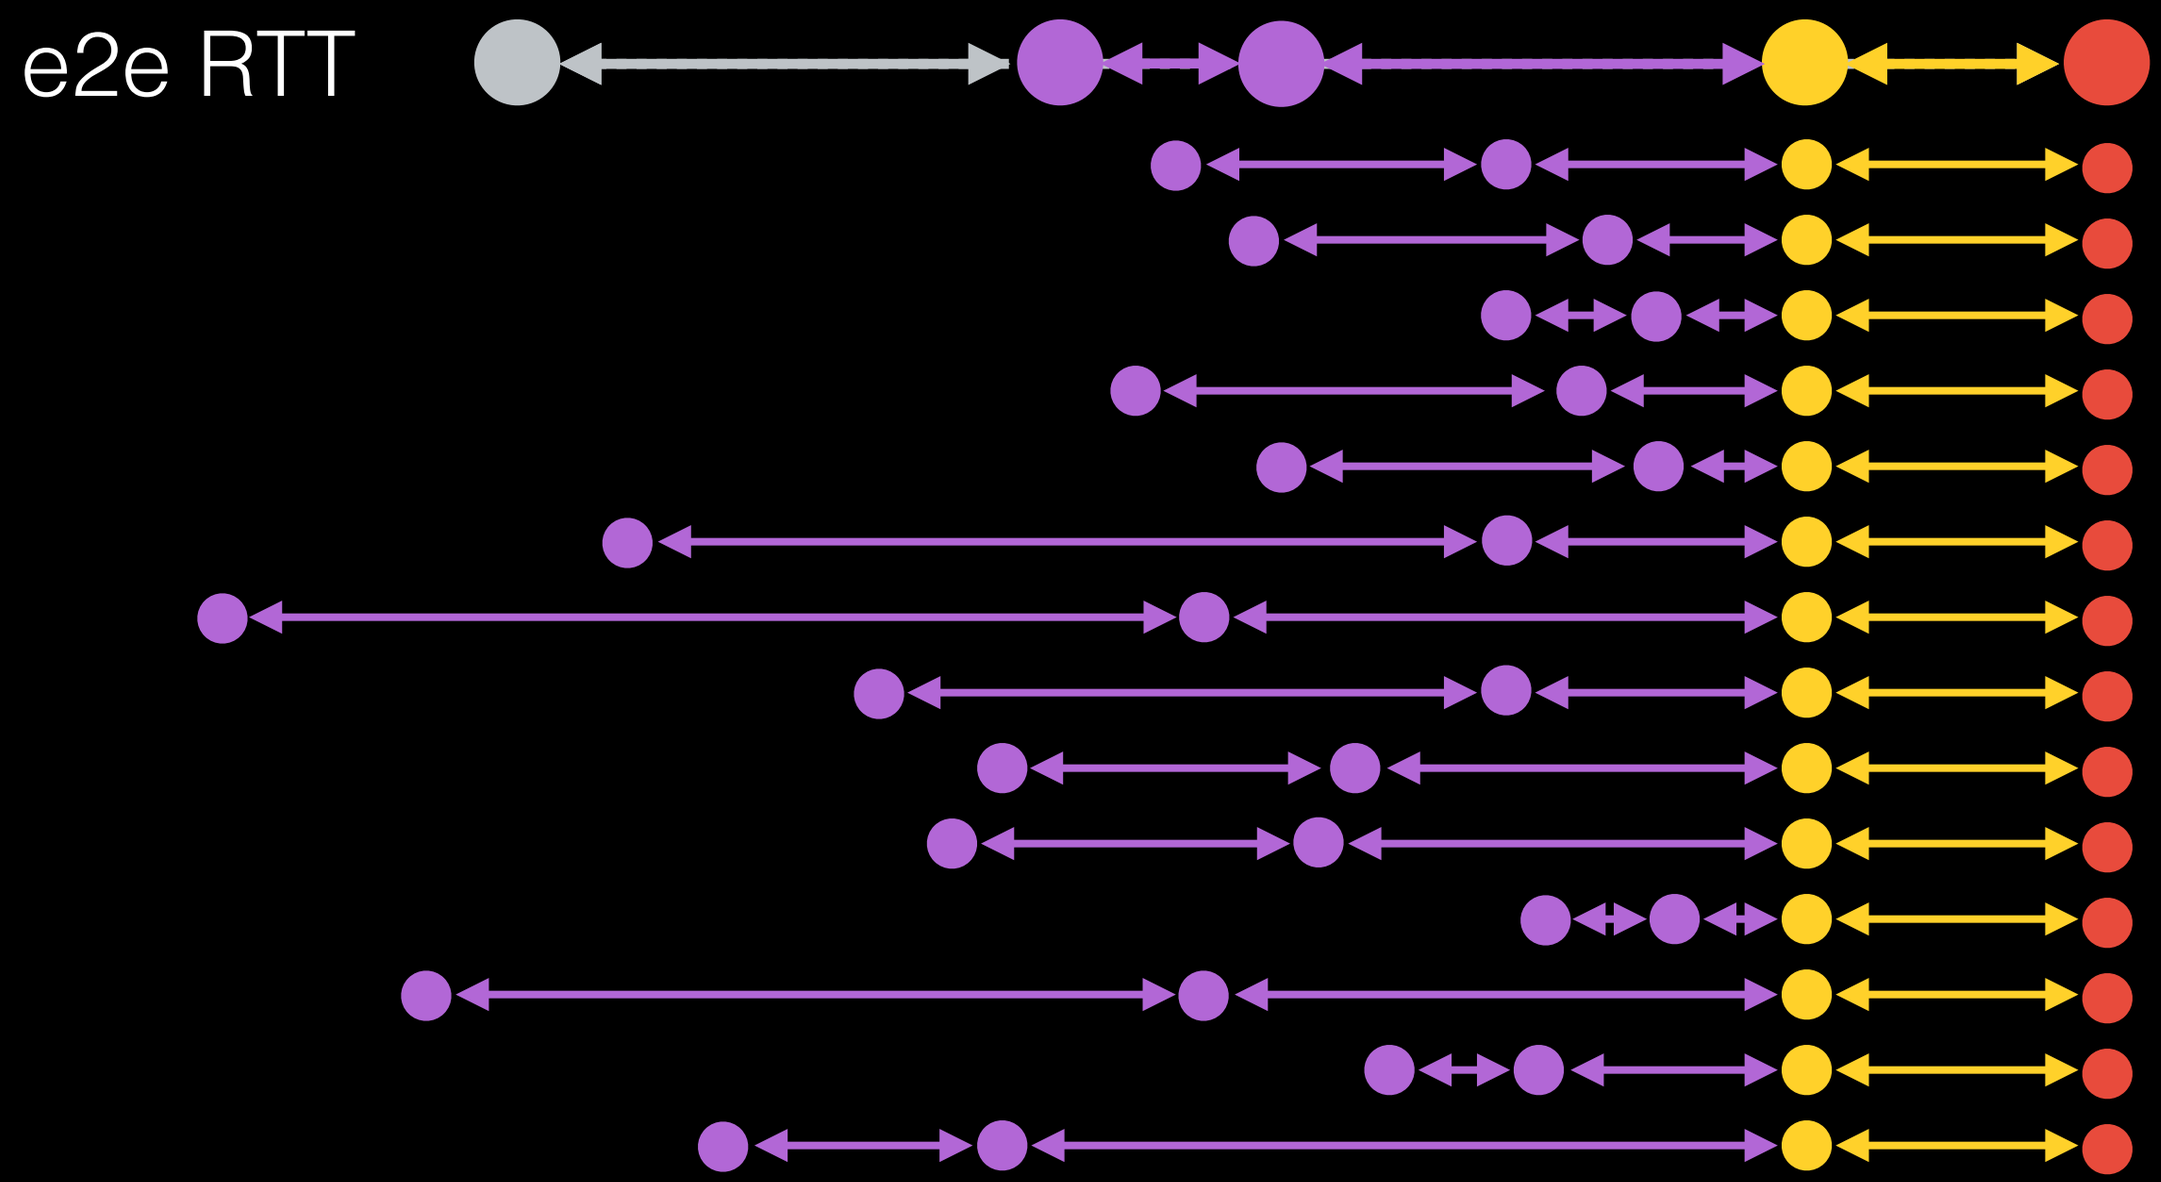

Faster deanonymization with Ting

Reason about what the client → entry RTT would have to be

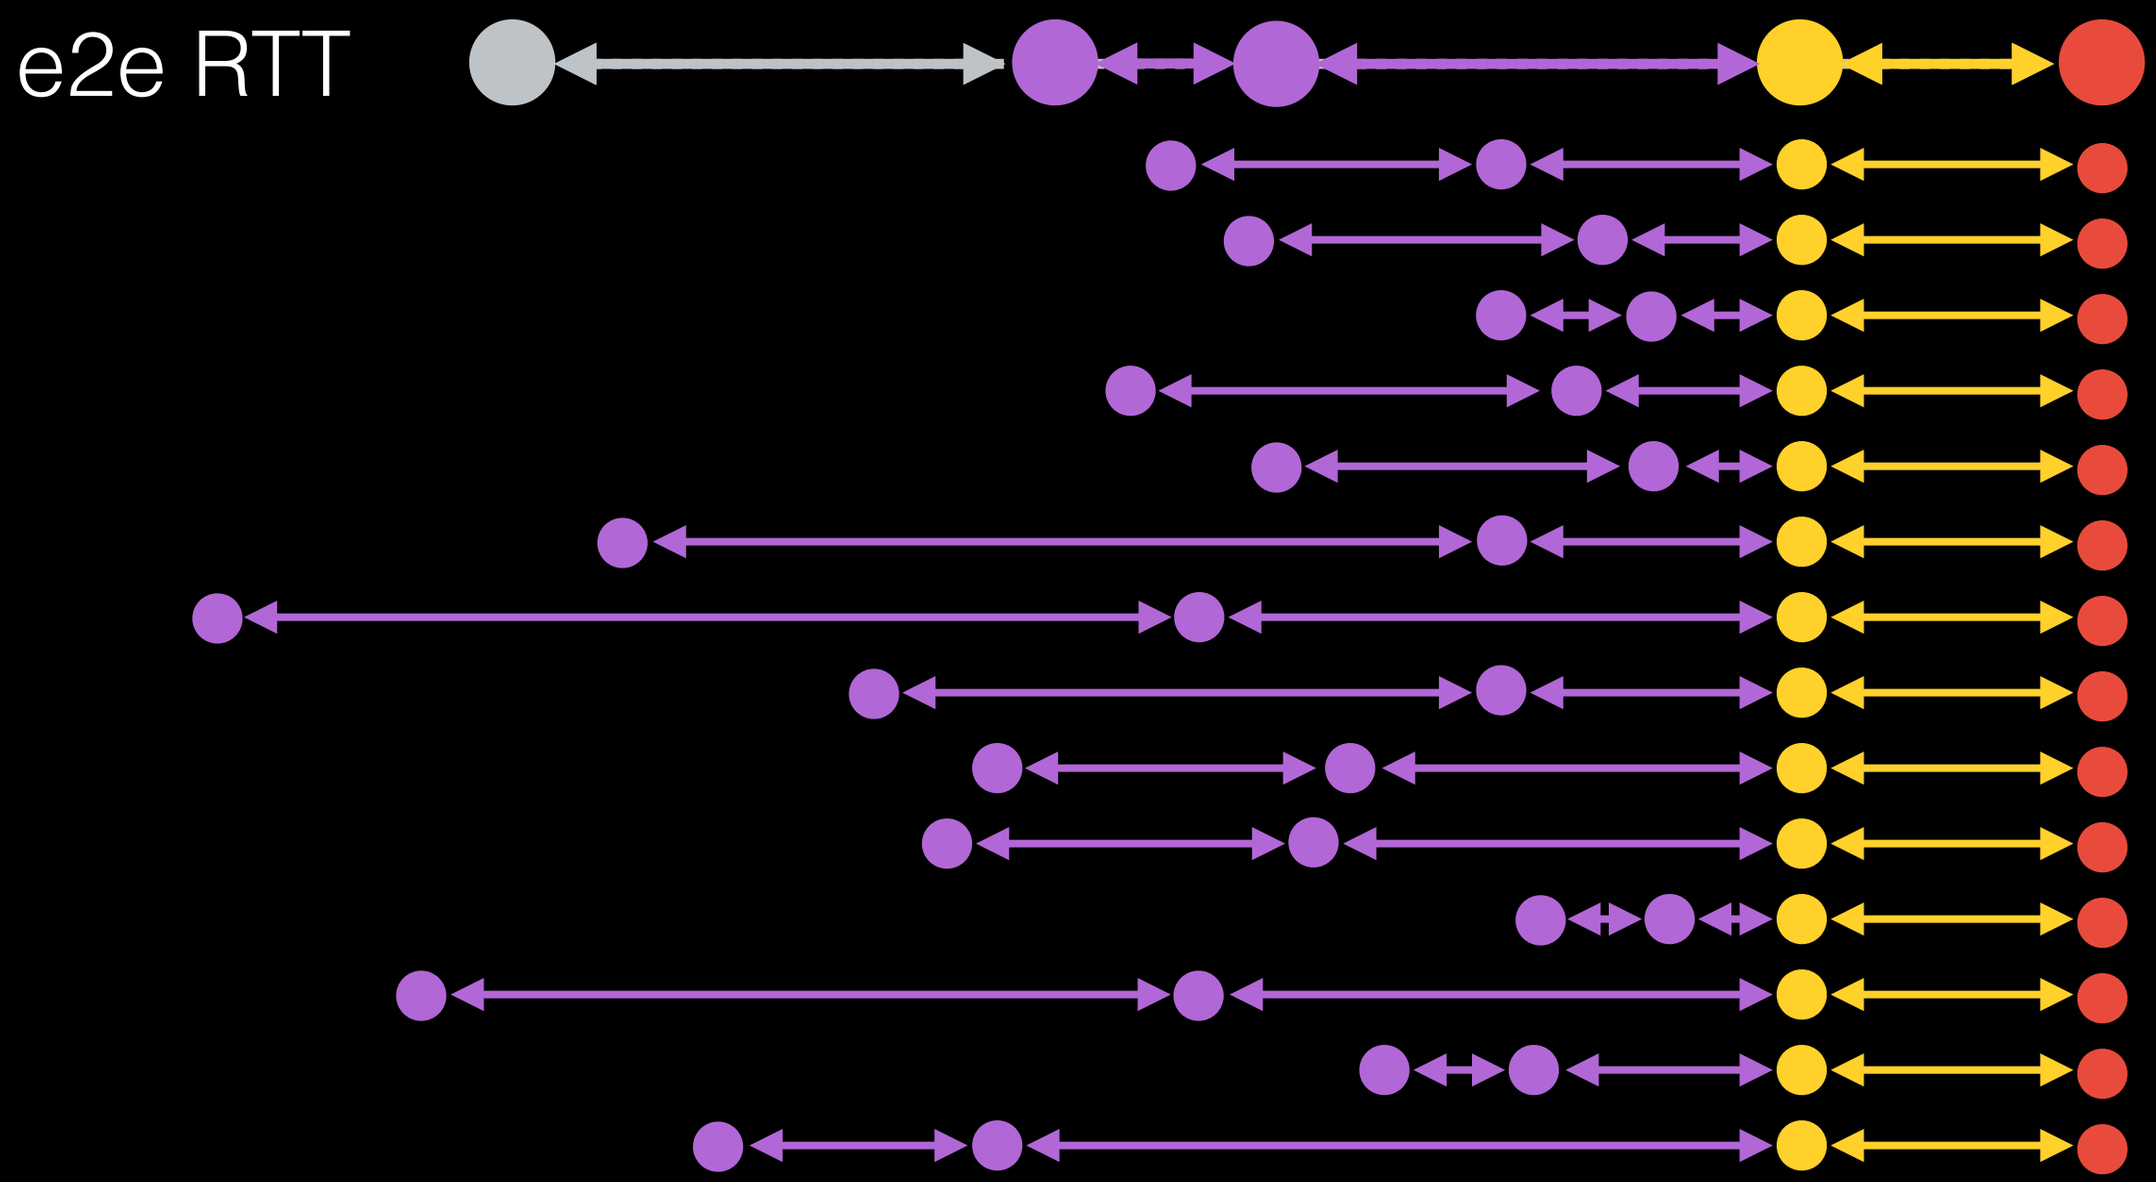

Faster deanonymization with Ting

Reason about what the client → entry RTT would have to be

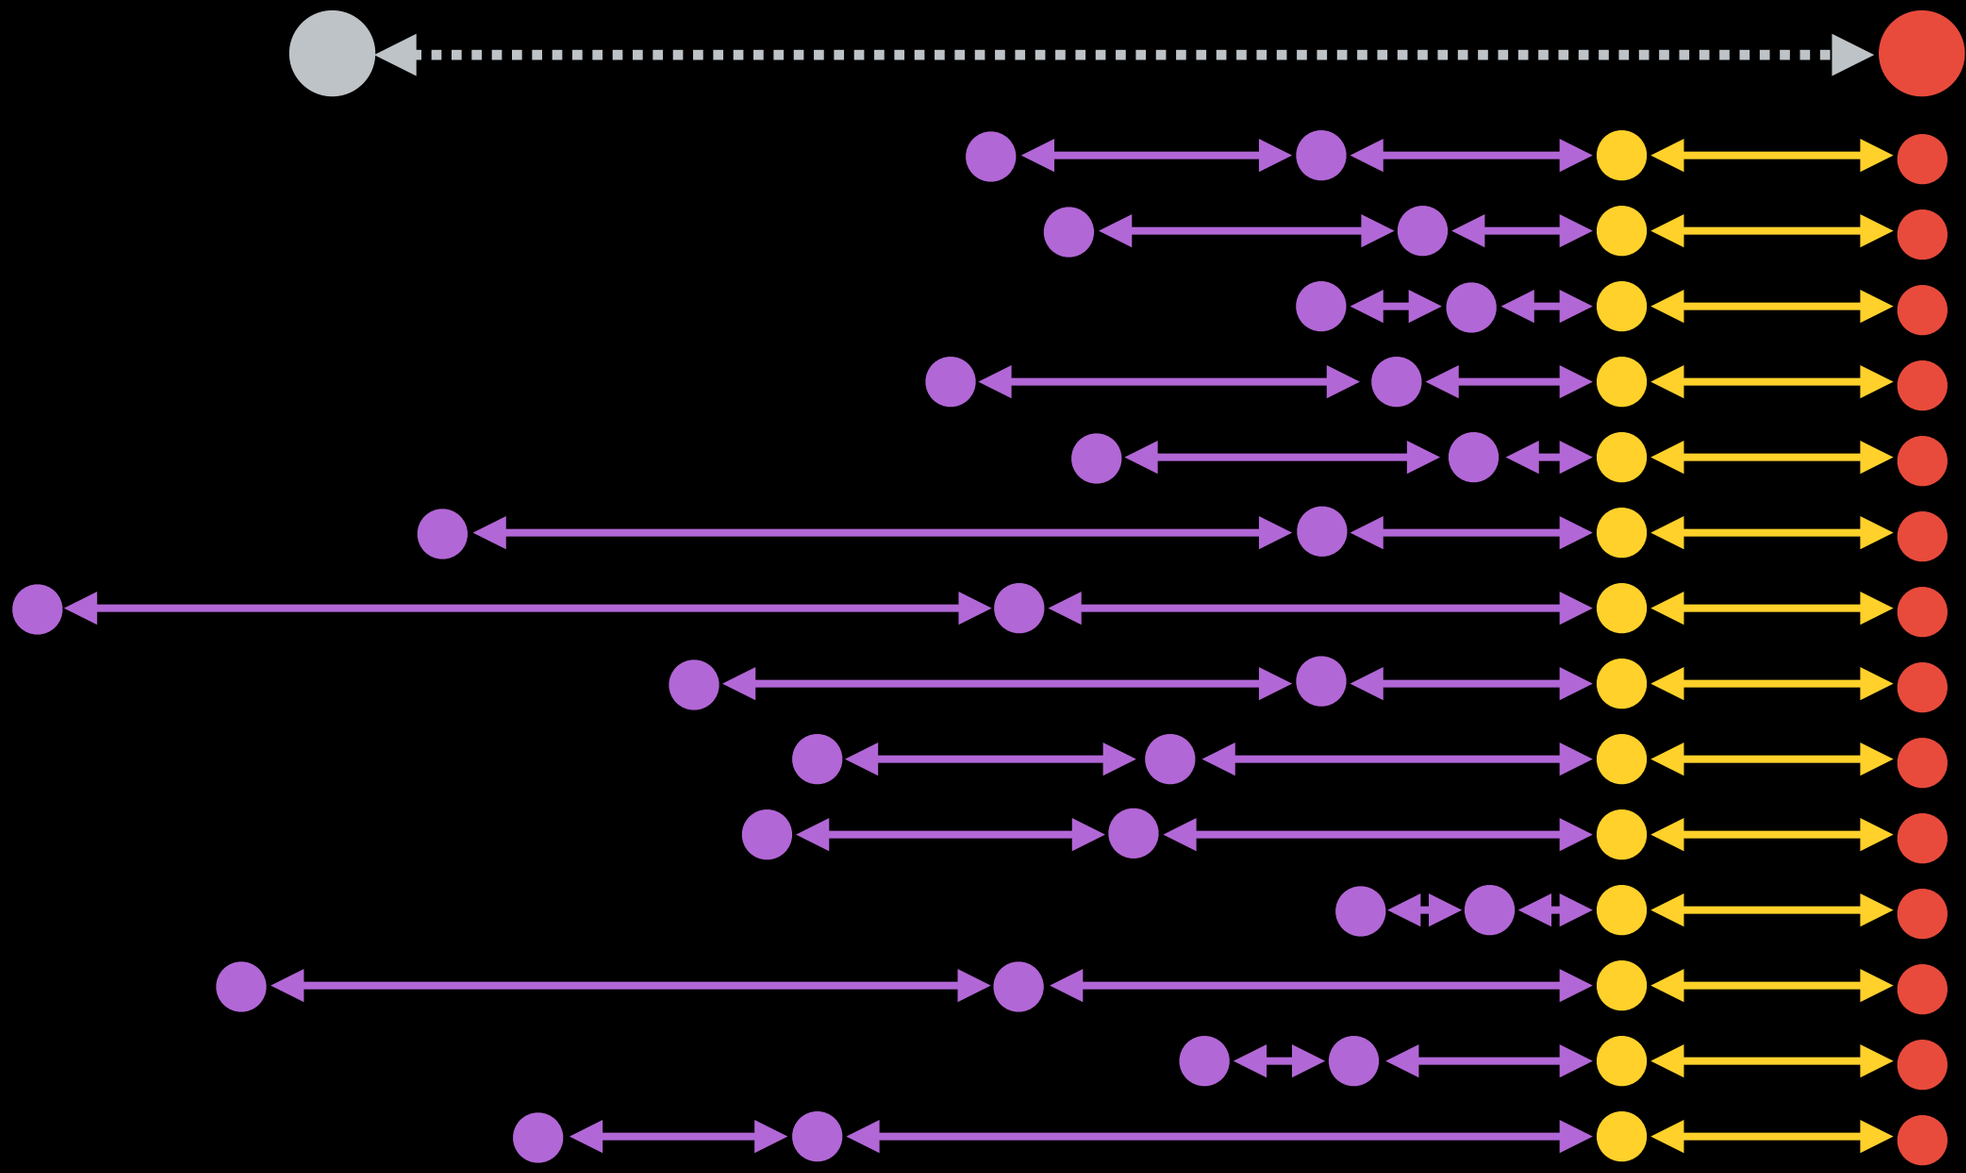

Client's RTT to the entry node is unknown

Pre-measured with Ting

Directly measured

Faster deanonymization with Ting

Reason about what the client → entry RTT would have to be

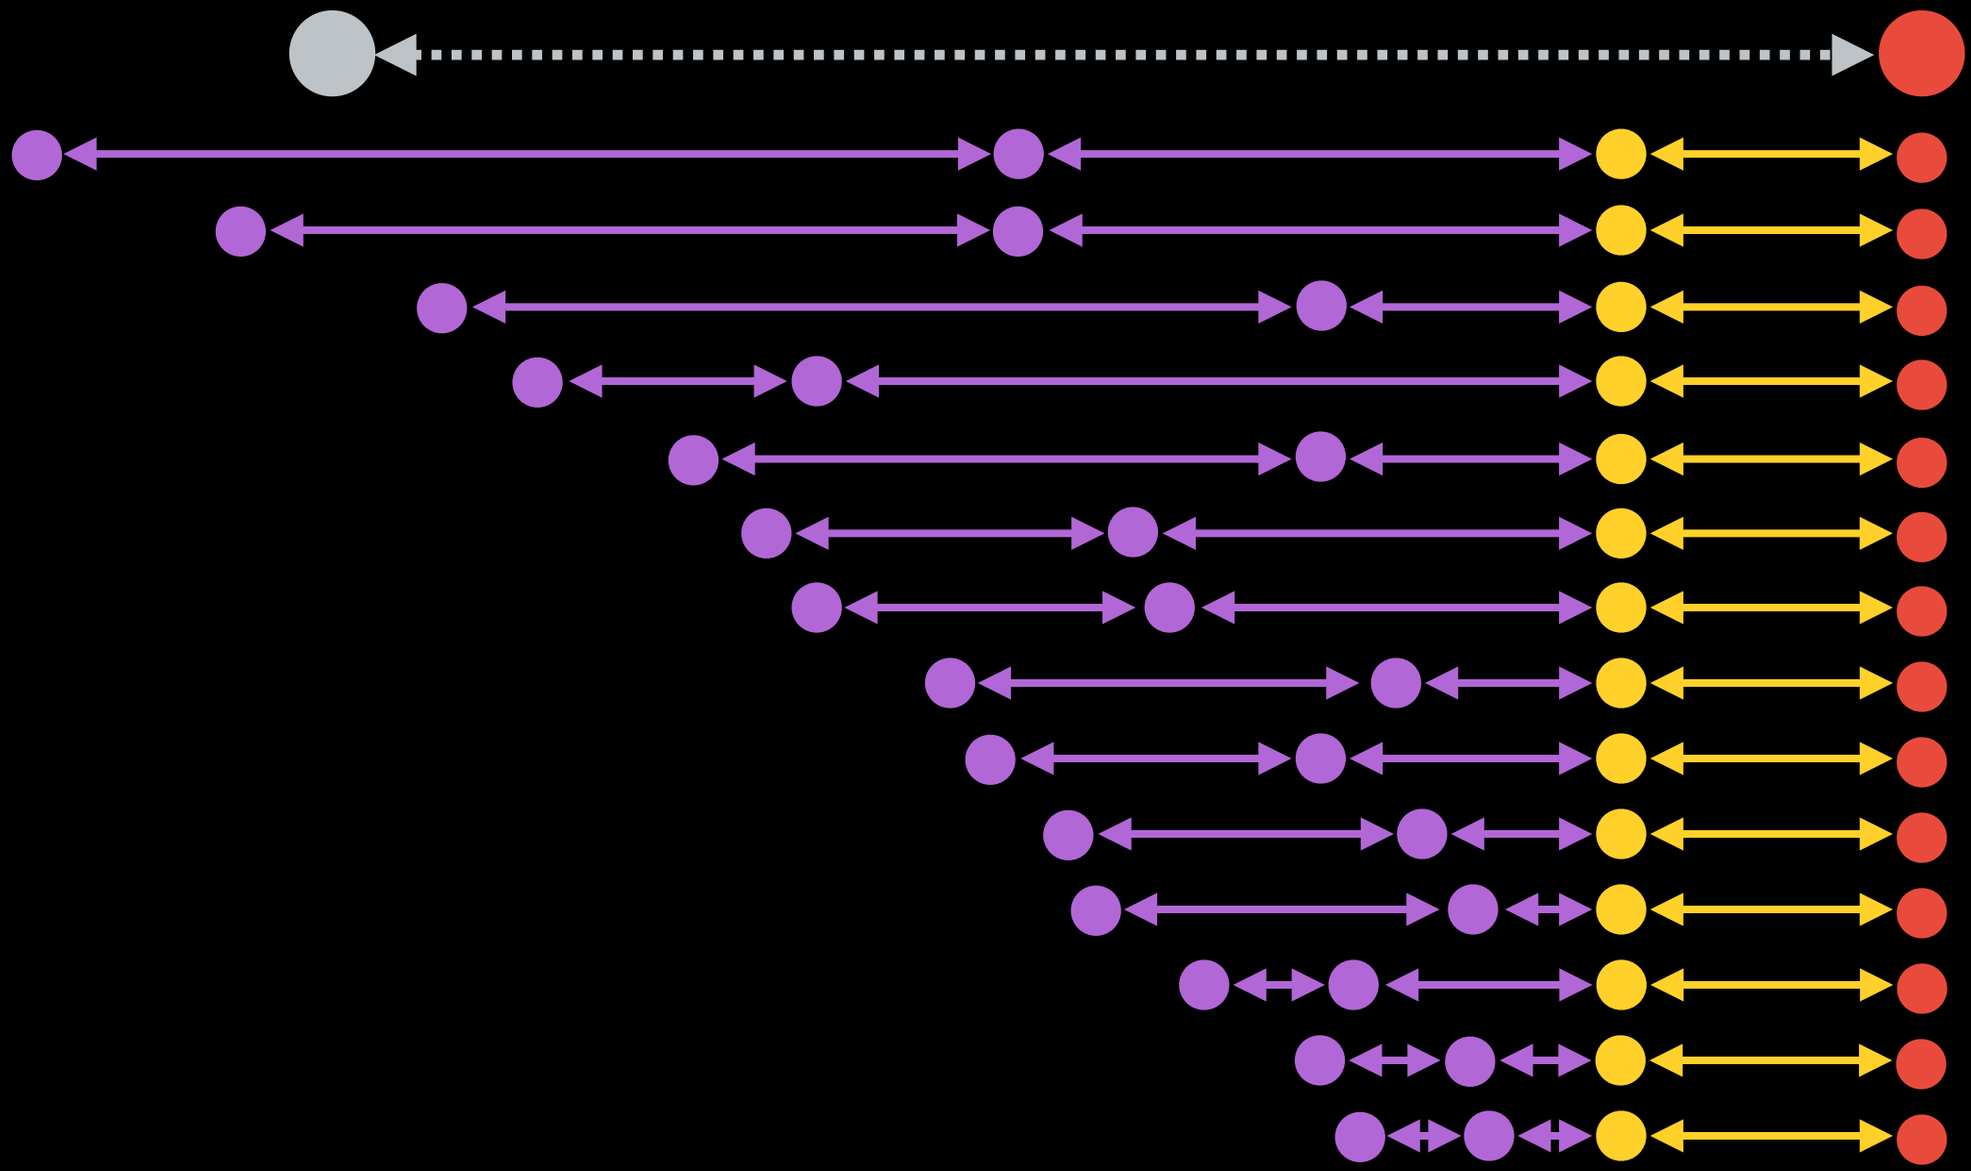

