LABORATOIRE BORDELAIS DE RECHERCHE EN INFORMATIQUE

LaBRI

anr

université de BORDEAUX

Byzantine-Resilient Peer Sampling for Large-Scale Distributed Systems

Thesis defended on April 24, 2026, by **Augusta Mukam**

Before a committee consisting of :

Mme Sonia Ben Mokhtar	Research Director, CNRS	Reviewer
M. Yérom-David Bromberg	Professor, University of Rennes	Reviewer
Mme Patricia Thébault	Professor, University of Bordeaux	Examiner
M. Léo Mendiboure	Associate Professor, University of Pau and the Pays de l'Adour	Examiner
M. Laurent Réveillère	**Full Professor, University of Bordeaux**	**Thesis Advisor**
M. Joachim Bruneau-Queyreix	**Associate Professor, Bordeaux INP**	**Co-Advisor**

Large scale distributed systems

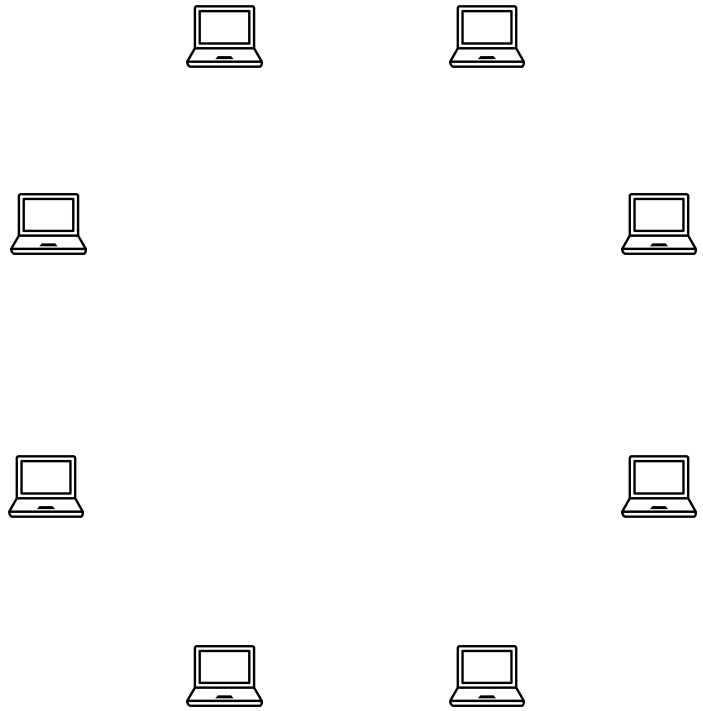

Large scale distributed systems

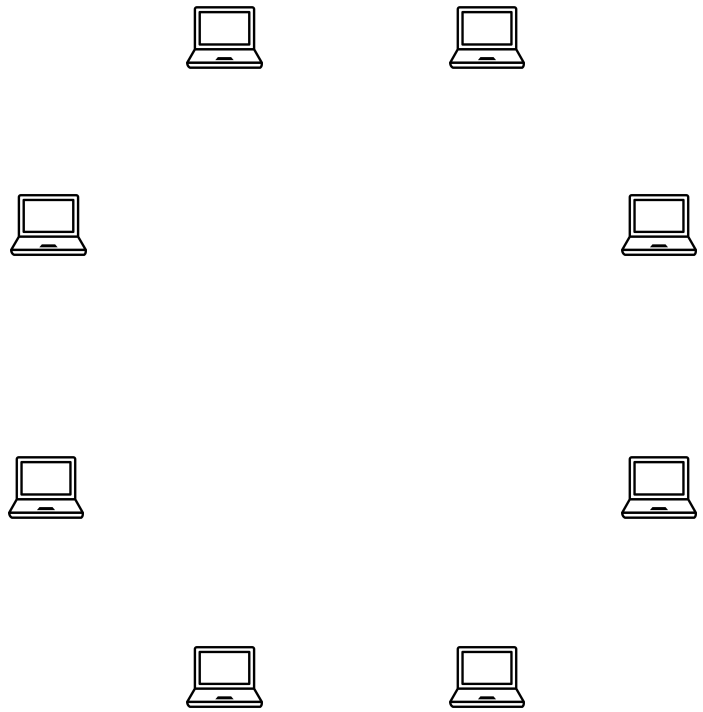

Characteristics

➔ large scale

Large scale distributed systems

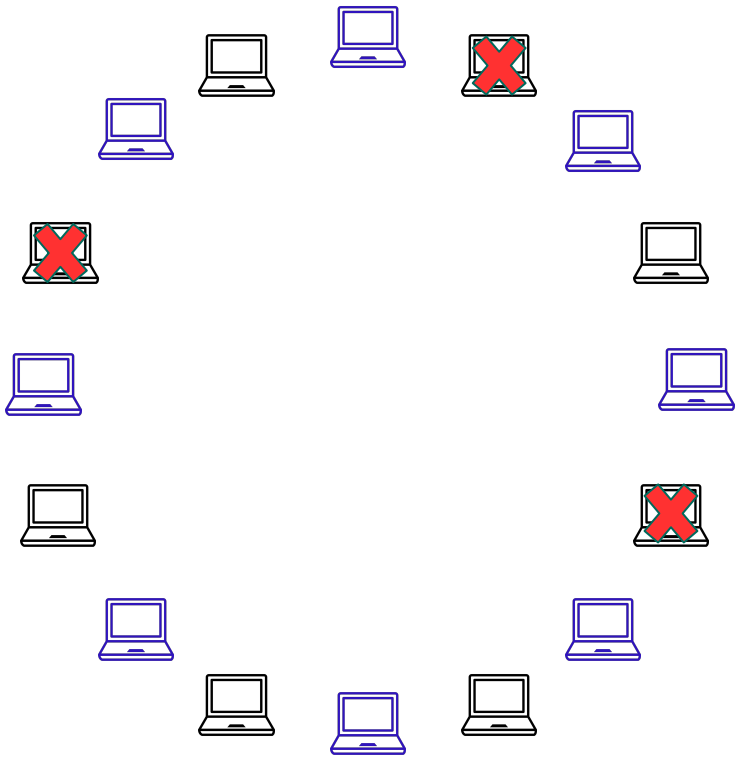

Characteristics

- ➔ large scale
- ➔ dynamic (churn, failures)

Large scale distributed systems

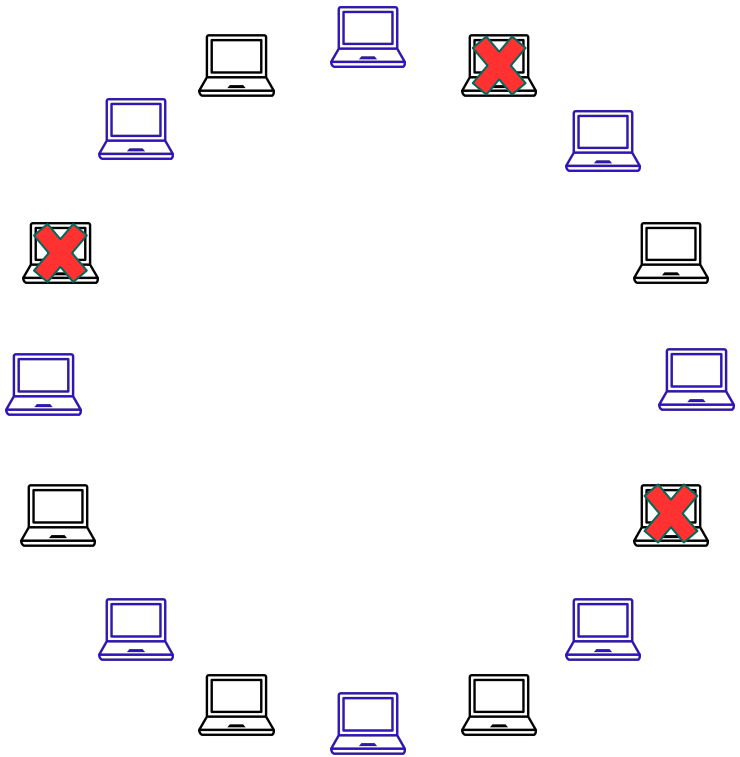

Characteristics

- ➔ large scale
- ➔ dynamic (churn, failures)
- ➔ decentralized

How node discover themselves?

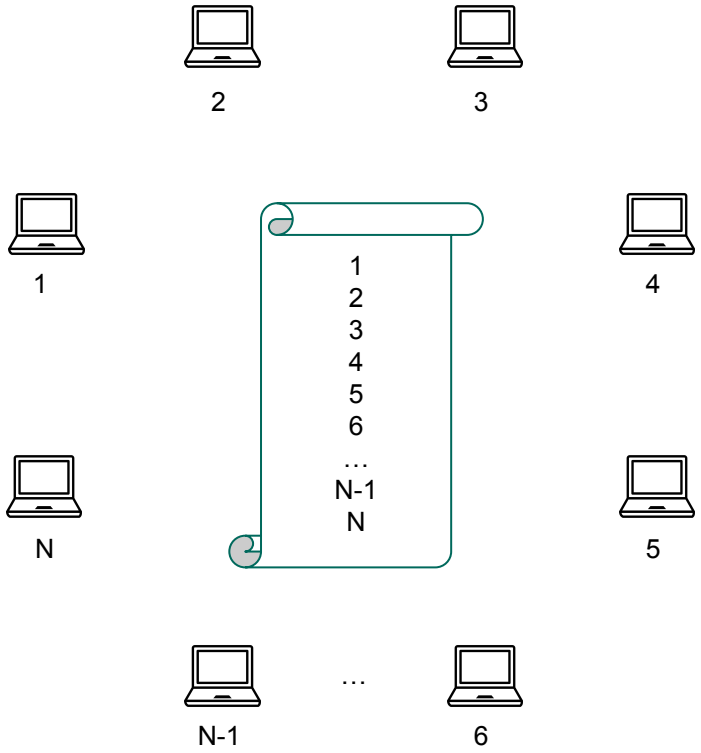

How node discover themselves?

How does a node choose its neighbors?

Jelasity *et al.* (2004)