Client's RTT to the entry node is unknown

Pre-measured with Ting

Directly measured

Ruling out nodes without probing them

Reason about what the client → entry RTT would have to be

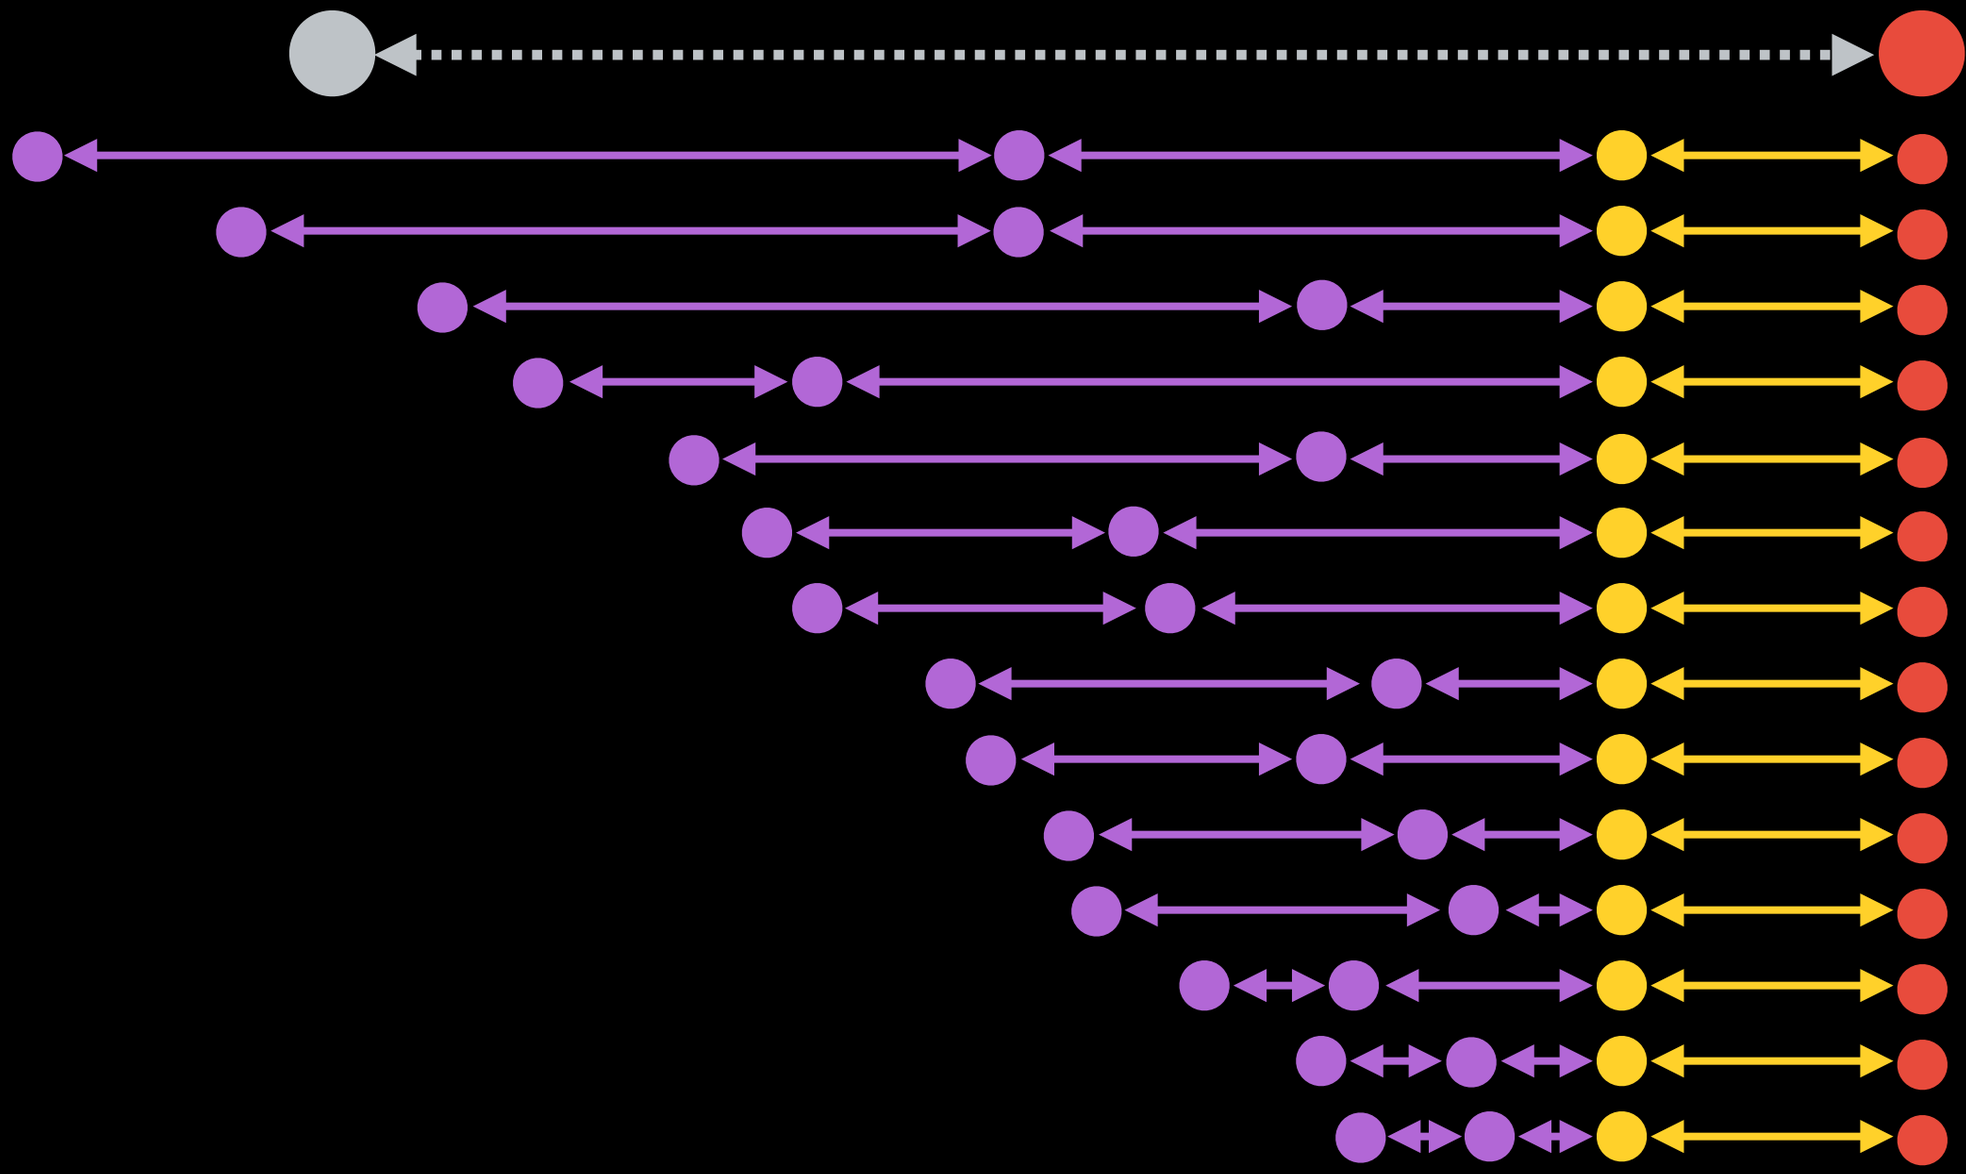

Client's RTT to the entry node is unknown

Pre-measured with Ting

Directly measured

Ruling out nodes without probing them

Reason about what the client → entry RTT would have to be

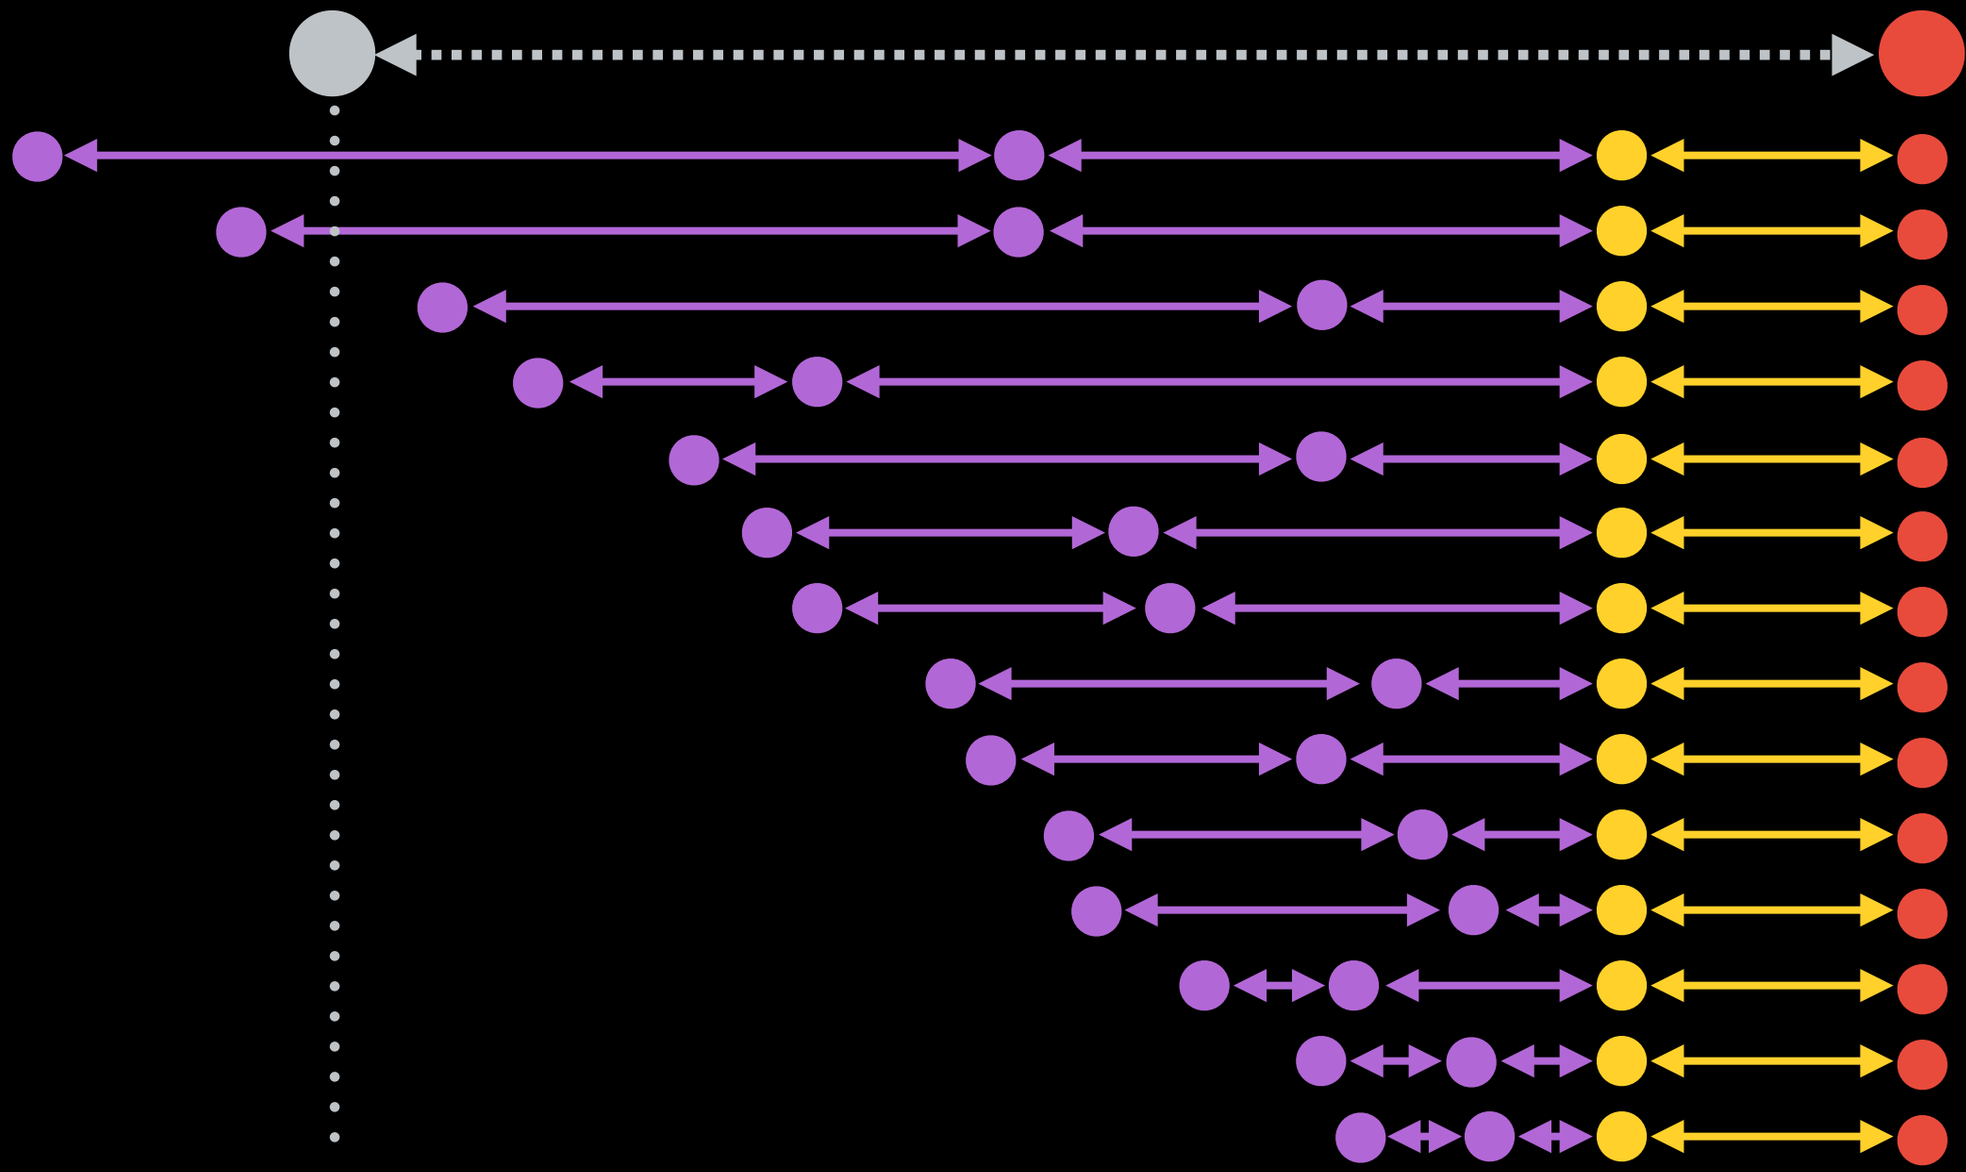

Client's RTT to the entry node is unknown

Pre-measured with Ting

Directly measured