The peer sampling service

Objectif

➔ uniform random samples

The peer sampling service

Objectif

➔ uniform random samples

Applications

➔ information dissemination
➔ overlay topology

The peer sampling service

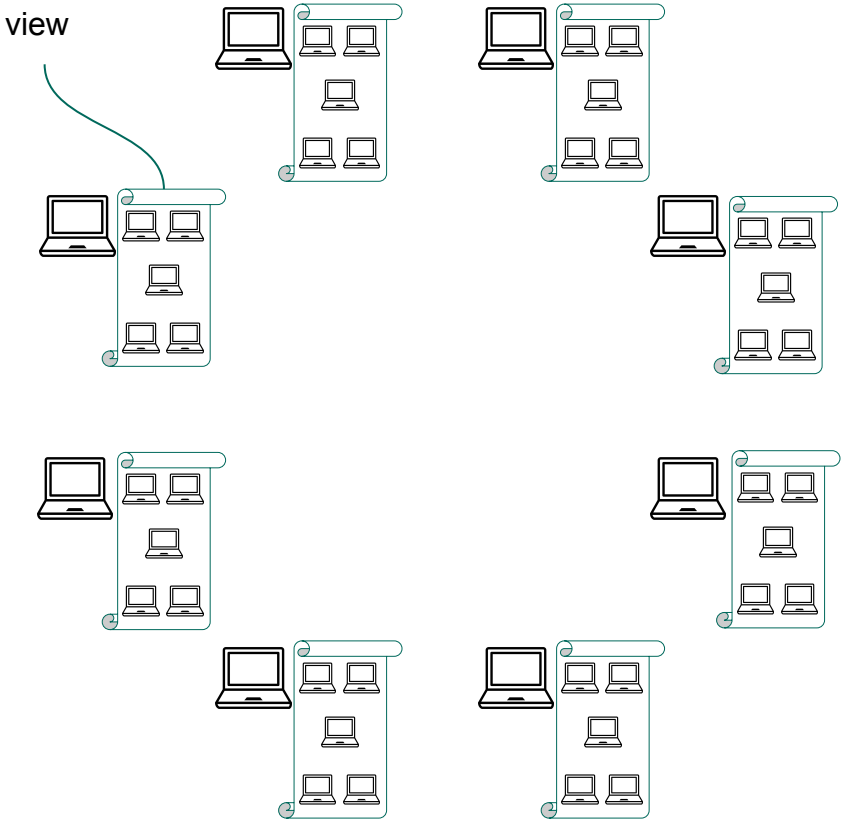

The peer sampling service

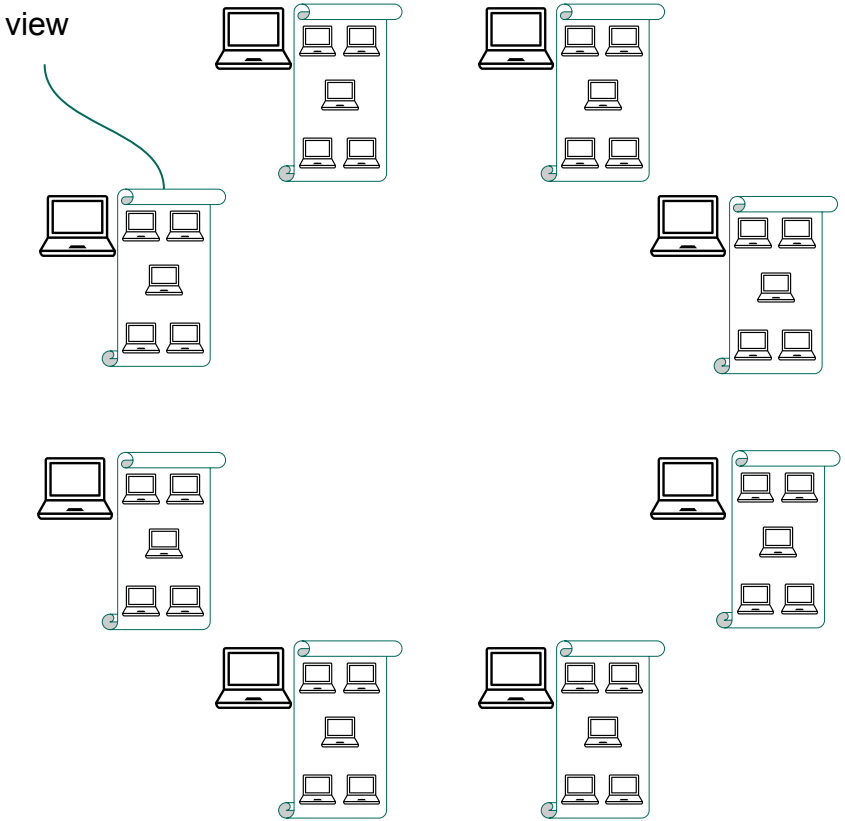

Protocol

➔ gossip-based propagation

The peer sampling service