Ruling out nodes without probing them

Reason about what the client → entry RTT would have to be

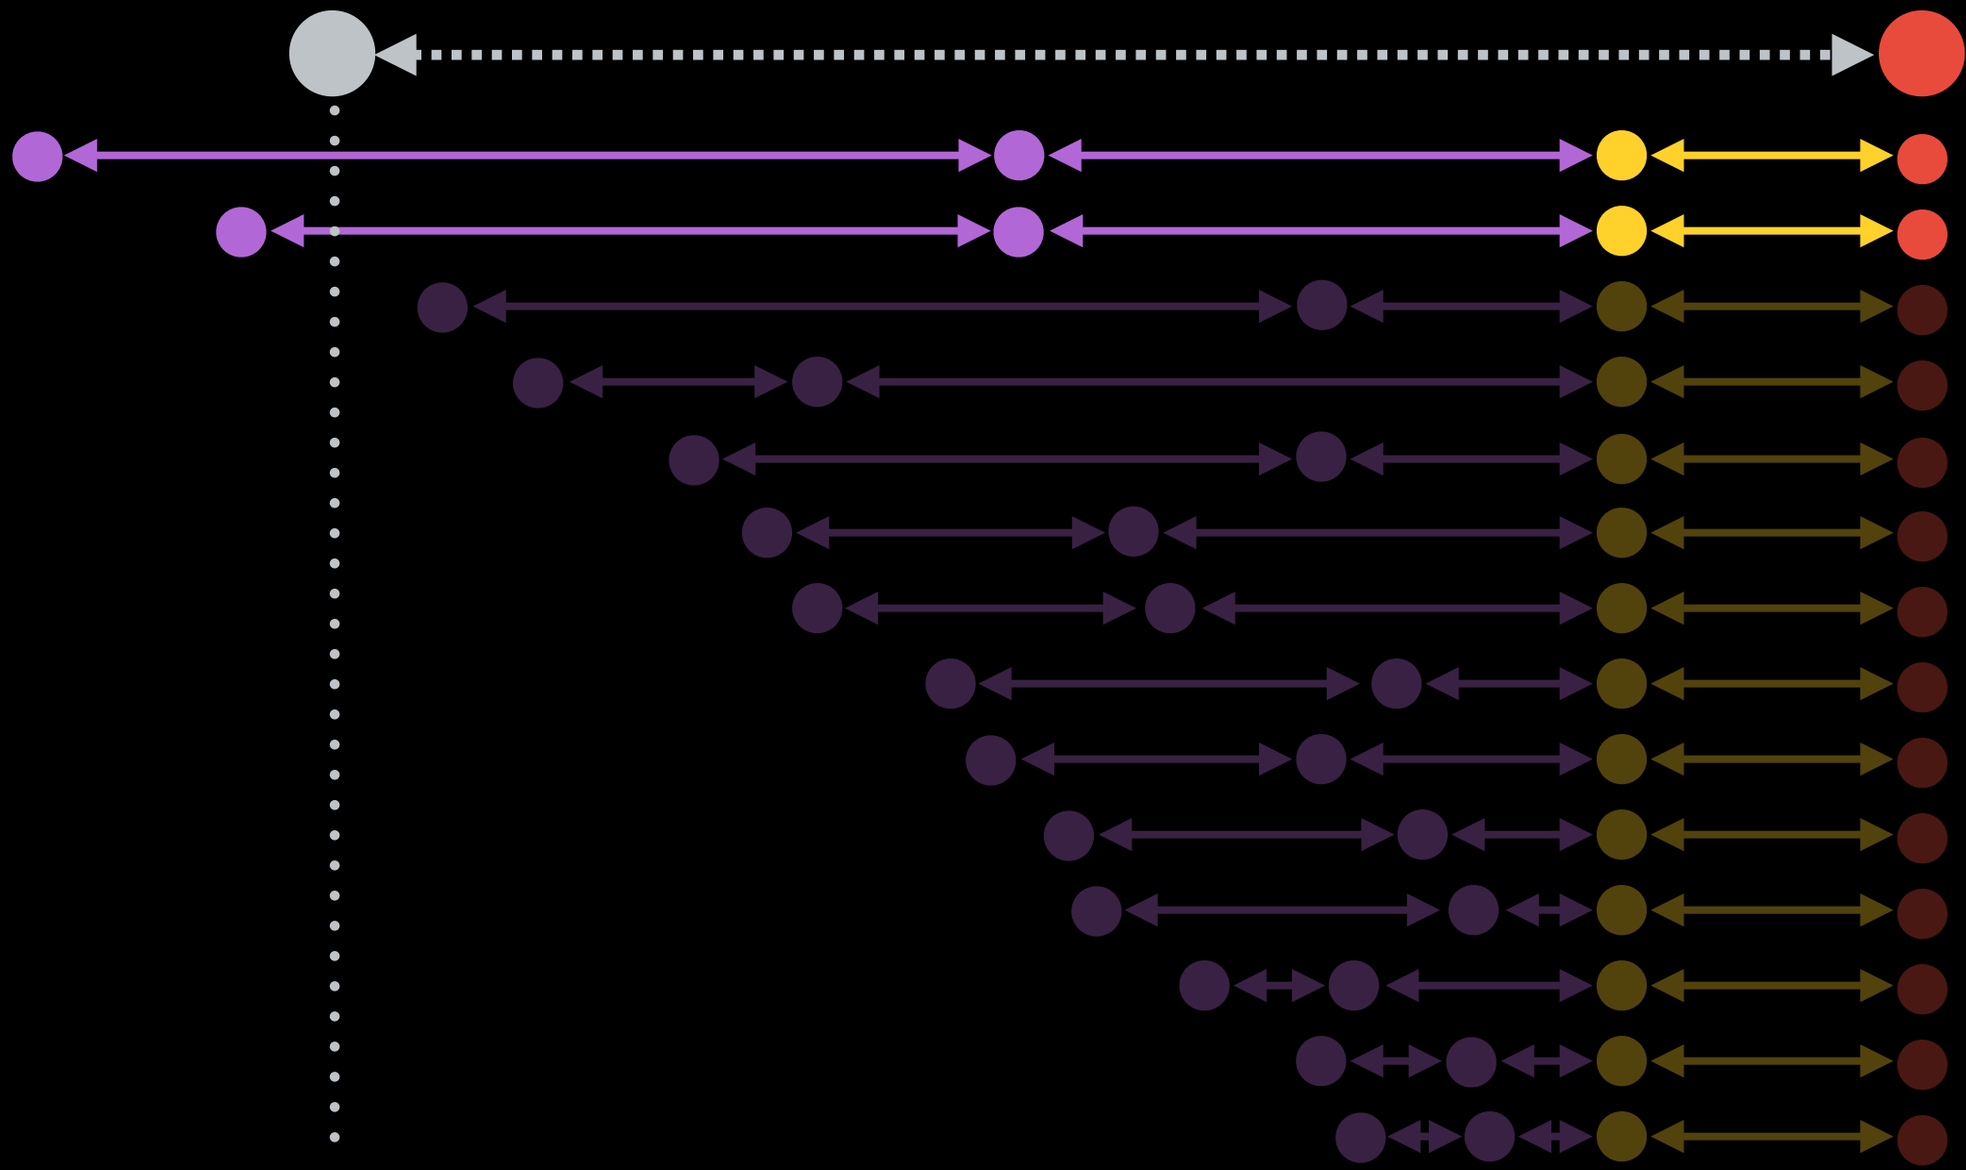

Client's RTT to the entry node is unknown

Pre-measured with Ting

Directly measured

Ruling out nodes without probing them

Reason about what the client → entry RTT would have to be

Client's RTT to the entry node is unknown

Pre-measured with Ting

Directly measured

Ruling out nodes without probing them

Reason about what the client → entry RTT would have to be

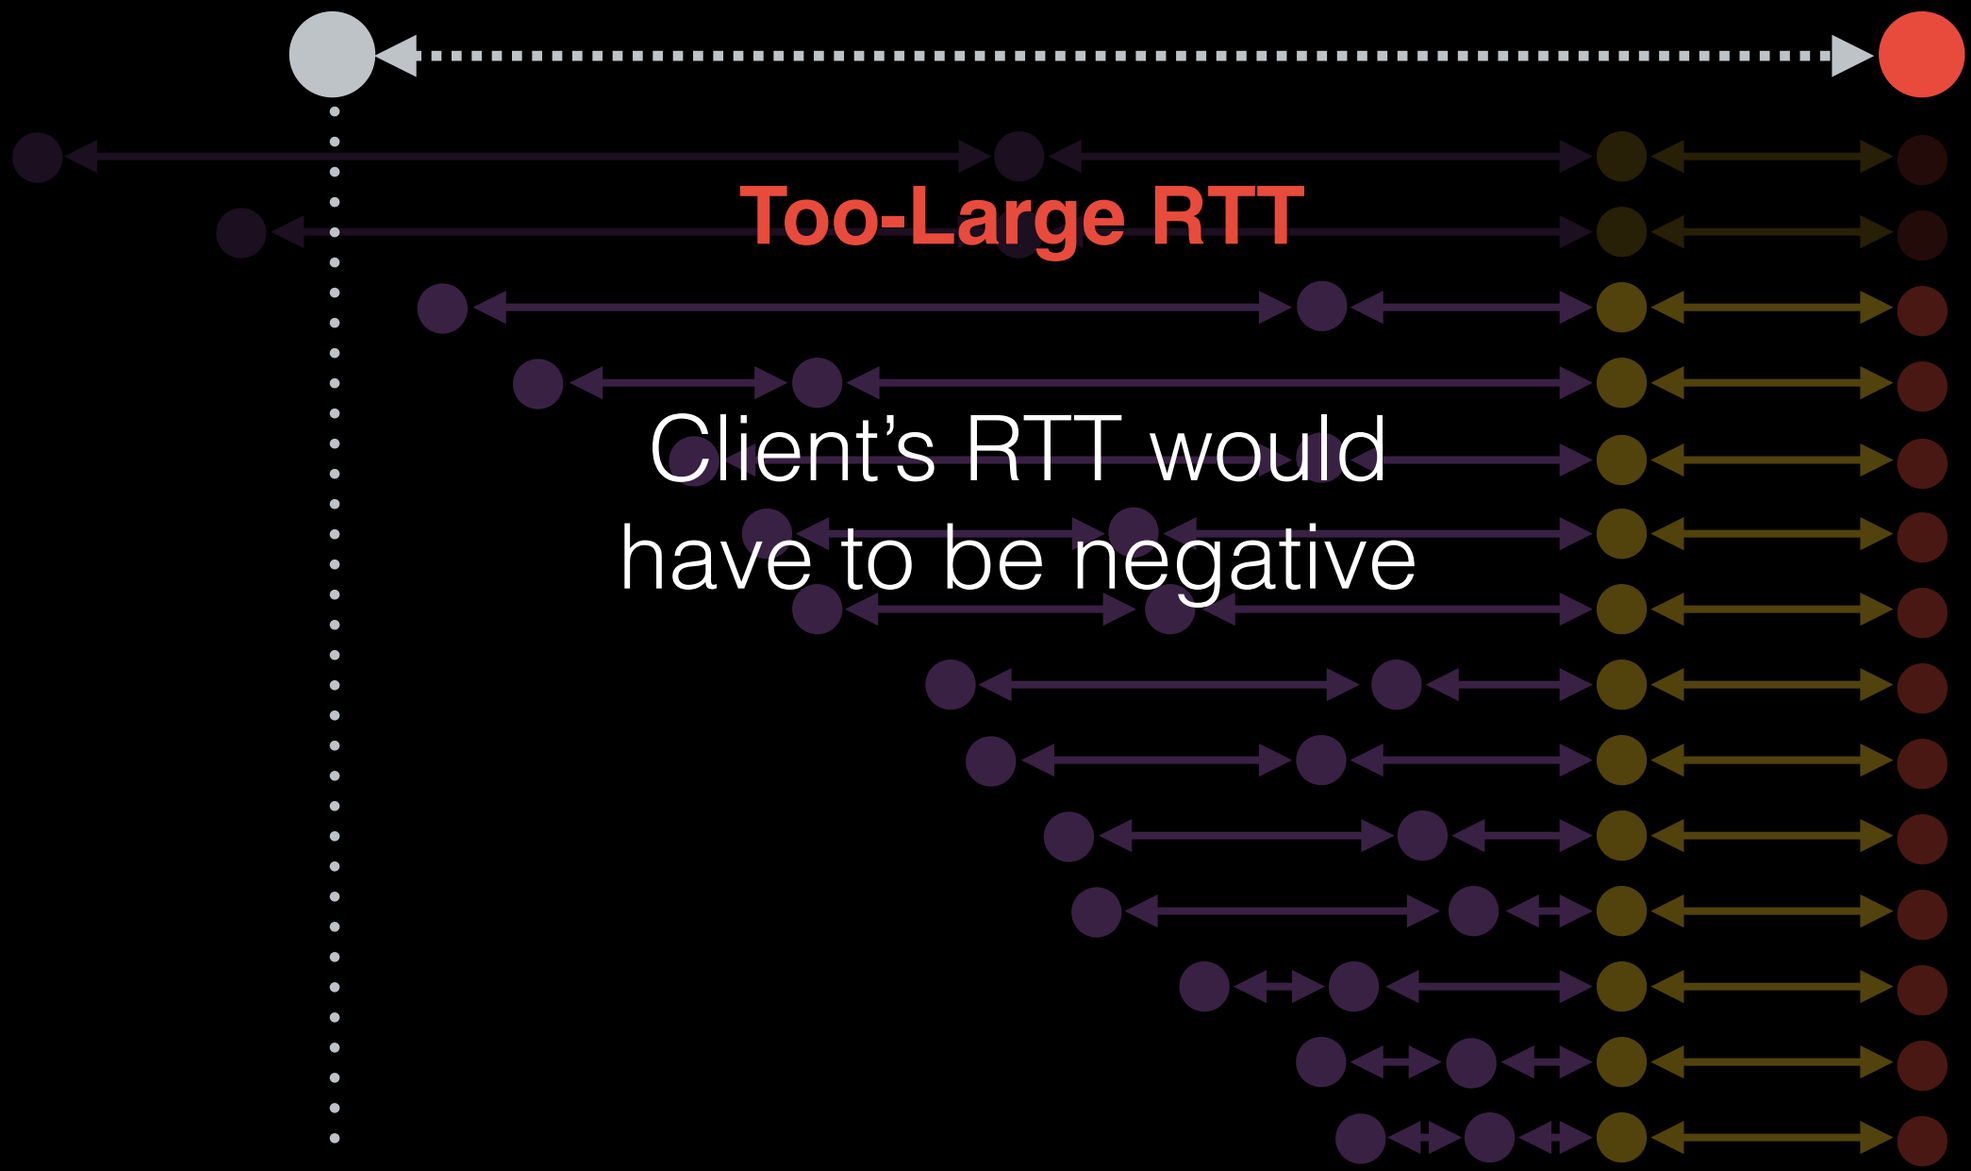

Client's RTT to the entry node is unknown

Pre-measured with Ting

Directly measured

Ruling out too-large RTTs

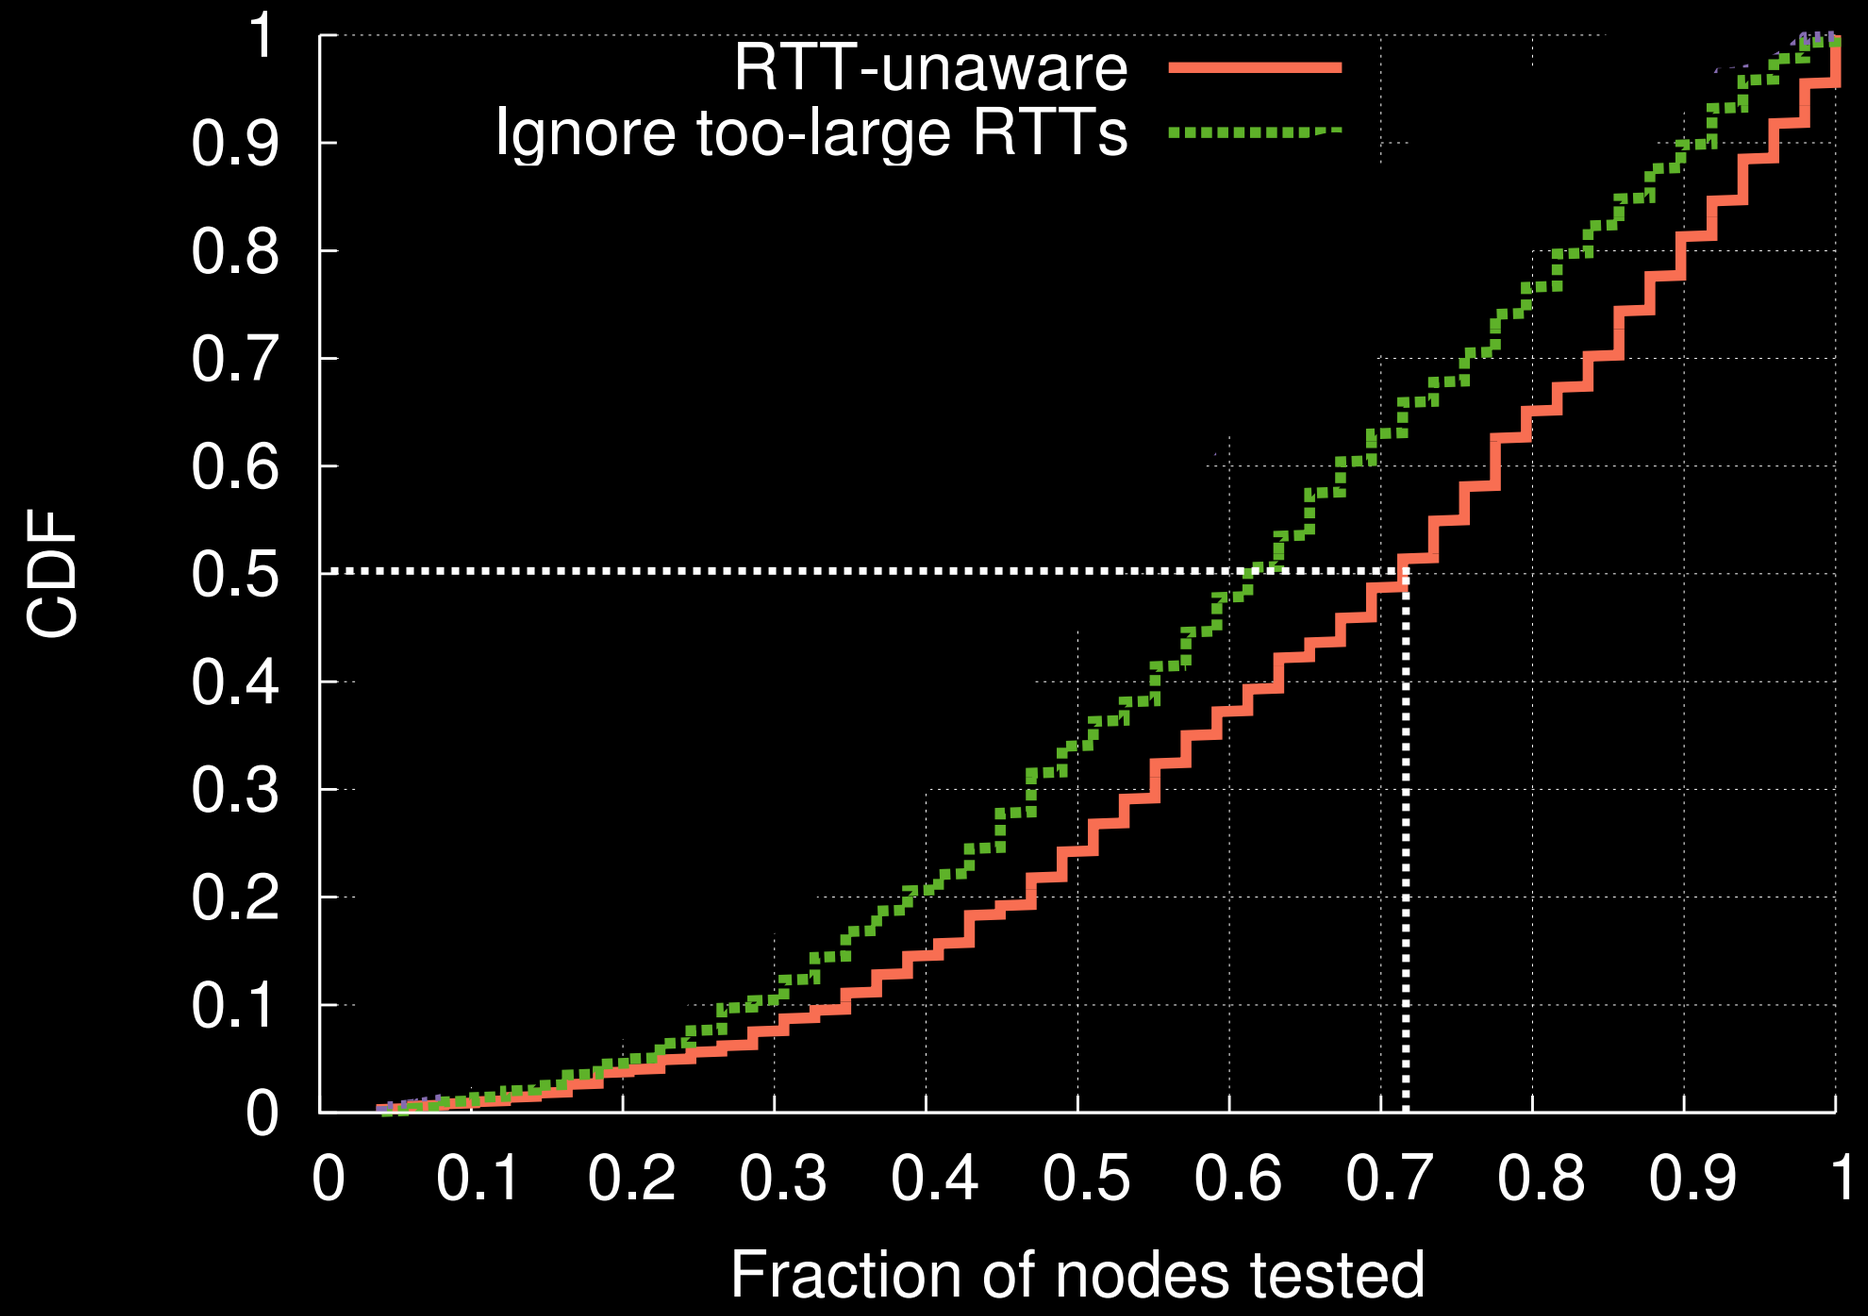

Ruling out too-large RTTs

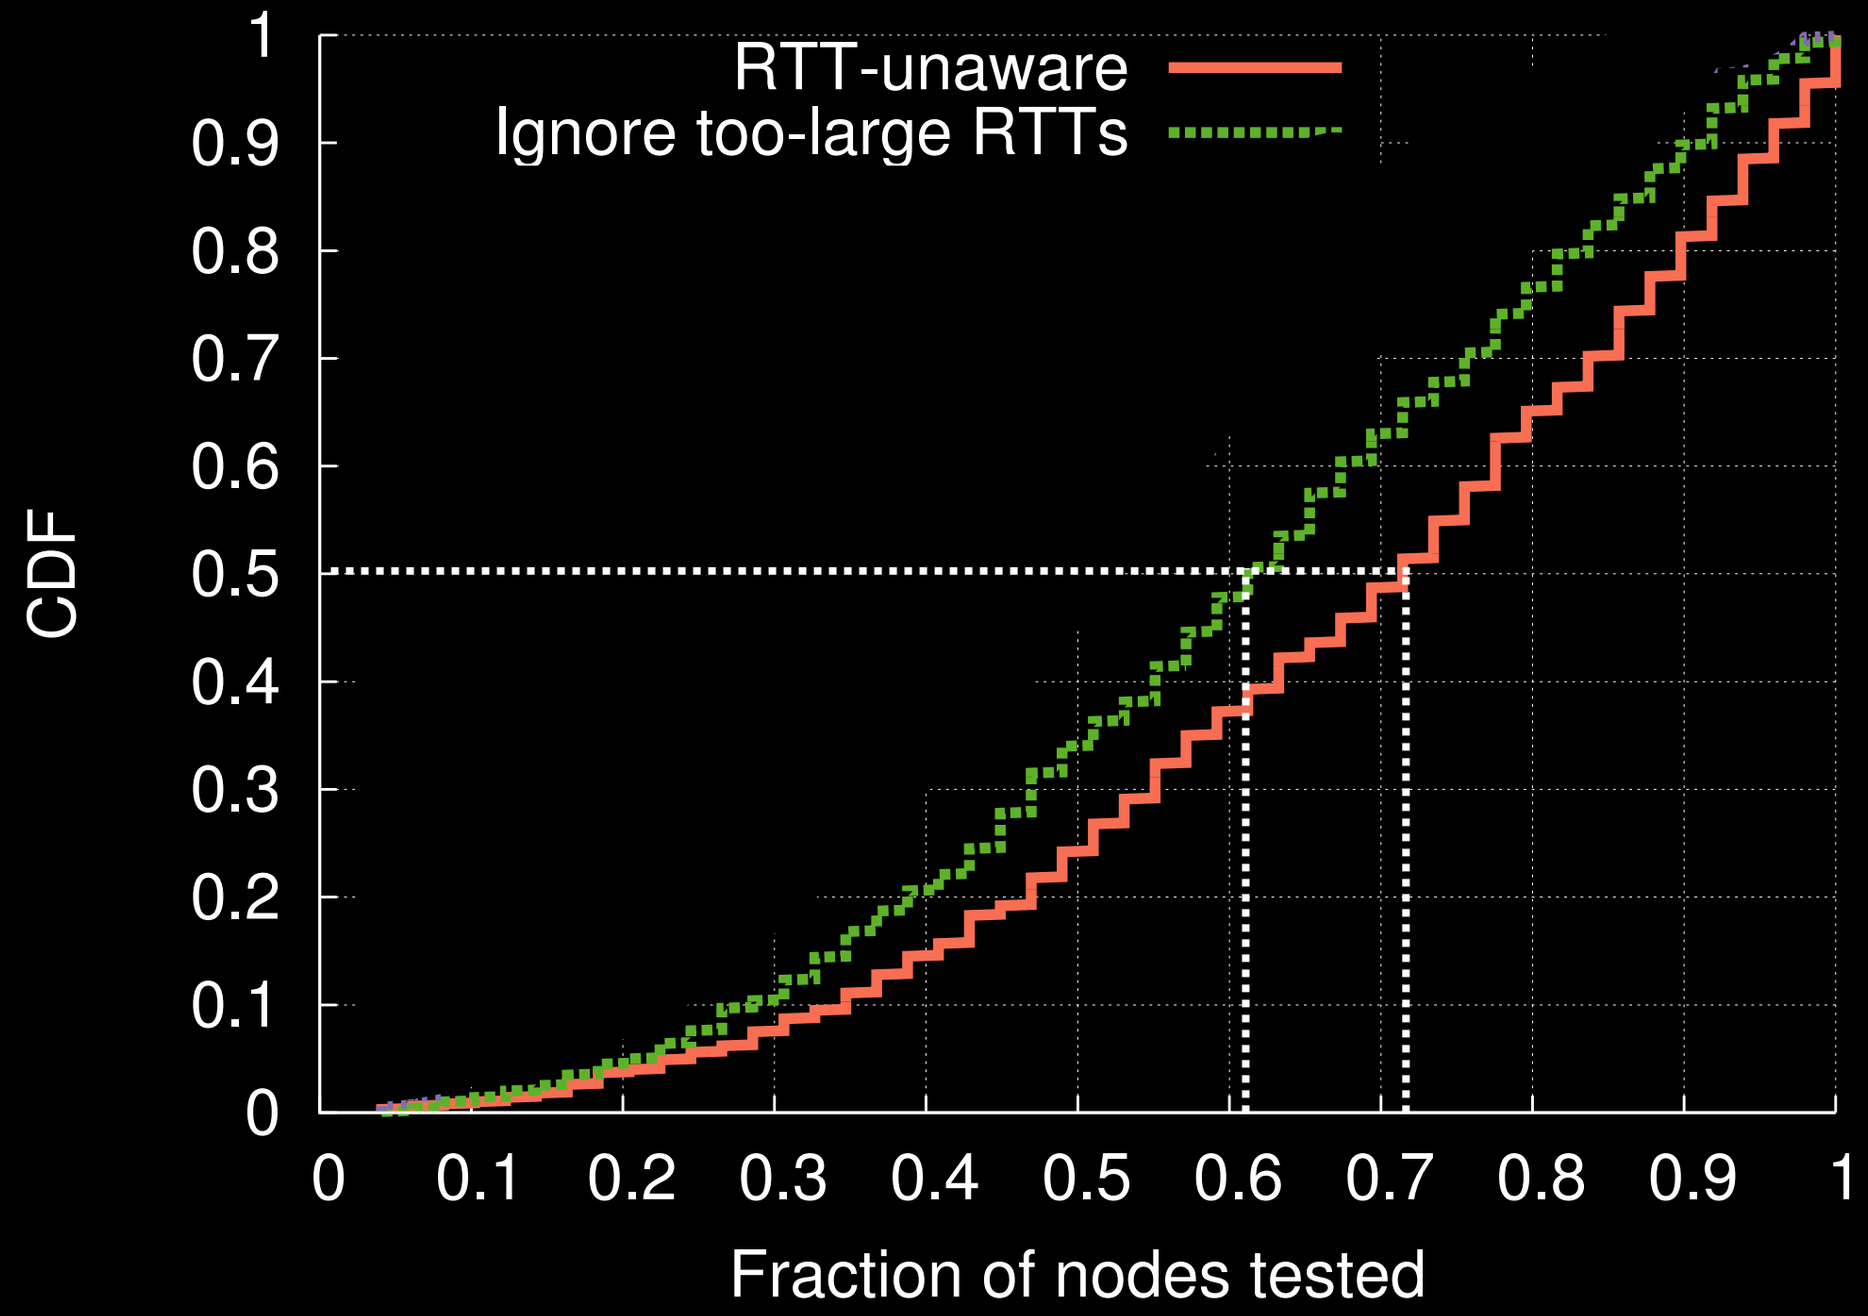

Informed target selection

Probe nodes according to probability that they are on the **circuit**

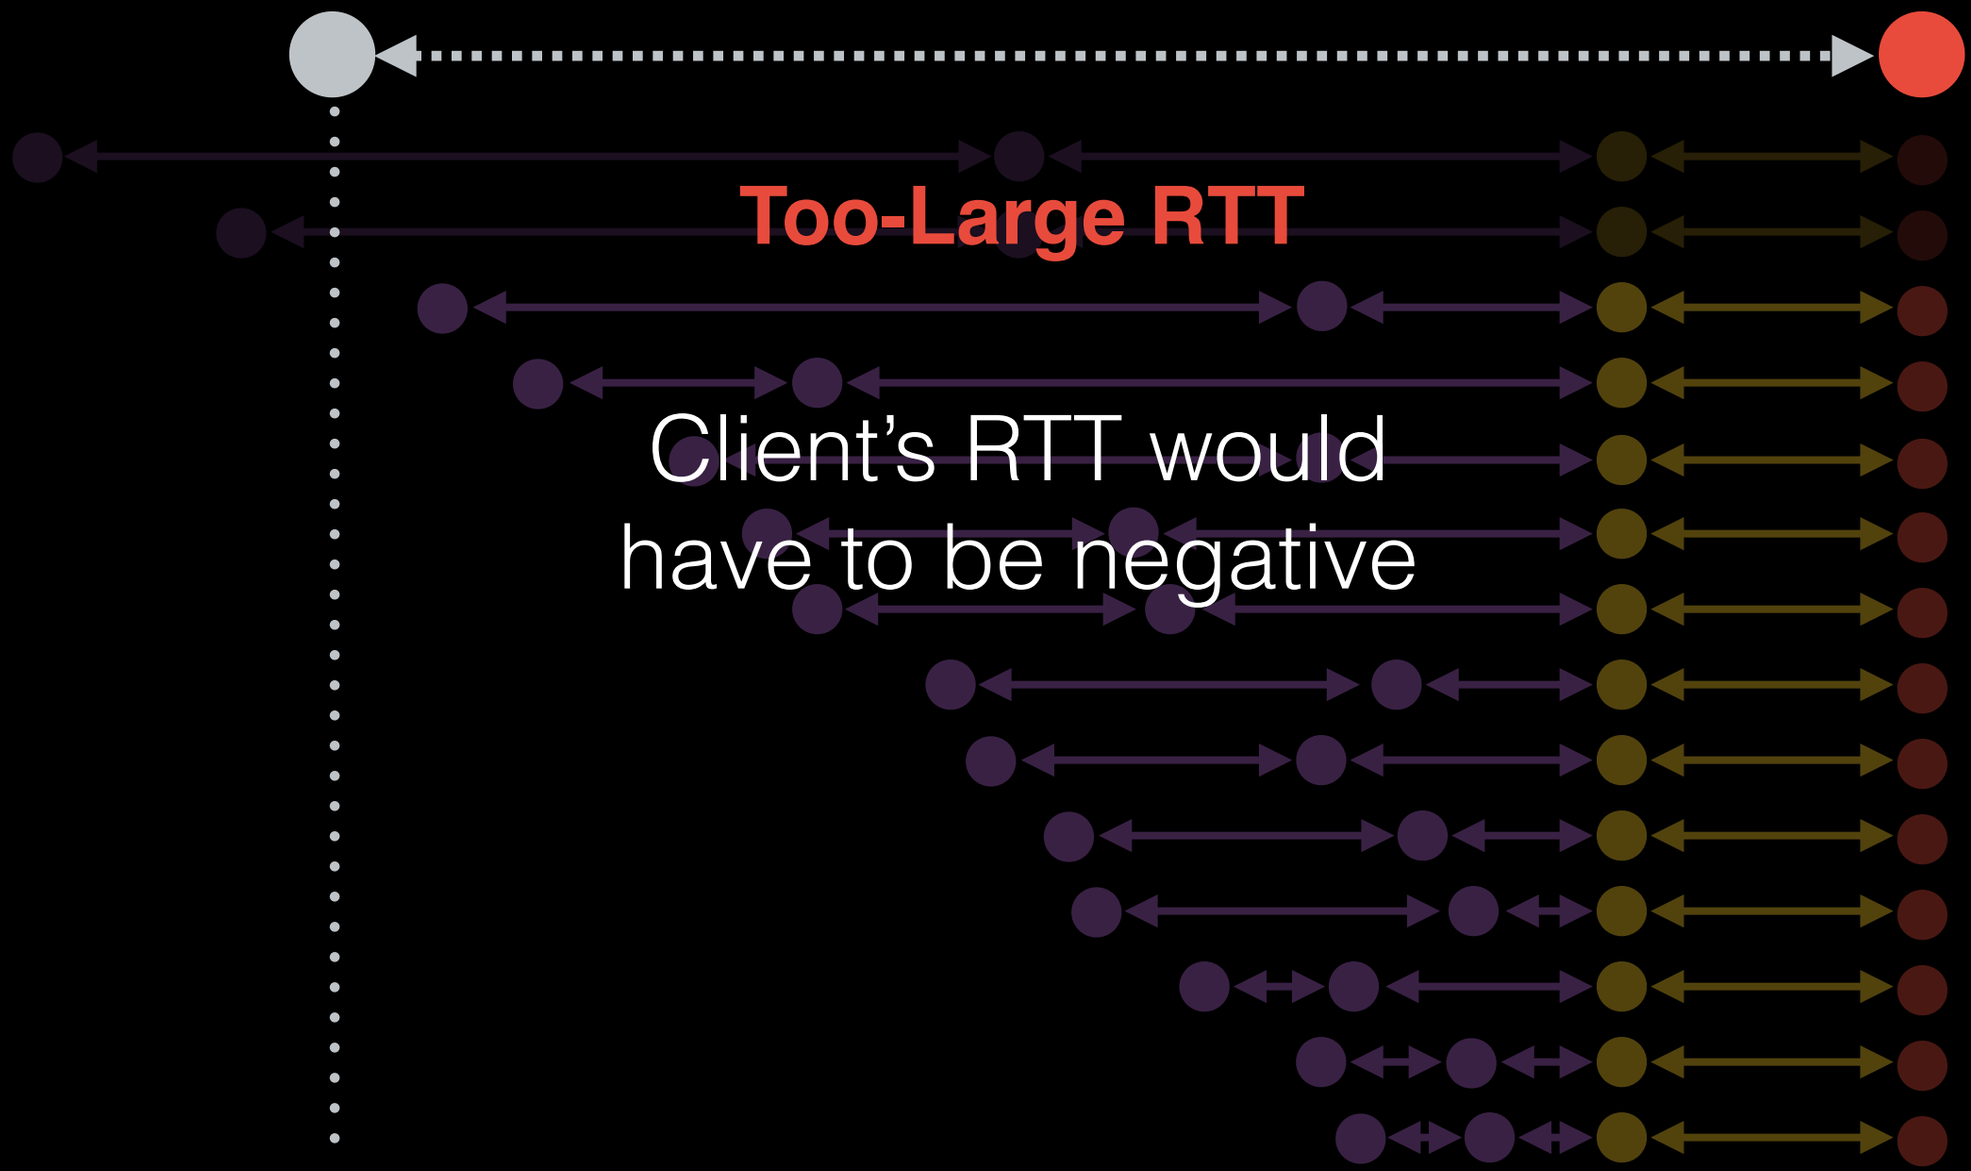

Client's RTT to the entry node is unknown

Pre-measured with Ting

Directly measured

Informed target selection

Probe the more likely circuits first

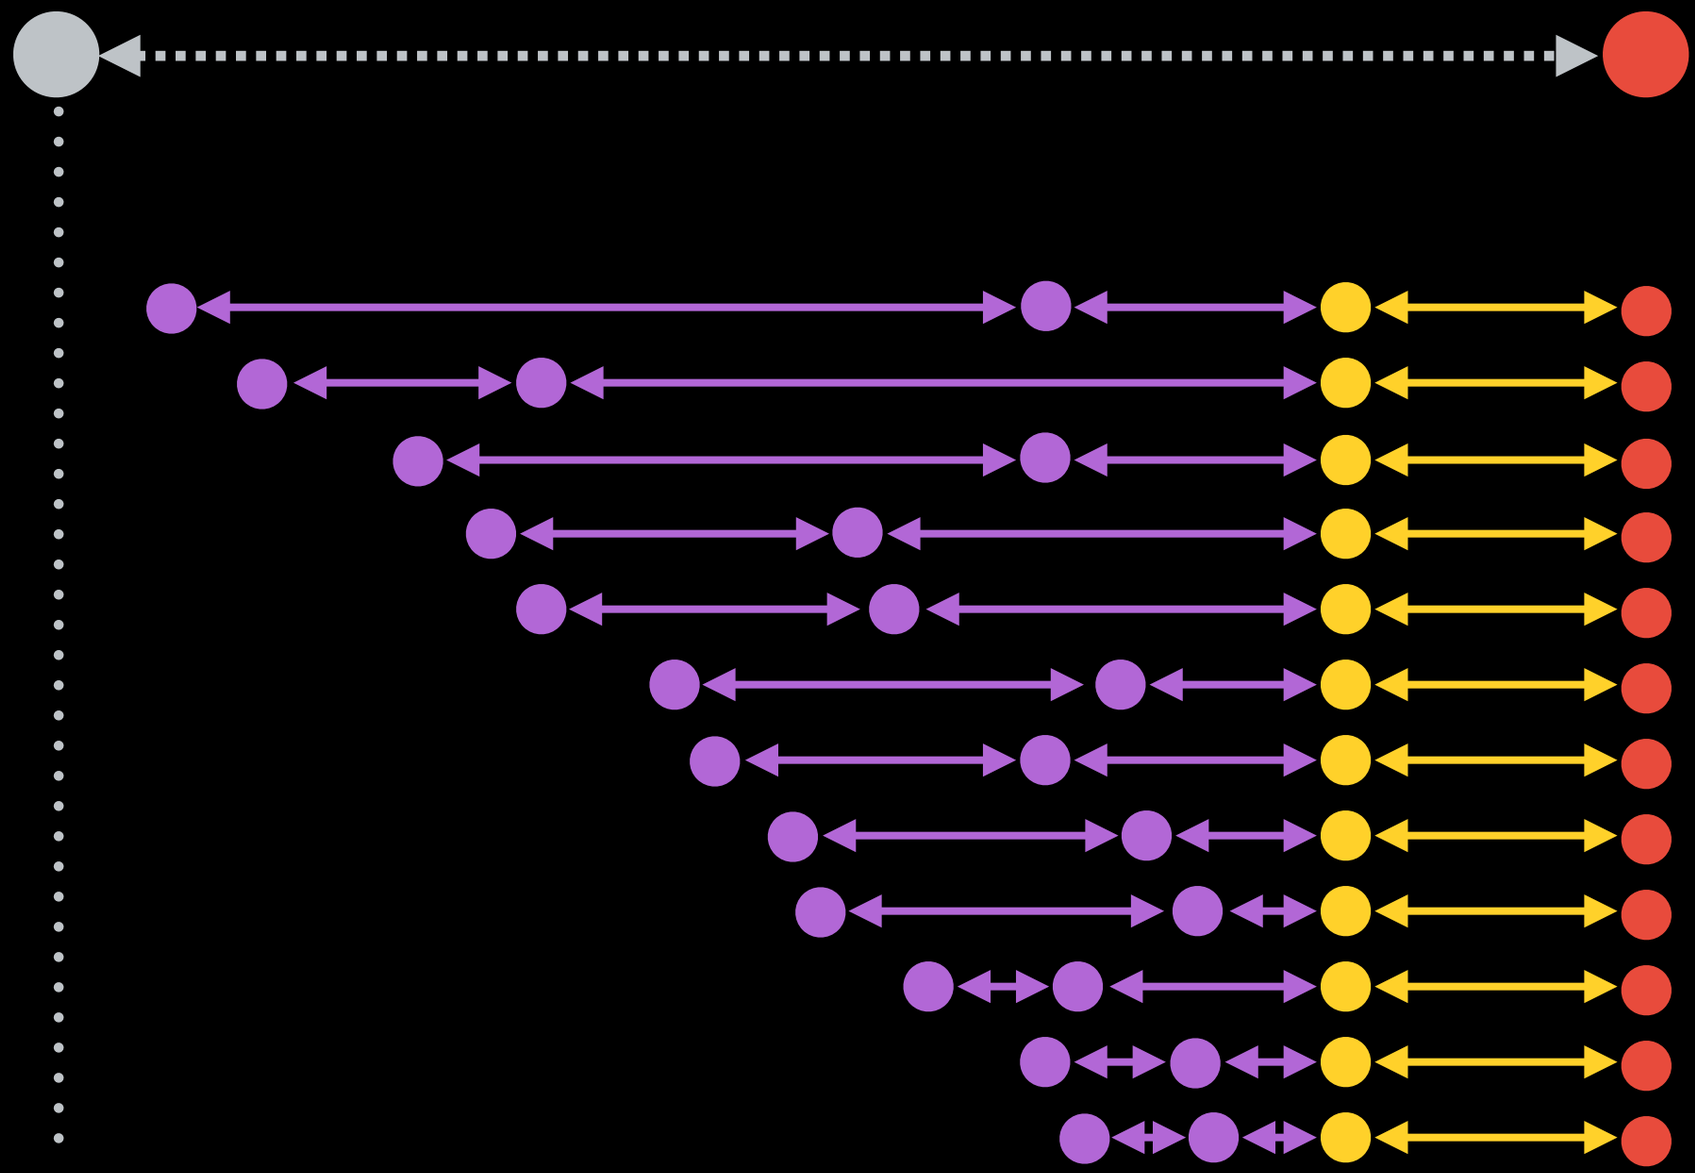