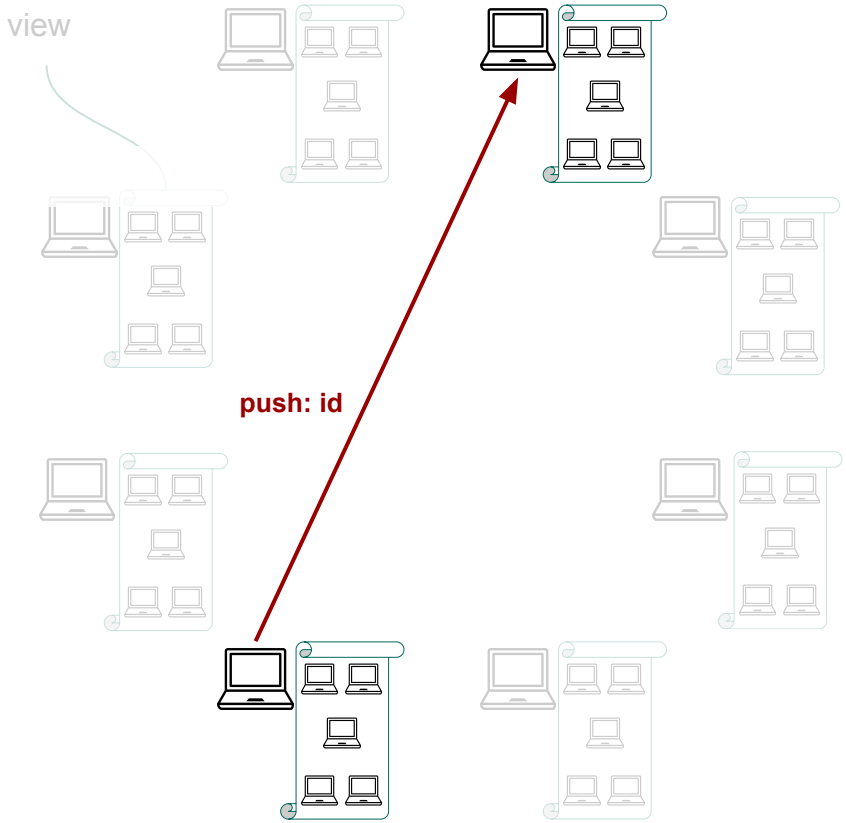

Protocol

➔ gossip-based propagation

push

The peer sampling service

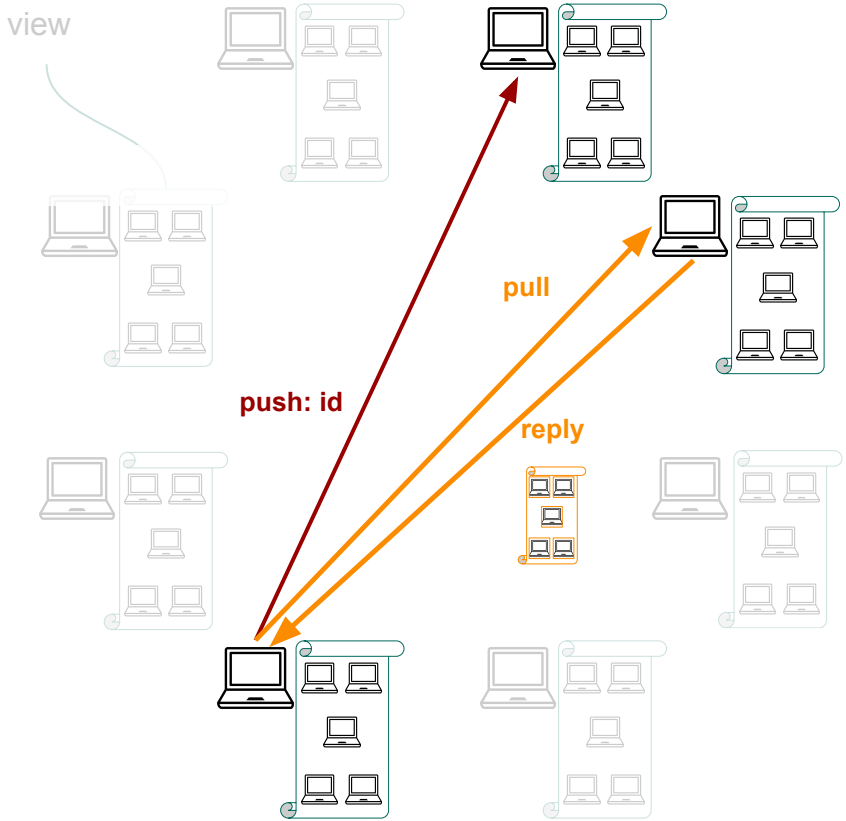

Protocol

➔ gossip-based propagation

push

pull

The peer sampling service

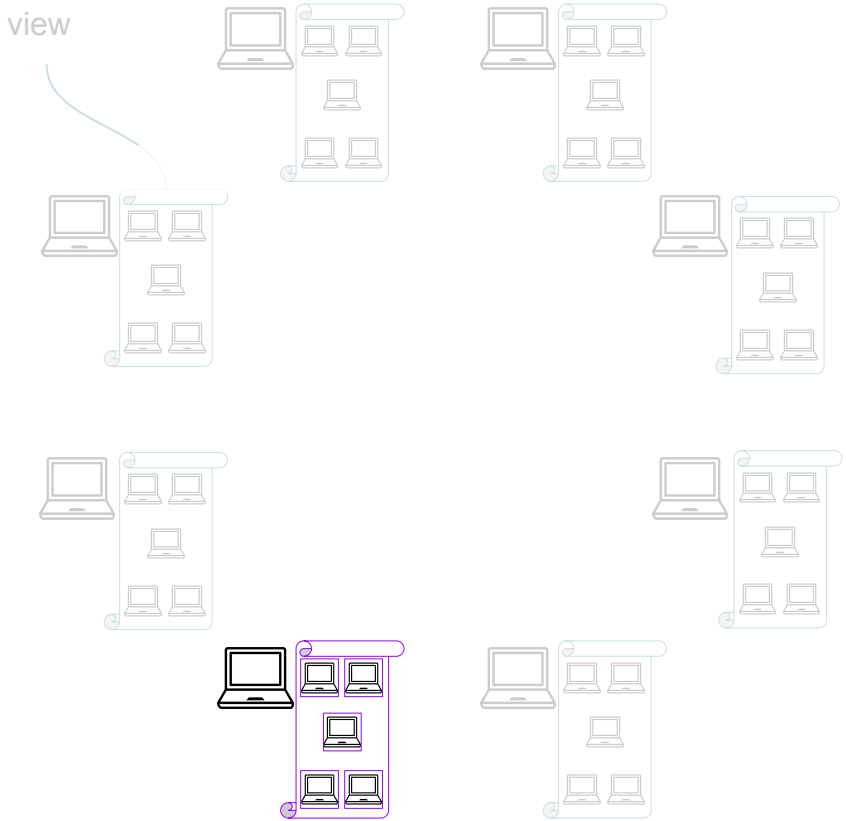

Protocol

- ➔ gossip-based propagation
- ➔ update view

push

pull

Under attack

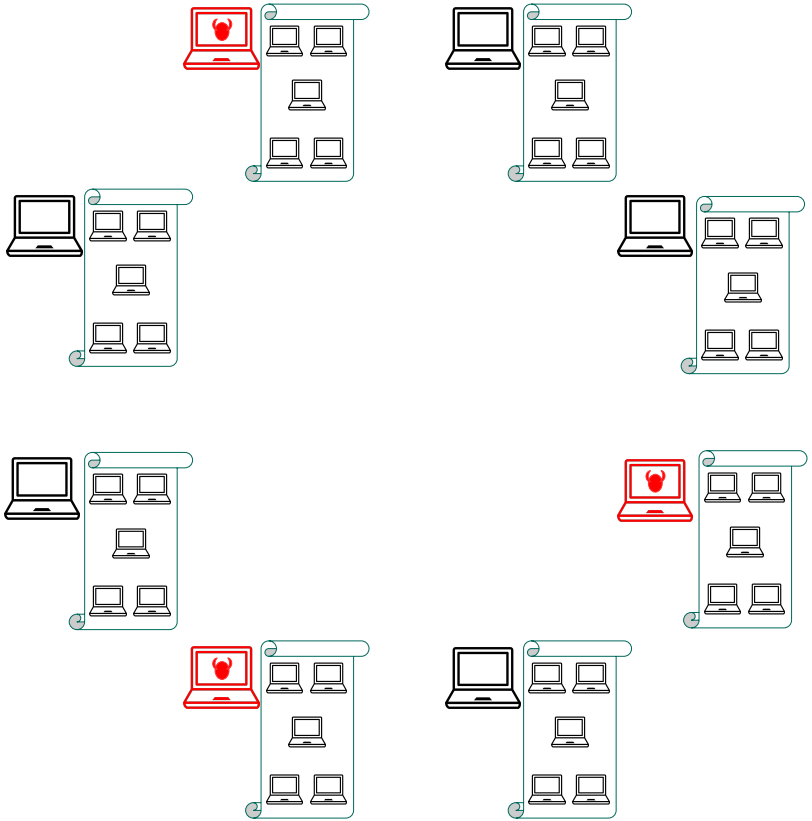