Client's RTT to the entry node is unknown

Pre-measured with Ting

Directly measured

Informed target selection

Probe the more likely circuits first

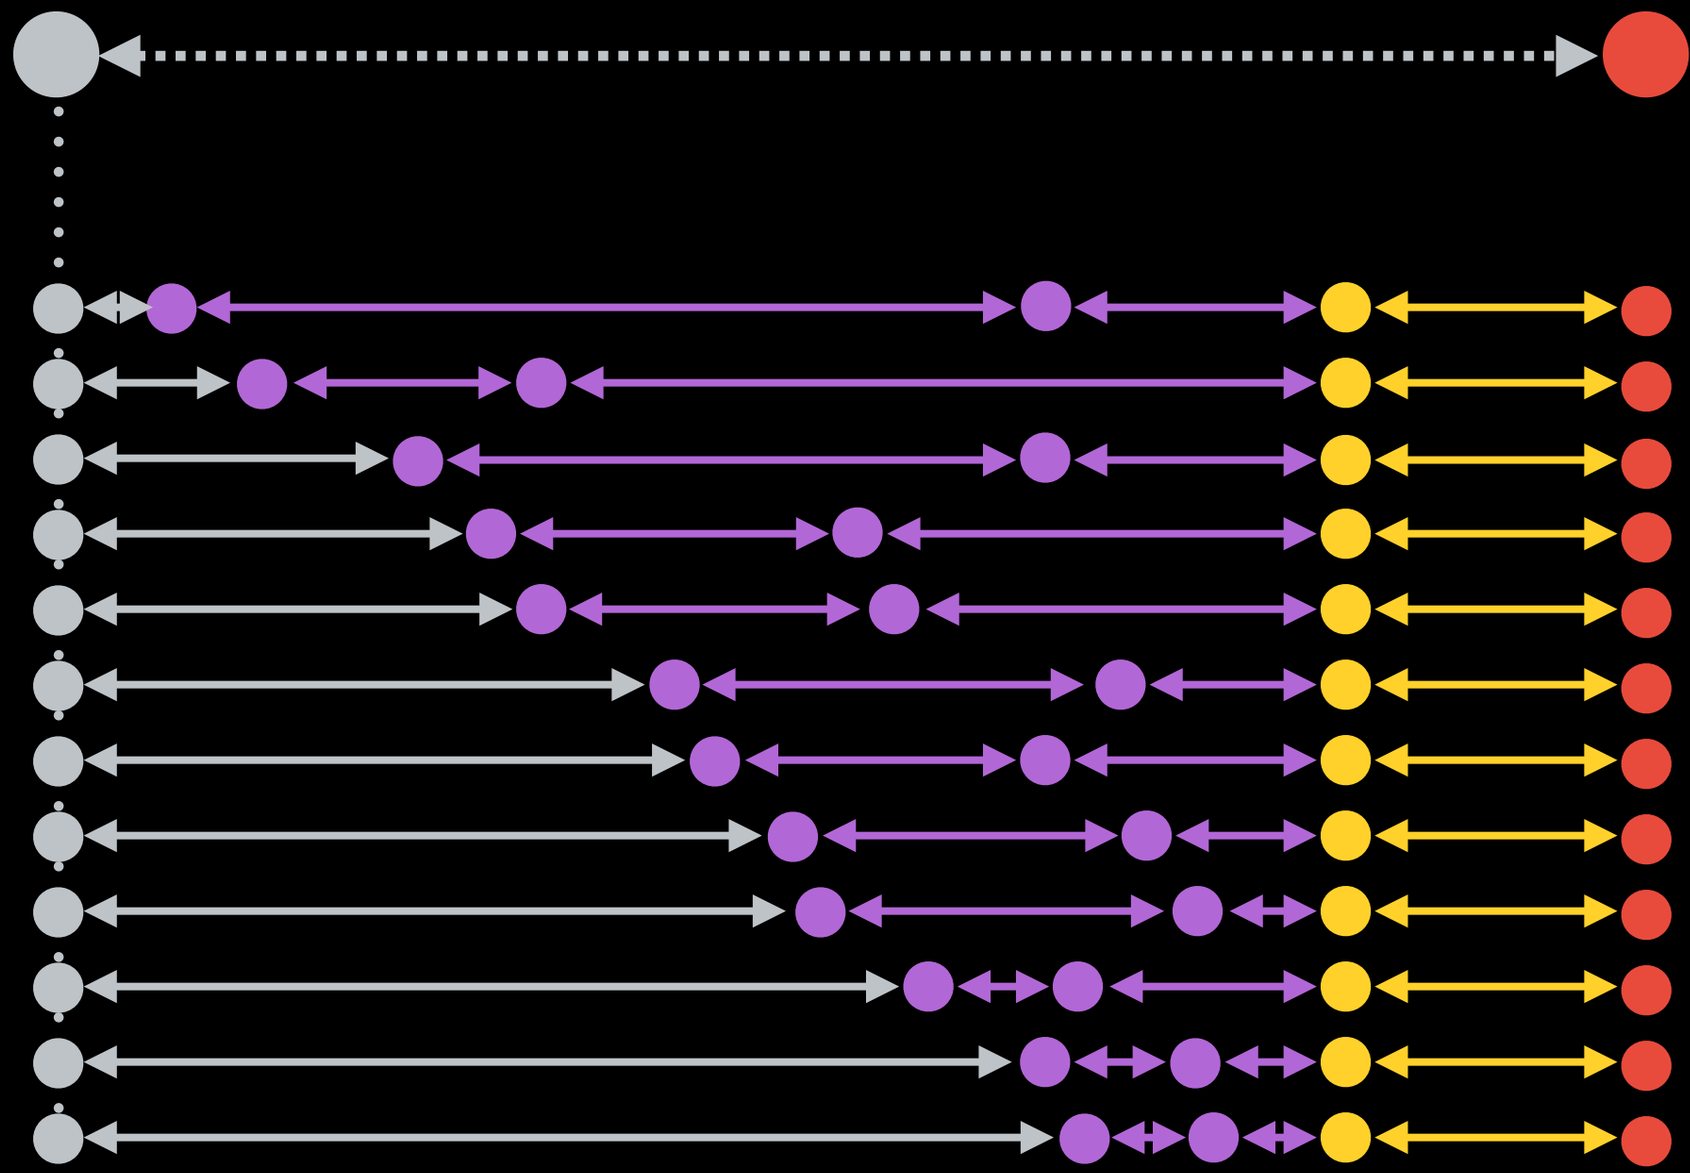

Client's RTT to the entry node is unknown

Pre-measured with Ting

Directly measured

Informed target selection

Probe the more likely circuits first

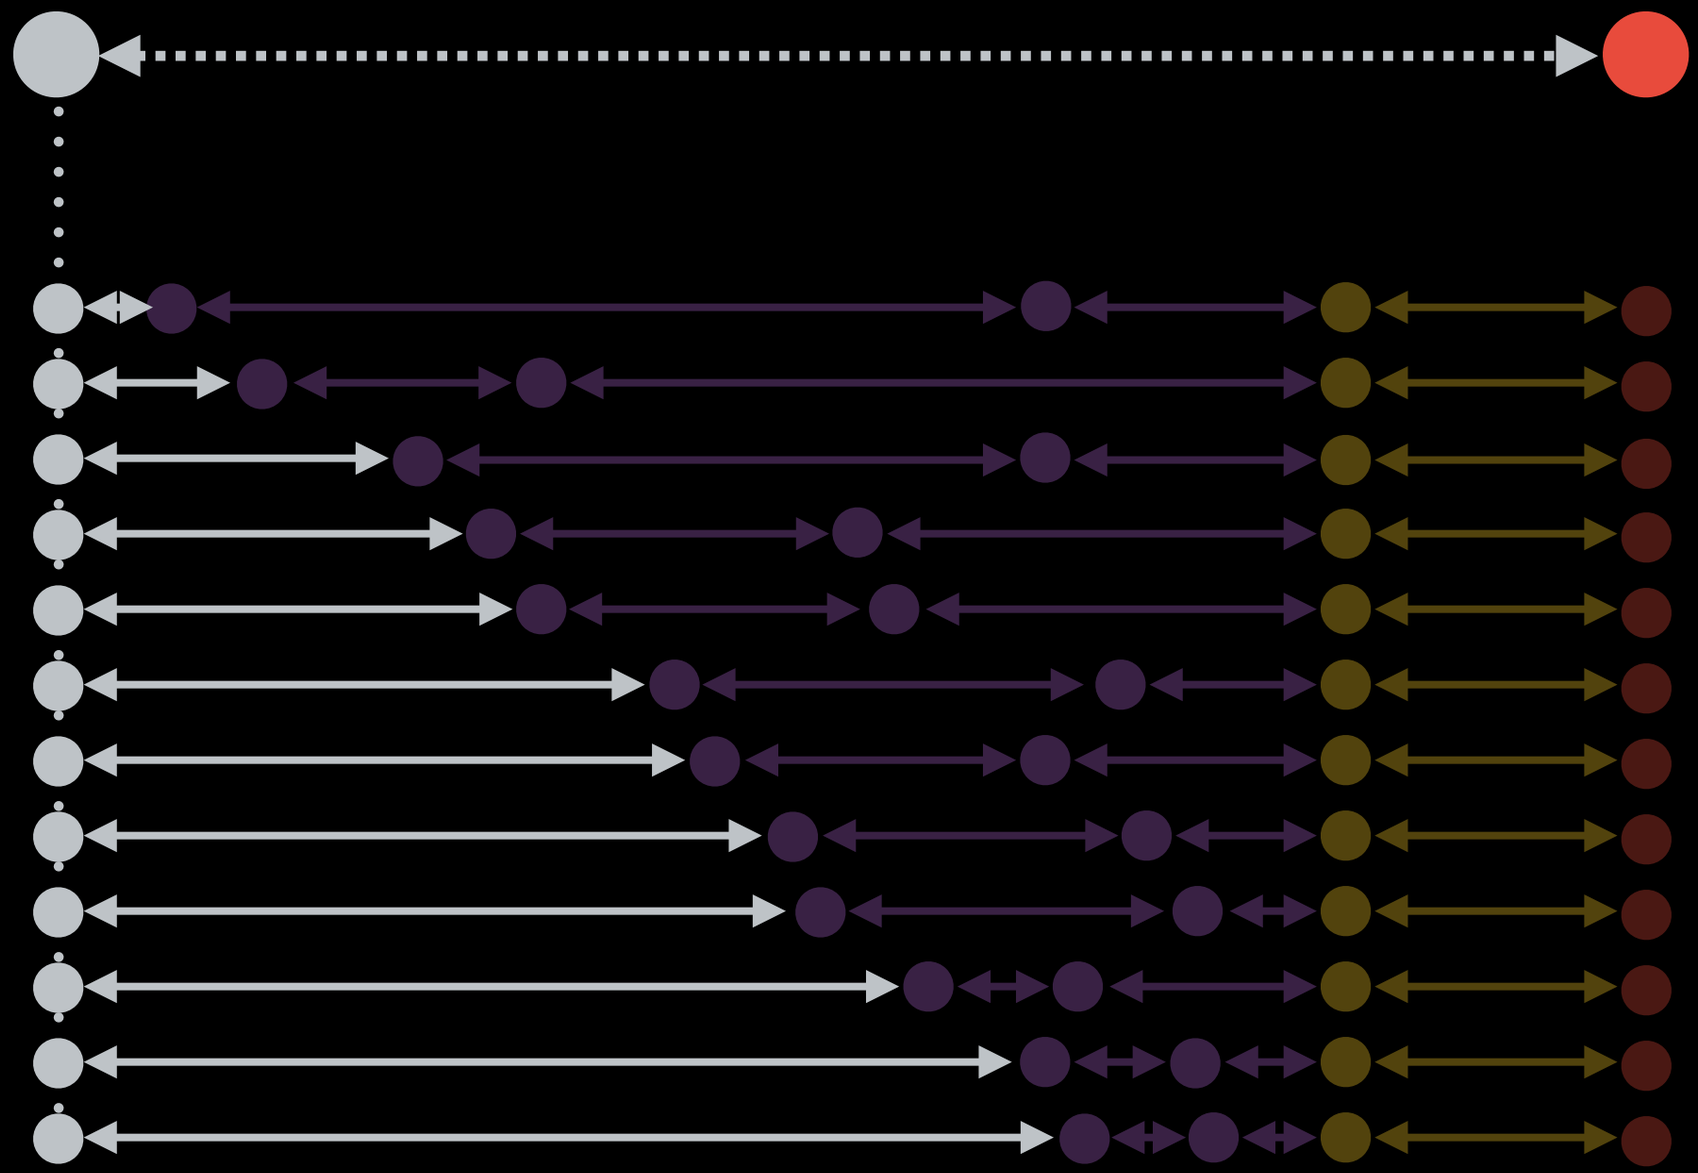

Client's RTT to the entry node is unknown

Pre-measured with Ting

Directly measured

Informed target selection

Probe the more likely circuits first