Under attack

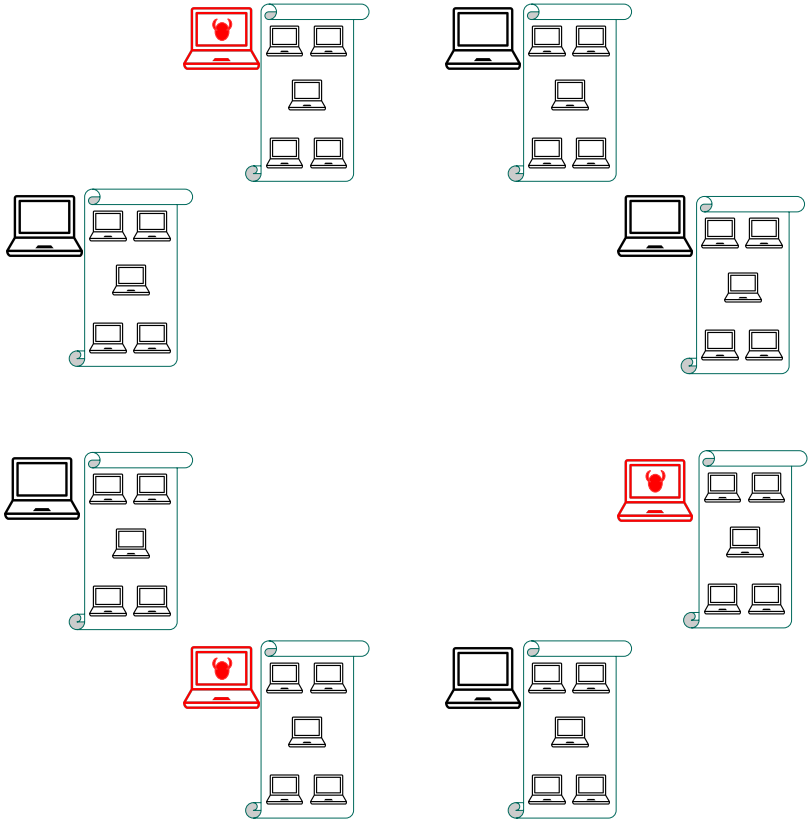

Objectif

➔ bias sampling process

Under attack

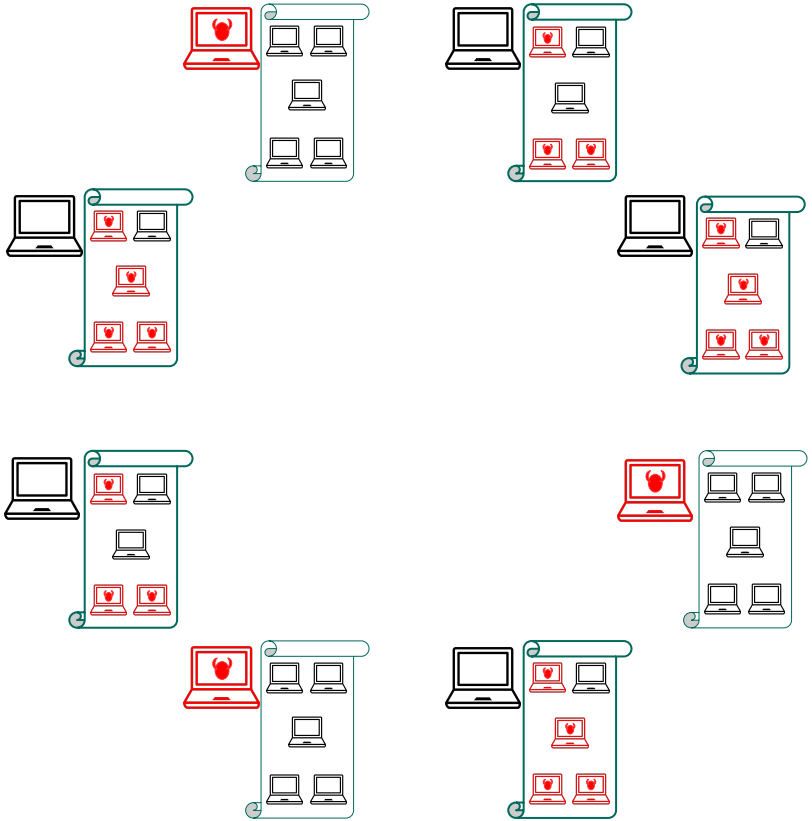

Objectif

- ➔ bias sampling process
- ➔ being overrepresented

Under attack

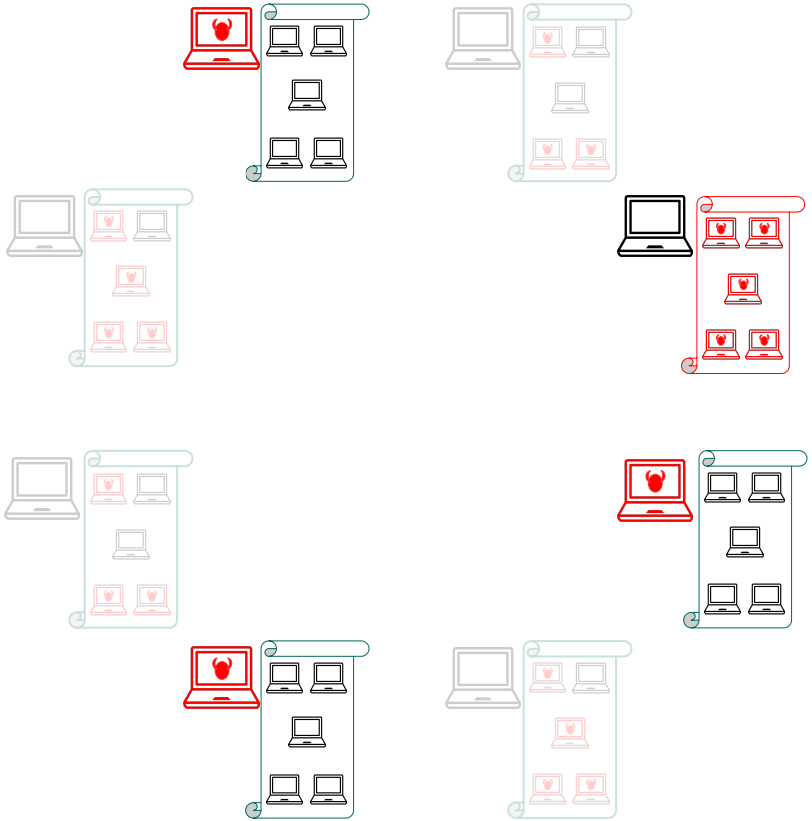

Objectif

- ➔ bias sampling process
- ➔ being overrepresented

Consequence

- ➔ Node eclipse

Under attack

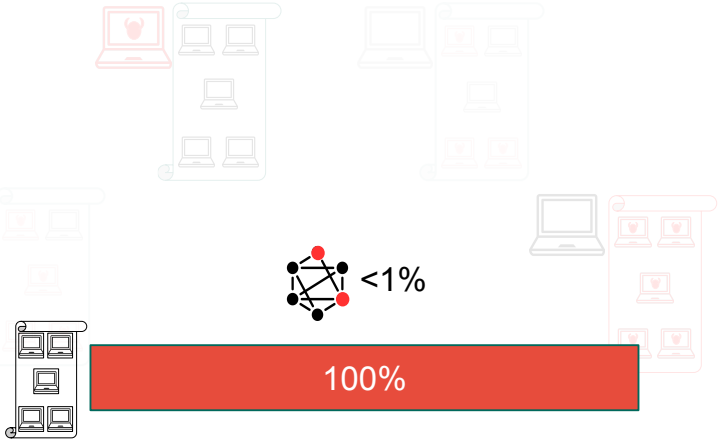

Objectif

- ➔ bias sampling process
- ➔ being overrepresented

Consequence

- ➔ Node eclipse

Under attack

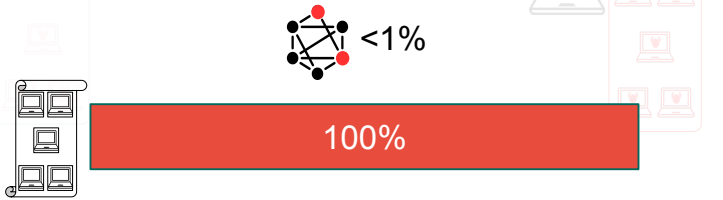

Objectif

- ➔ bias sampling process
- ➔ being overrepresented

Consequence

- ➔ Node eclipse
- ➔ Target higher level protocols

Heilman *et al.* (2015)

Brahms overview

Gossip component

Bortnikov *et al.* (2015)