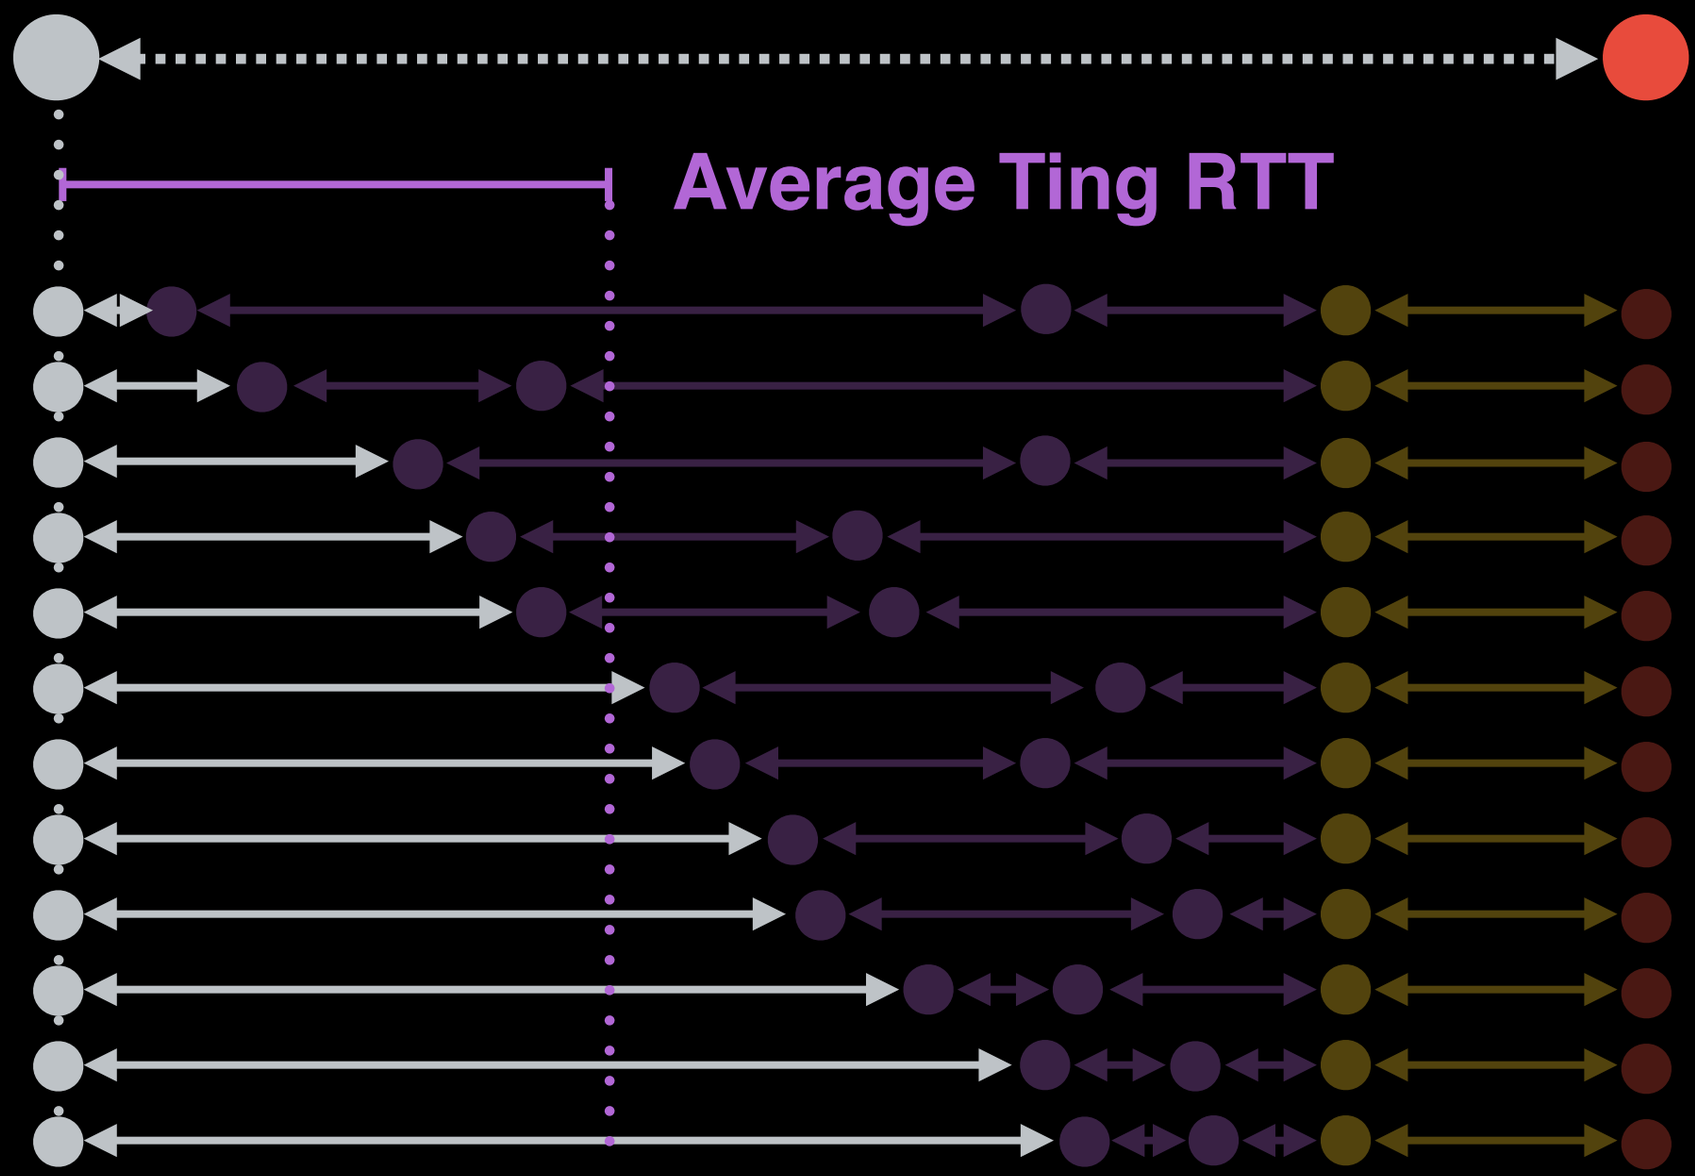

Client's RTT to the entry node is unknown

Pre-measured with Ting

Directly measured

Informed target selection

Probe the more likely circuits first

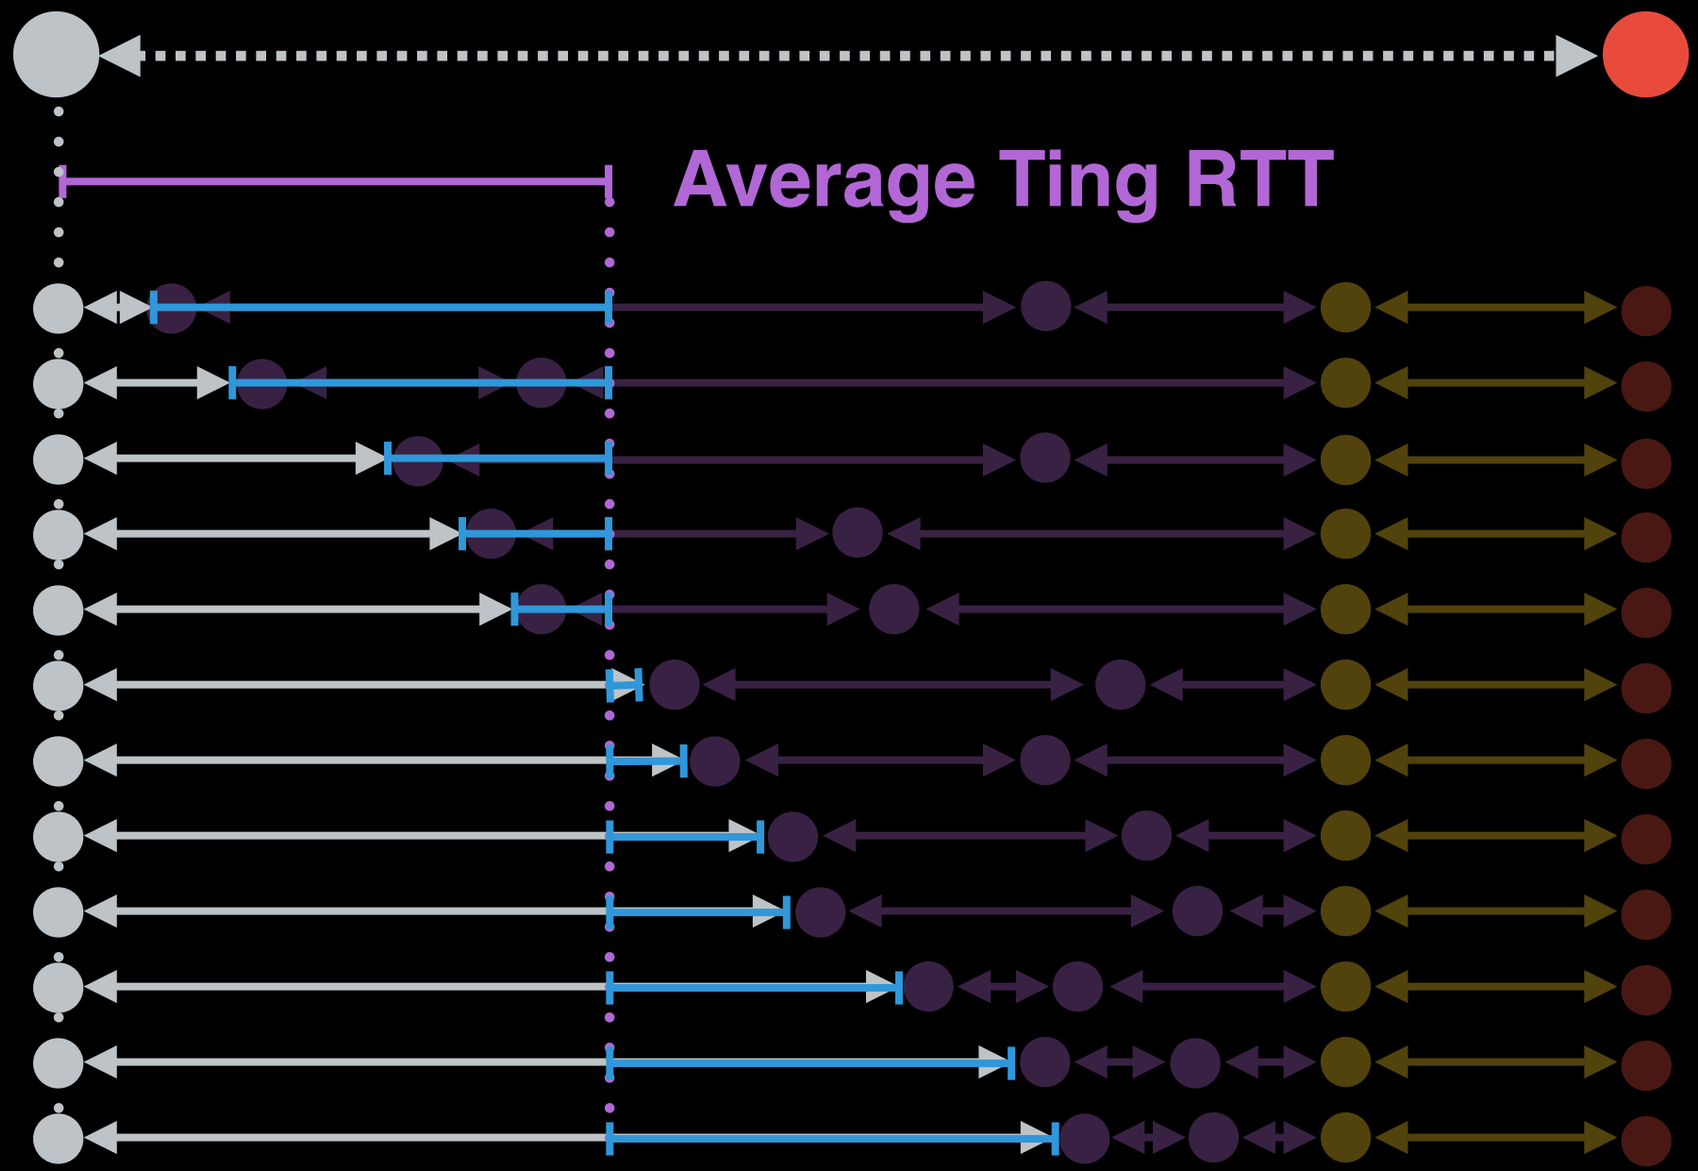

Client's RTT to the entry node is unknown

Pre-measured with Ting

Directly measured

Informed target selection

Probe the more likely circuits first

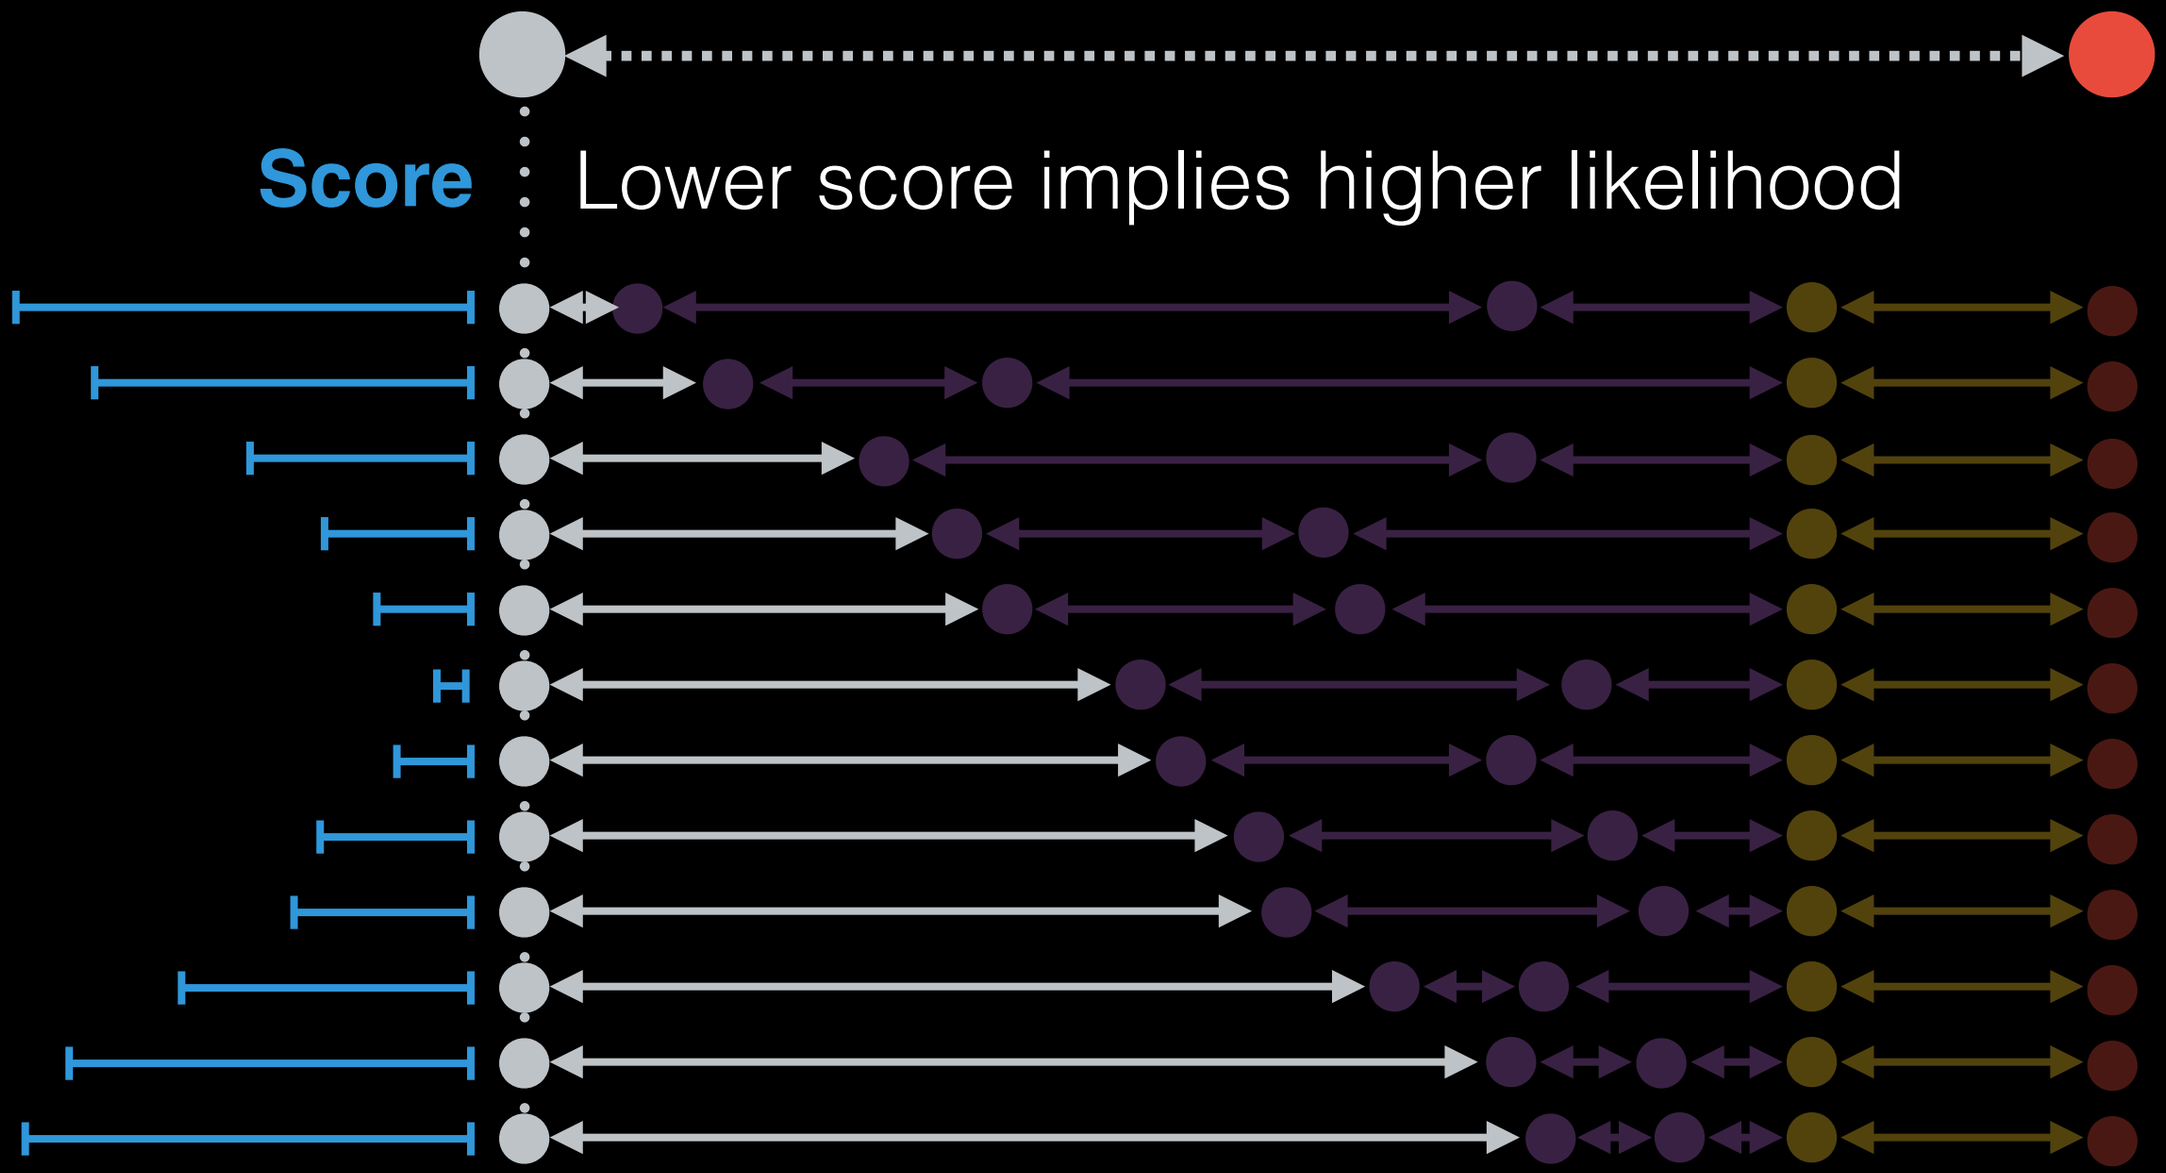

Informed target selection

Probe the more likely circuits first

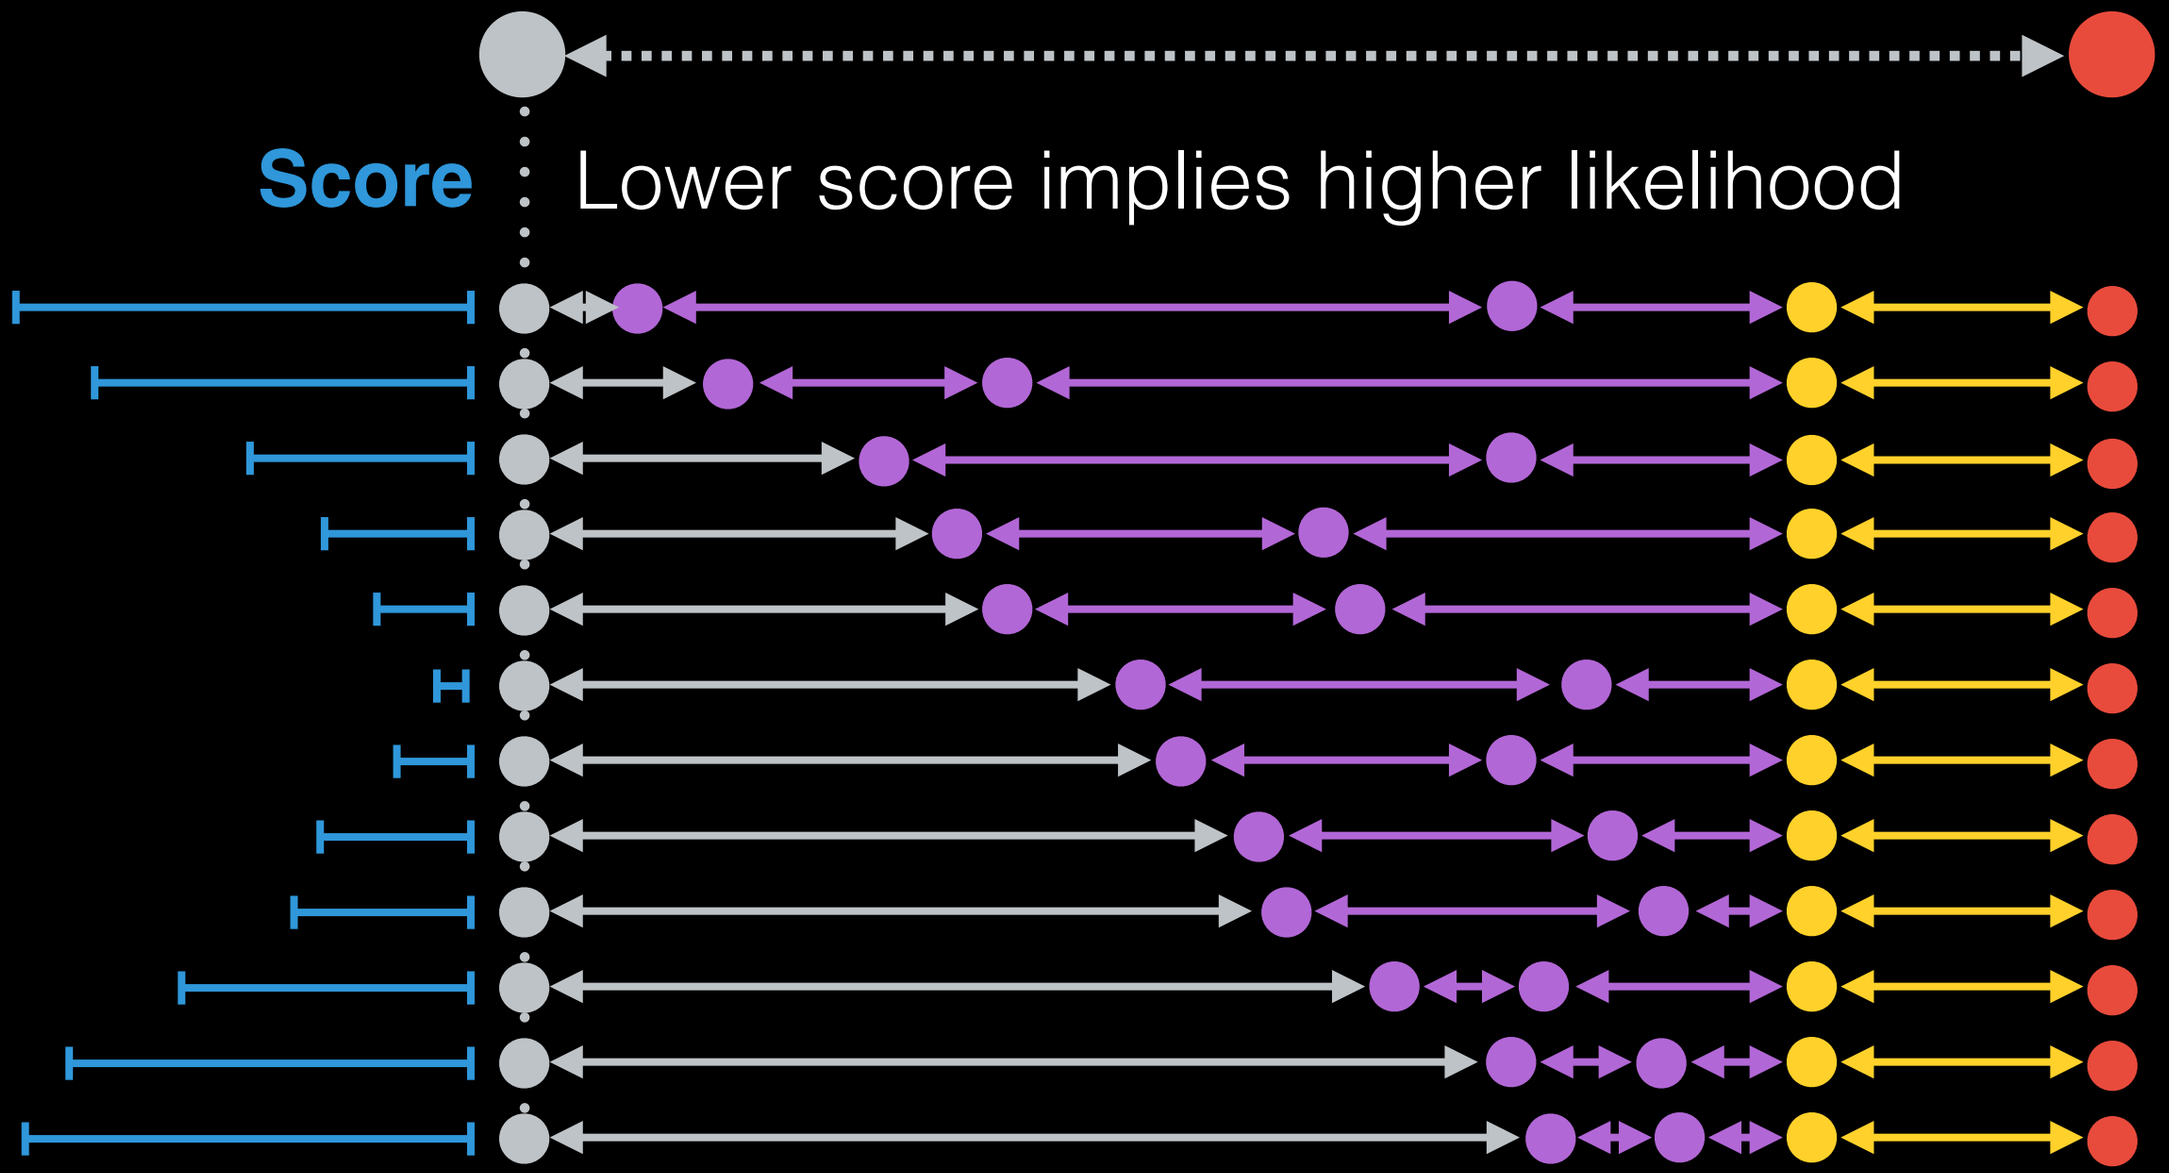

Informed target selection