- ➔ push & pull requests
- ➔ update view

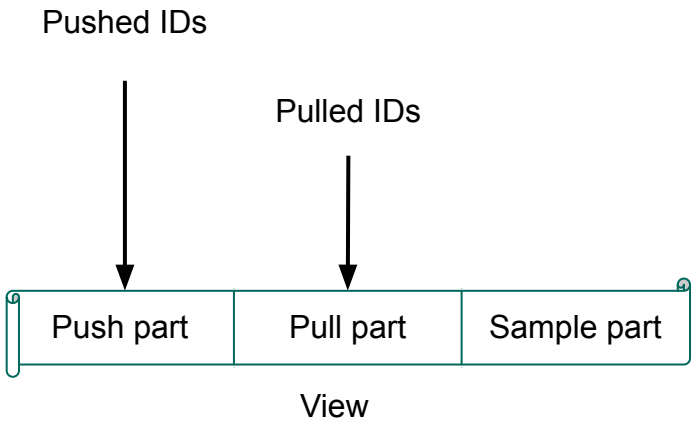

Brahms overview

Gossip component

- ➔ push & pull requests
- ➔ update view

Sampling component

- ➔ uniform sample of nodes

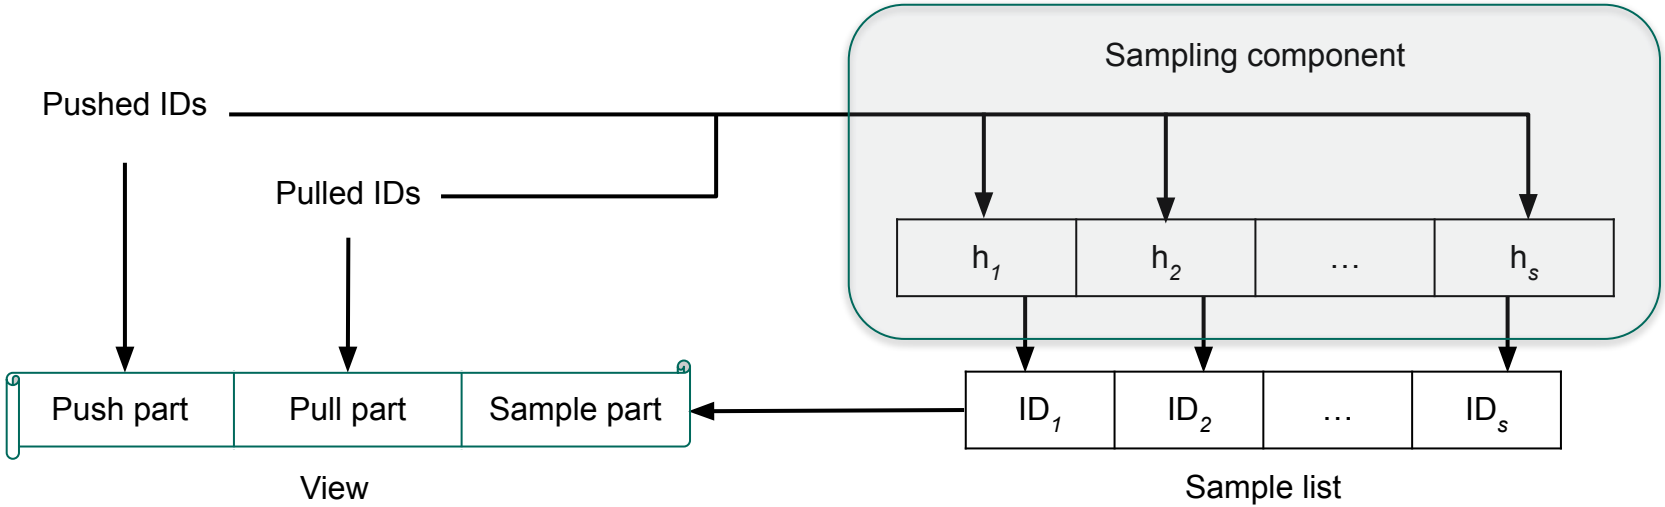

Under balanced attack

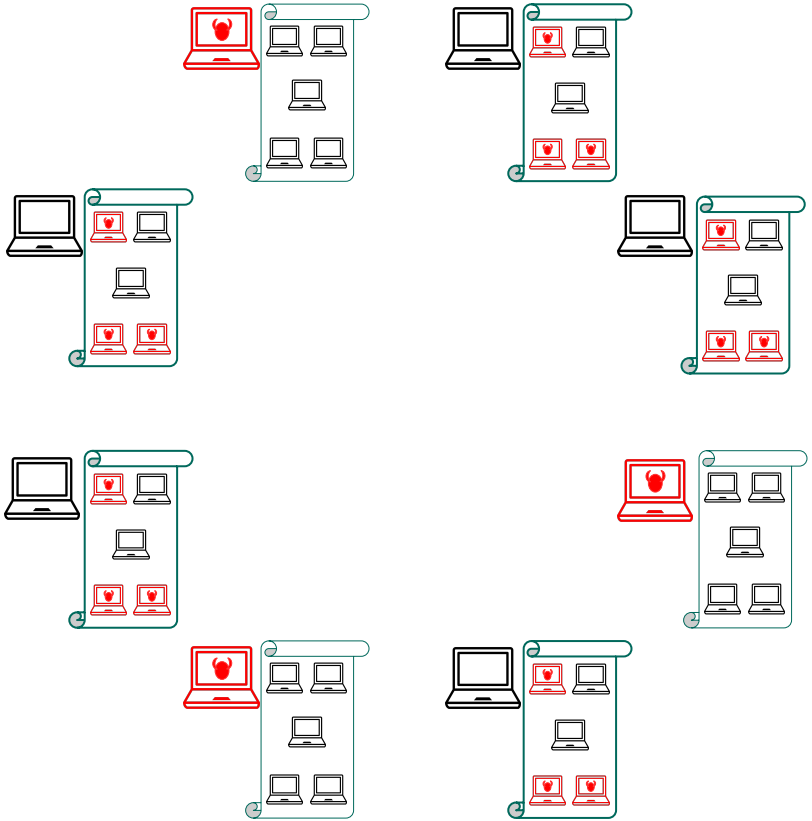

Under balanced attack

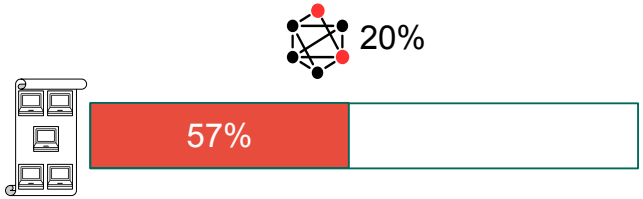

Byzantine-resilient peer sampling for large-scale distributed systems

1. Improve tolerance to Byzantine attacks

Built on top of Brahms
Occurrence counting
Collaborative counting

Mukam *et al.* (2024)

2. Count estimation for scalability

Comparative study of probabilistic datastructures
New performance metric
Bitmatcher

3. Unbounded stream count estimation

Decay and merge for probabilistic datastructures
Integration in Aupe

Mukam *et al.* (in preparation)