Probe the more likely circuits first

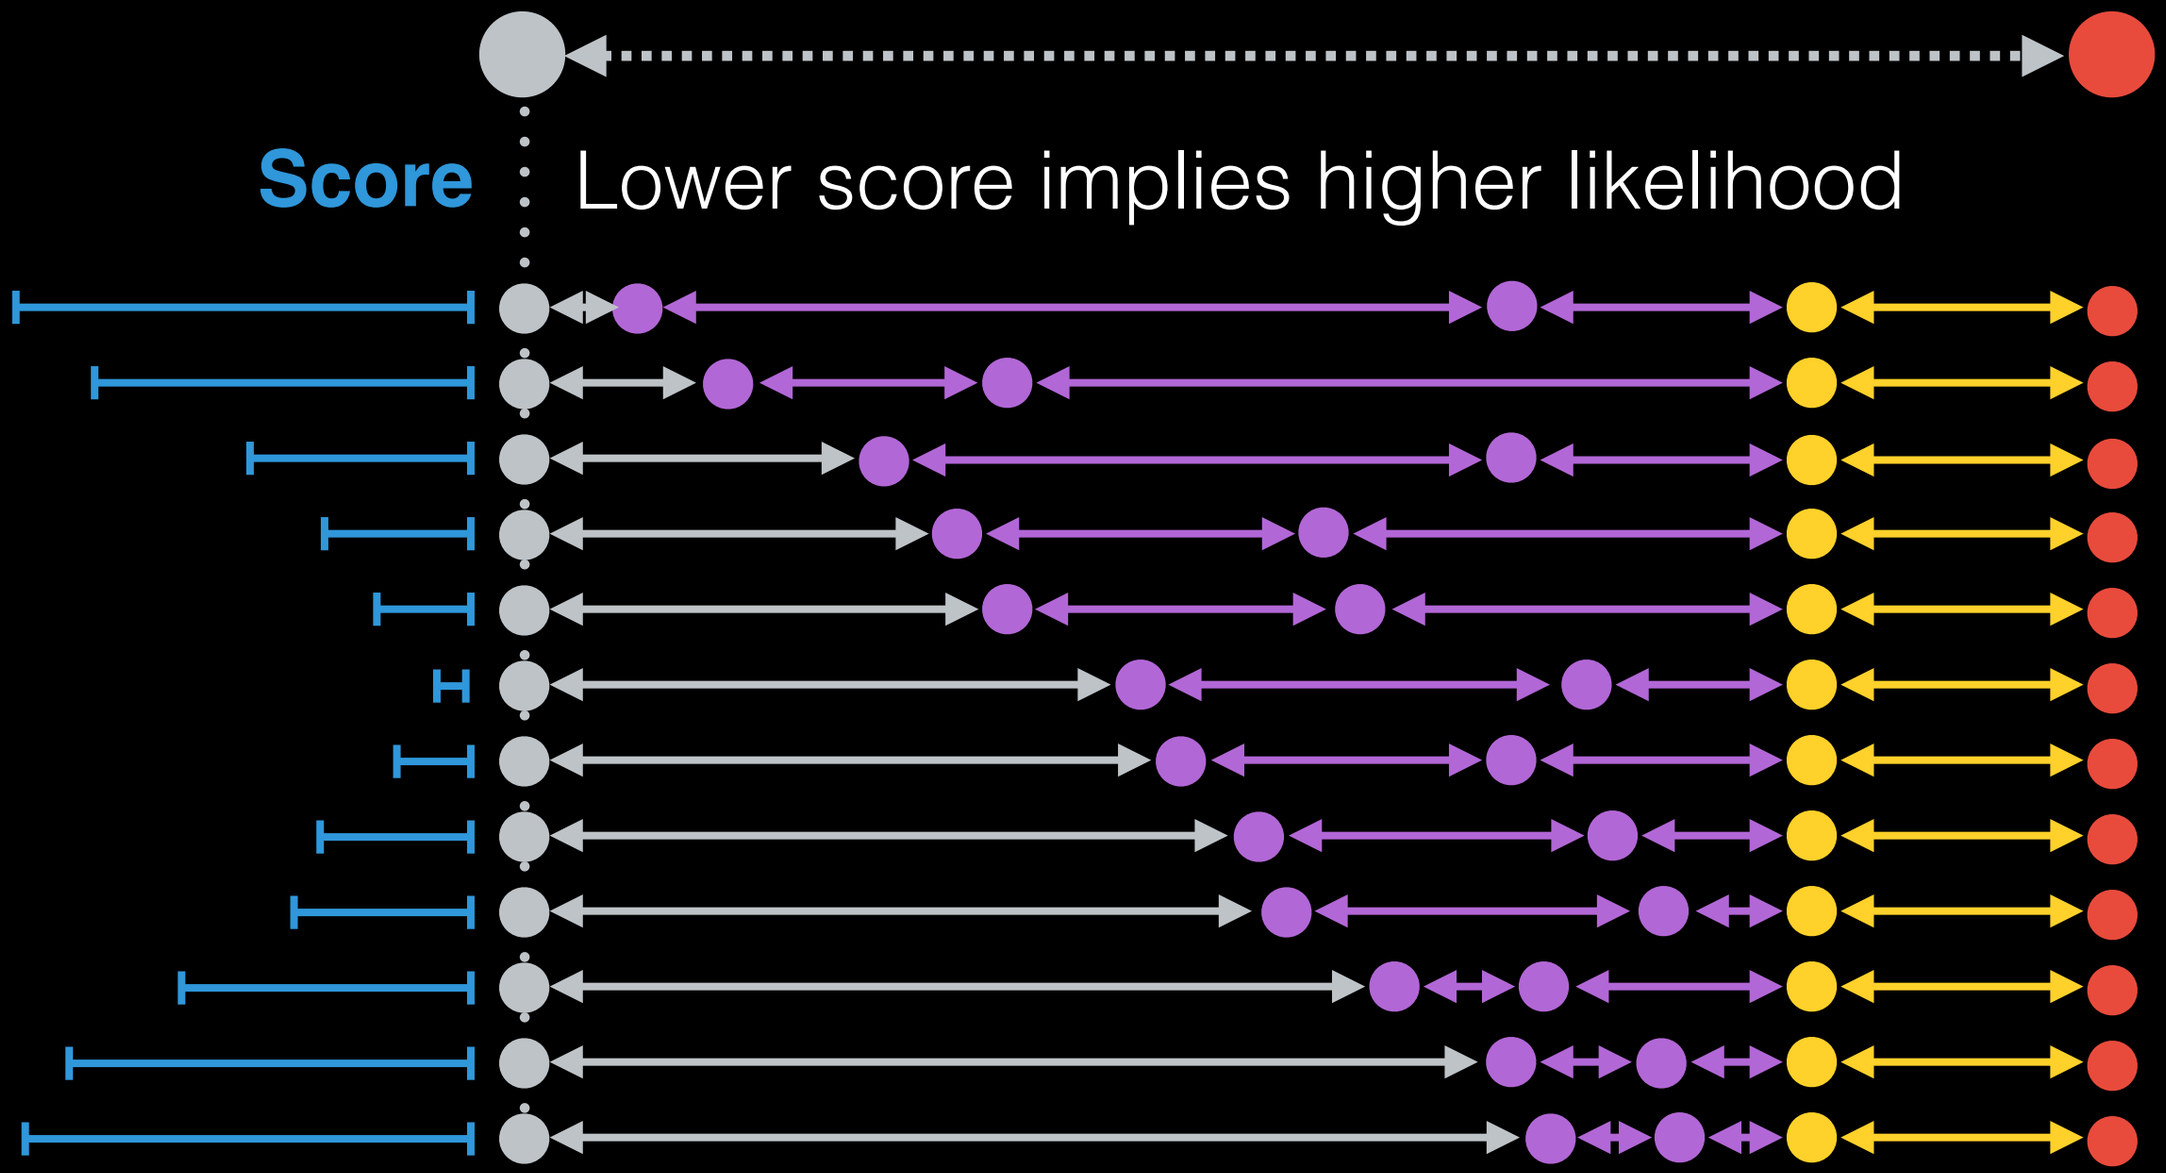

Informed target selection

Probe the more likely circuits first

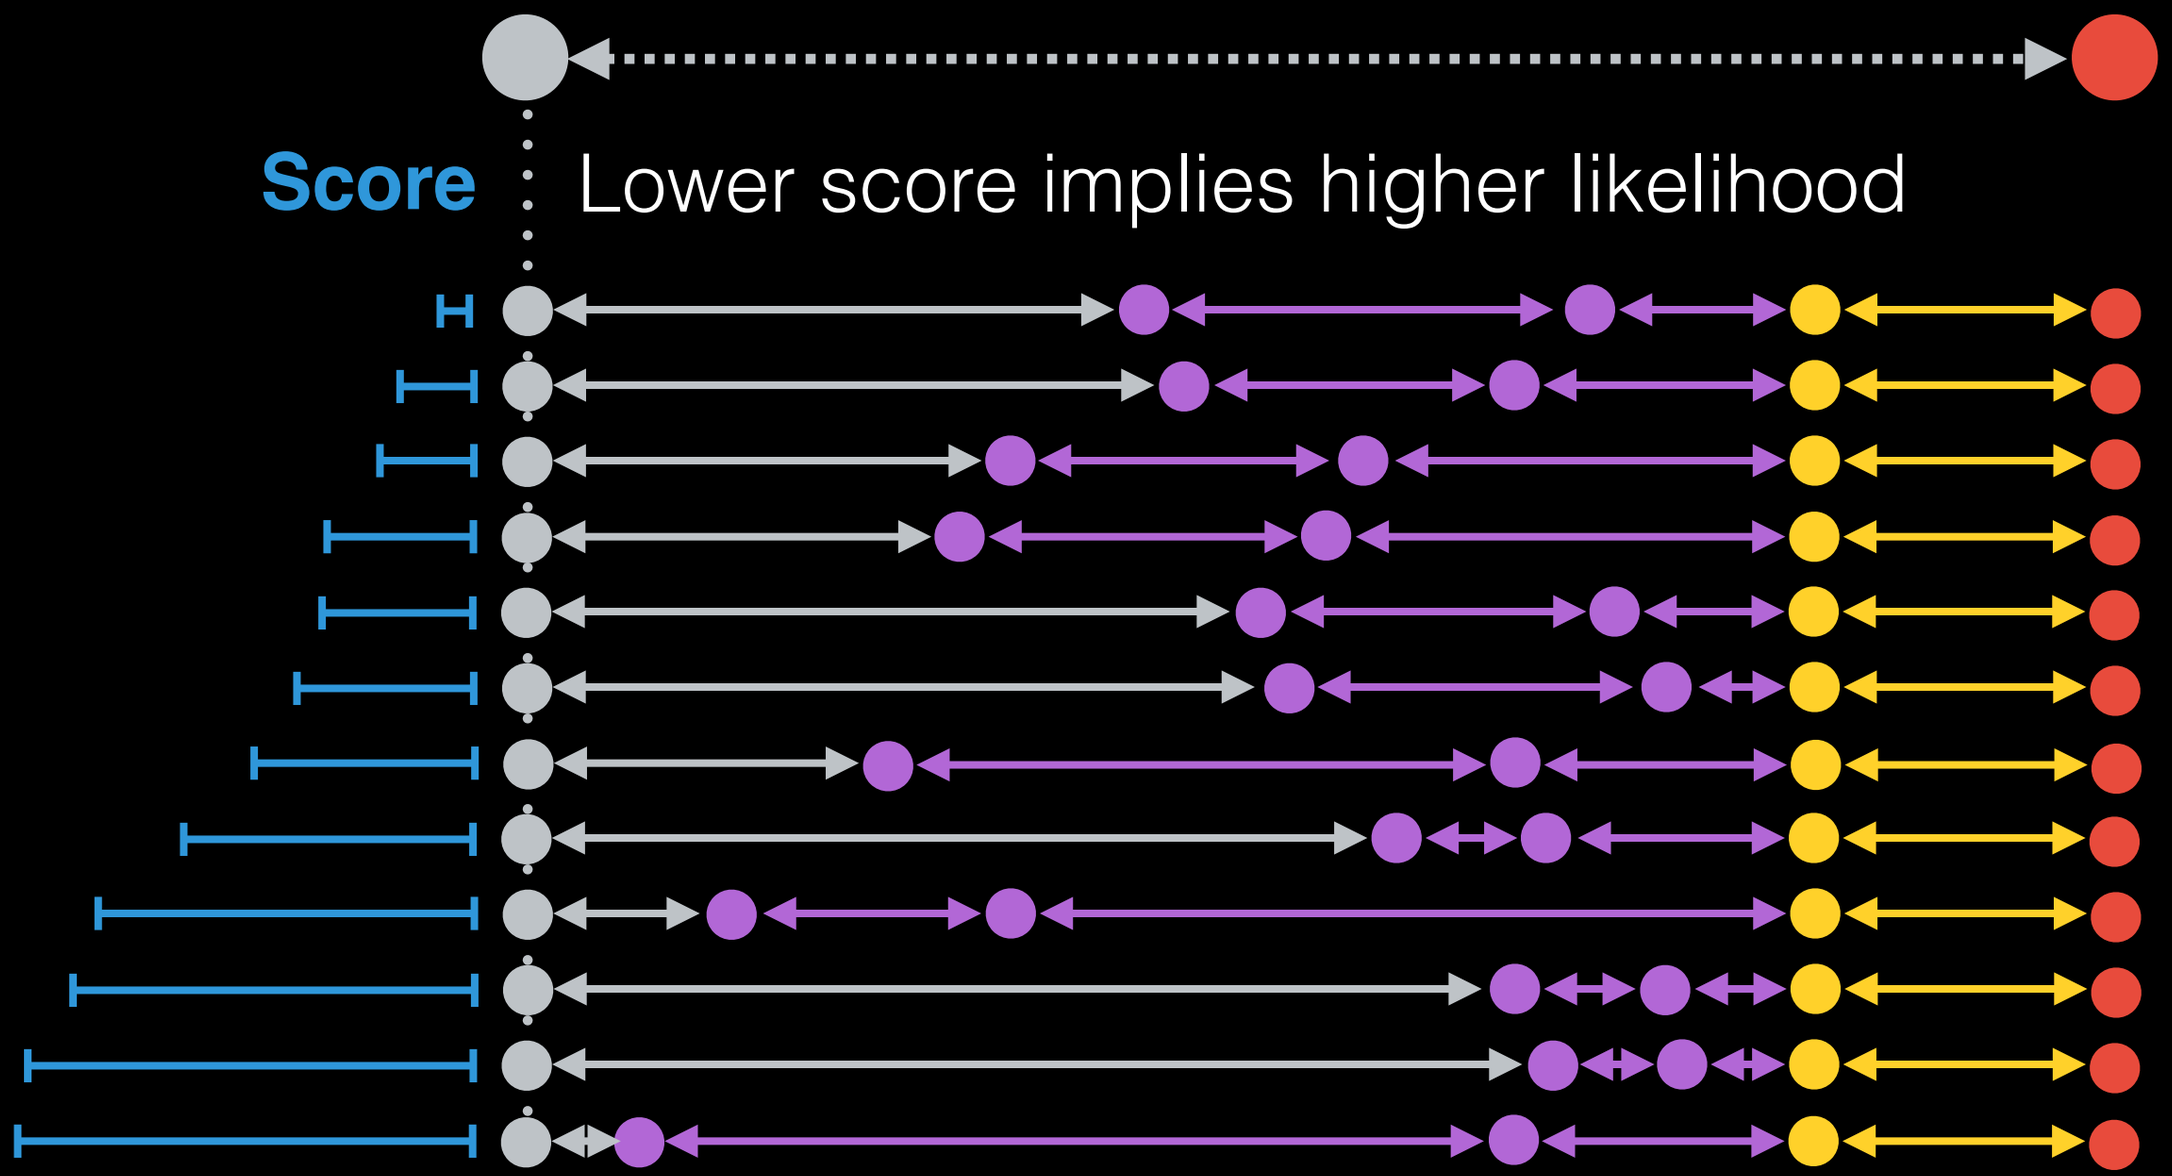

Faster deanonymization with Ting

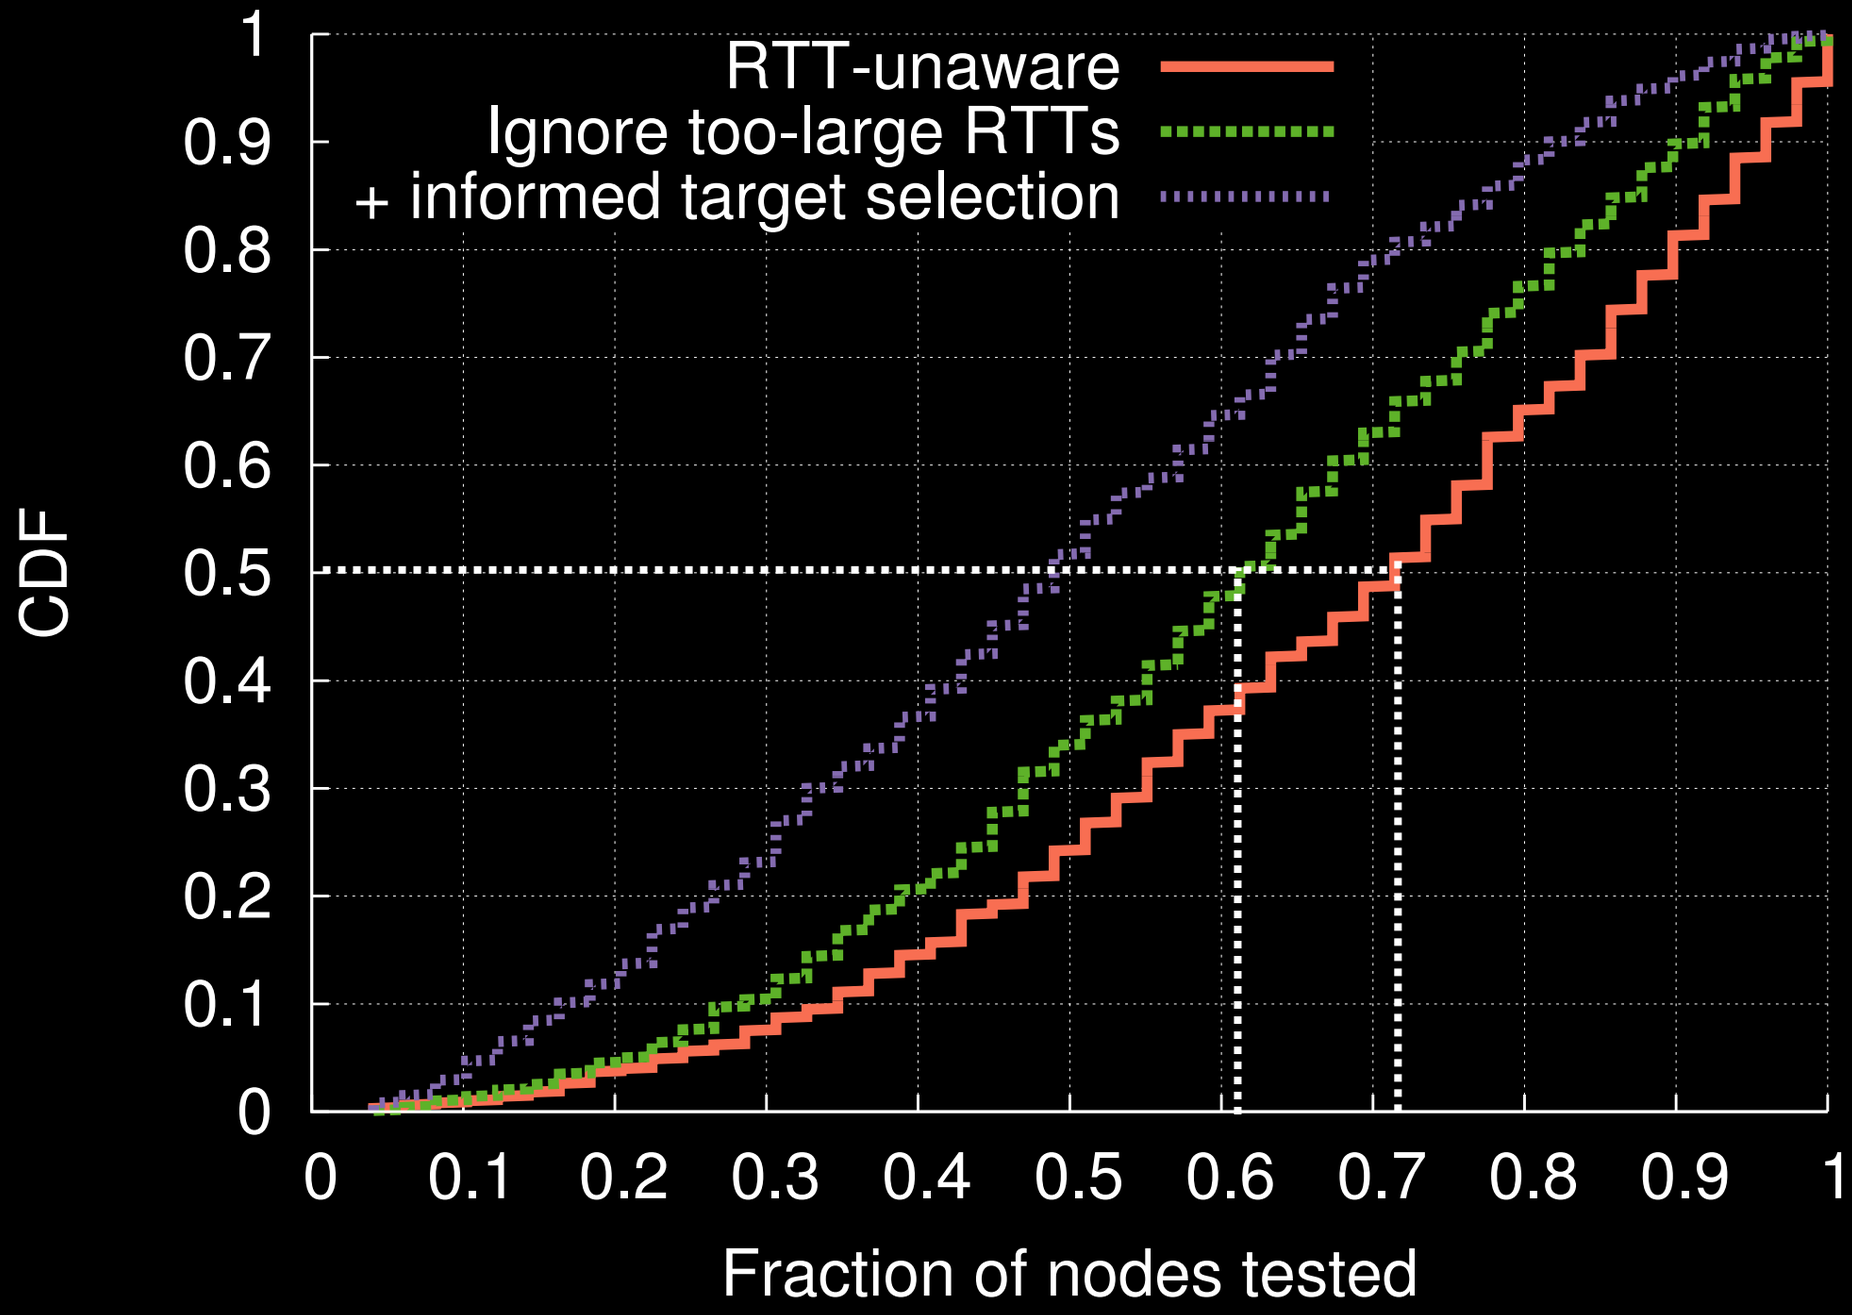

Informed target selection **decreases** search time by a median of **1.5×**

Faster deanonymization with Ting

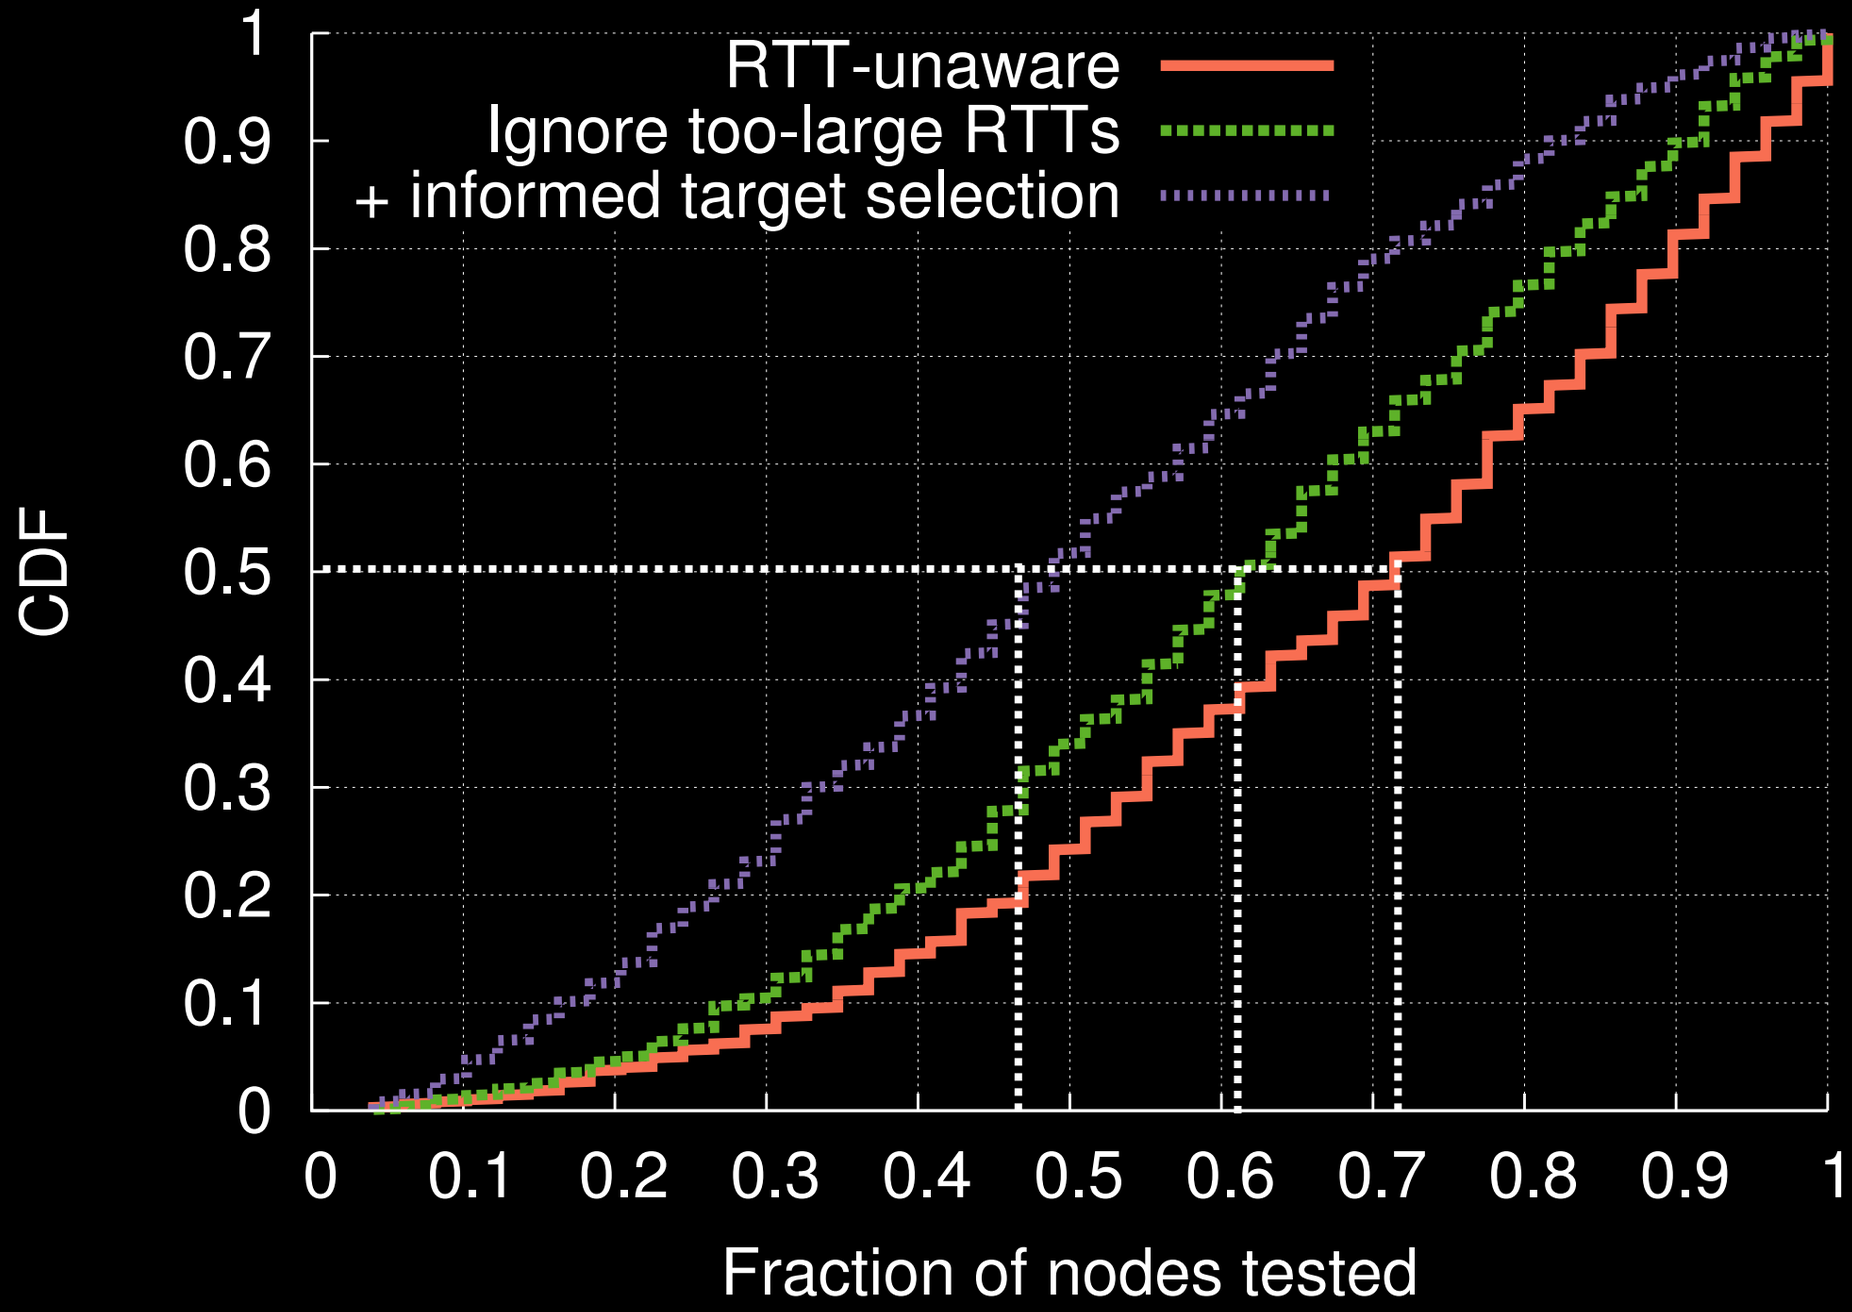

Informed target selection **decreases** search time by a median of **1.5×**

Faster deanonymization with Ting

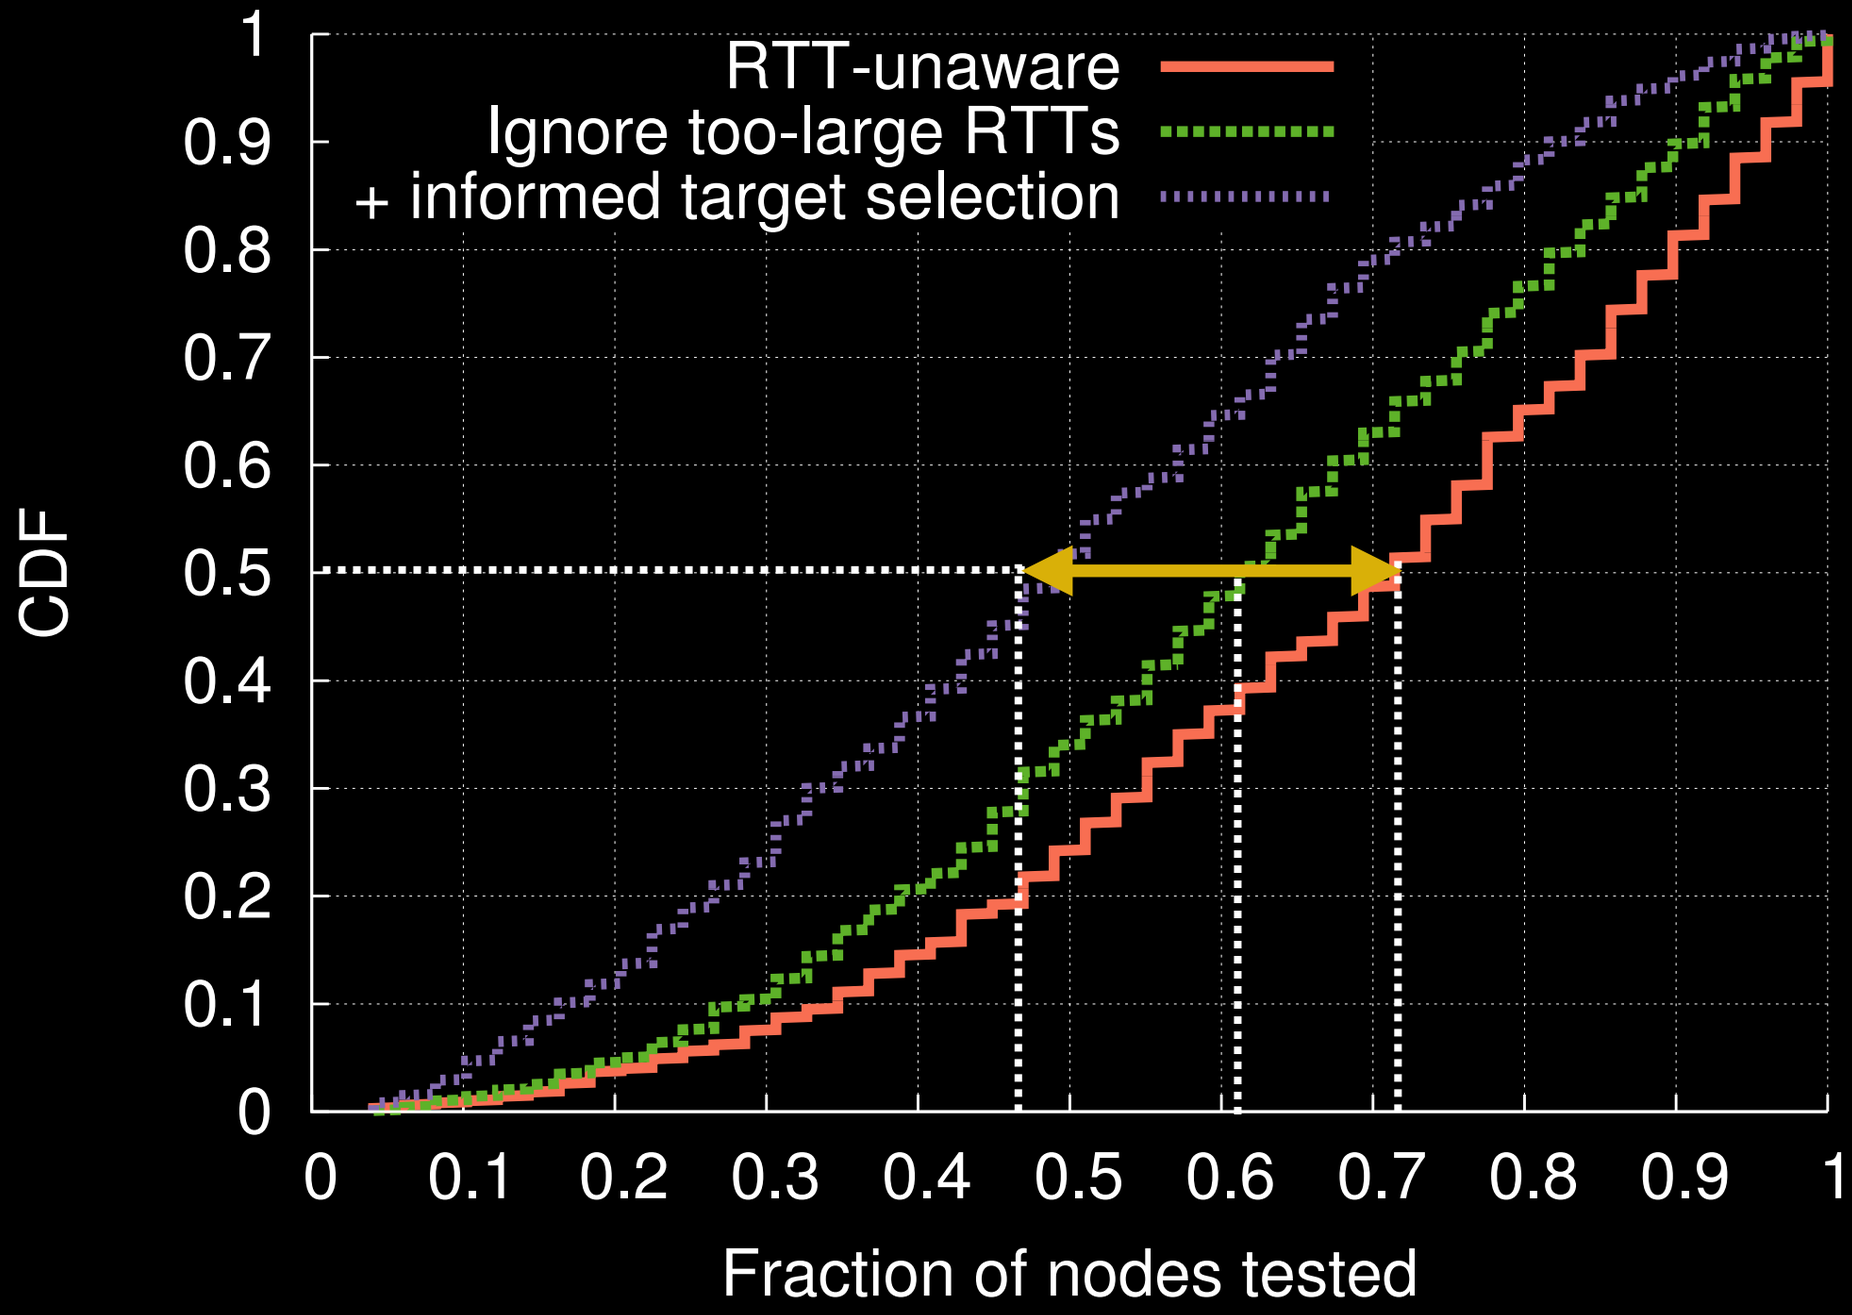

Informed target selection **decreases** search time by a median of **1.5×**

Summary

We lack a **practical** tool for measuring the latency between two arbitrary hosts

TING measures the latency between **Tor** nodes is fast, accurate, and practical

Source code and data available at:
www.cs.umd.edu/projects/ting

Implementation

Ting Client

Language: Python
Tor Controller: Stem
Tor-0.2.3.25-patched
SLOC: ~400

Test Relays

Tor-0.2.4.22 (latest)
PublishDescriptors 0
Restricted Exit Policy
Uptime: > 1 month