AUPE
Collaborative Byzantine fault-tolerant peer-sampling

Aupe protocol

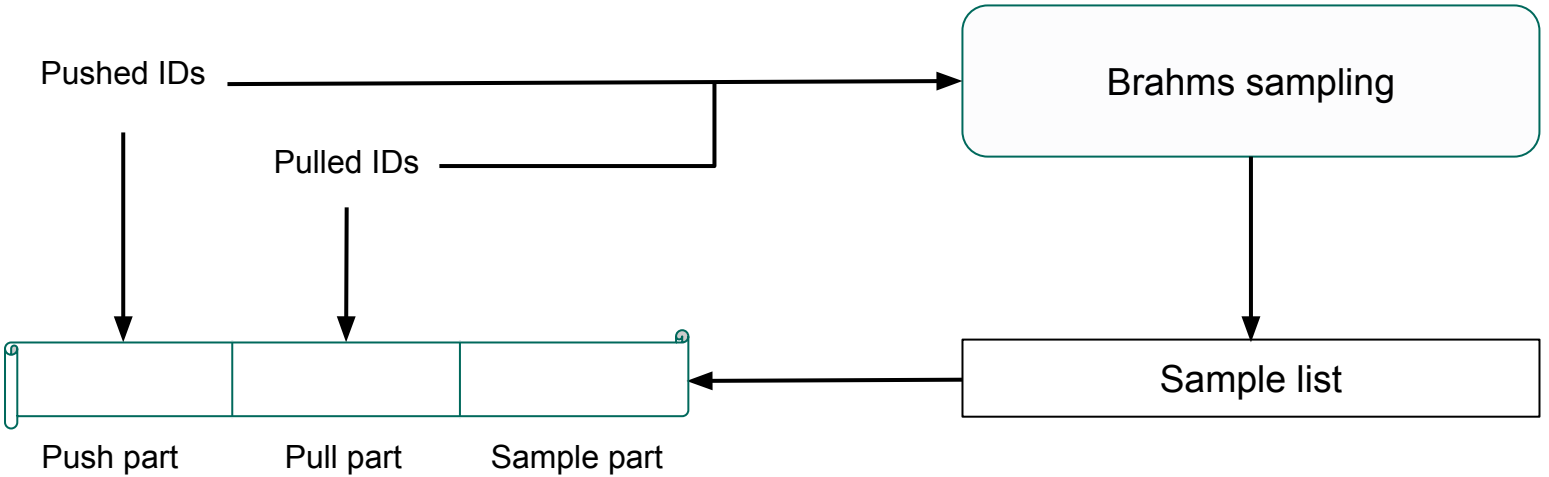

Aupe protocol

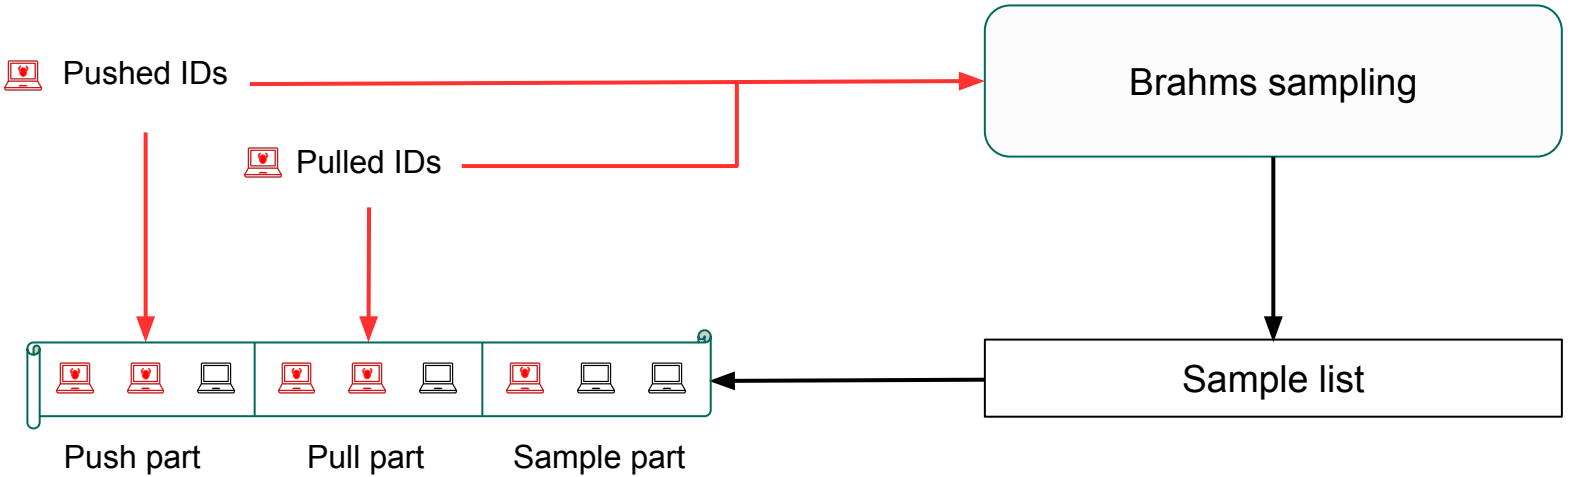

Aupe protocol

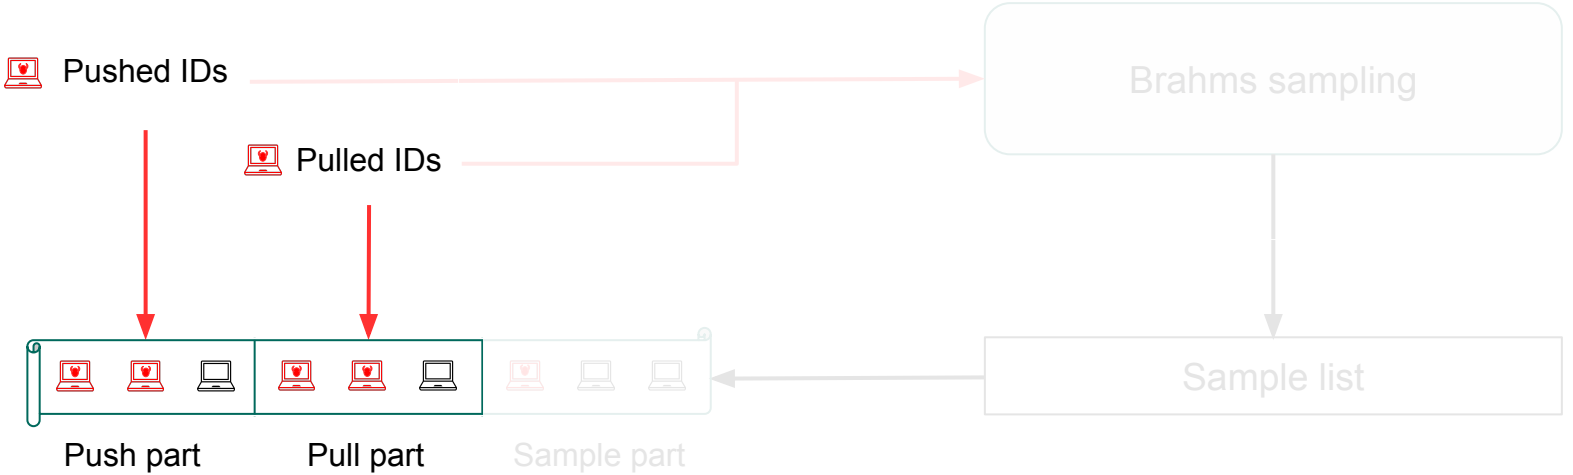

Aupe protocol

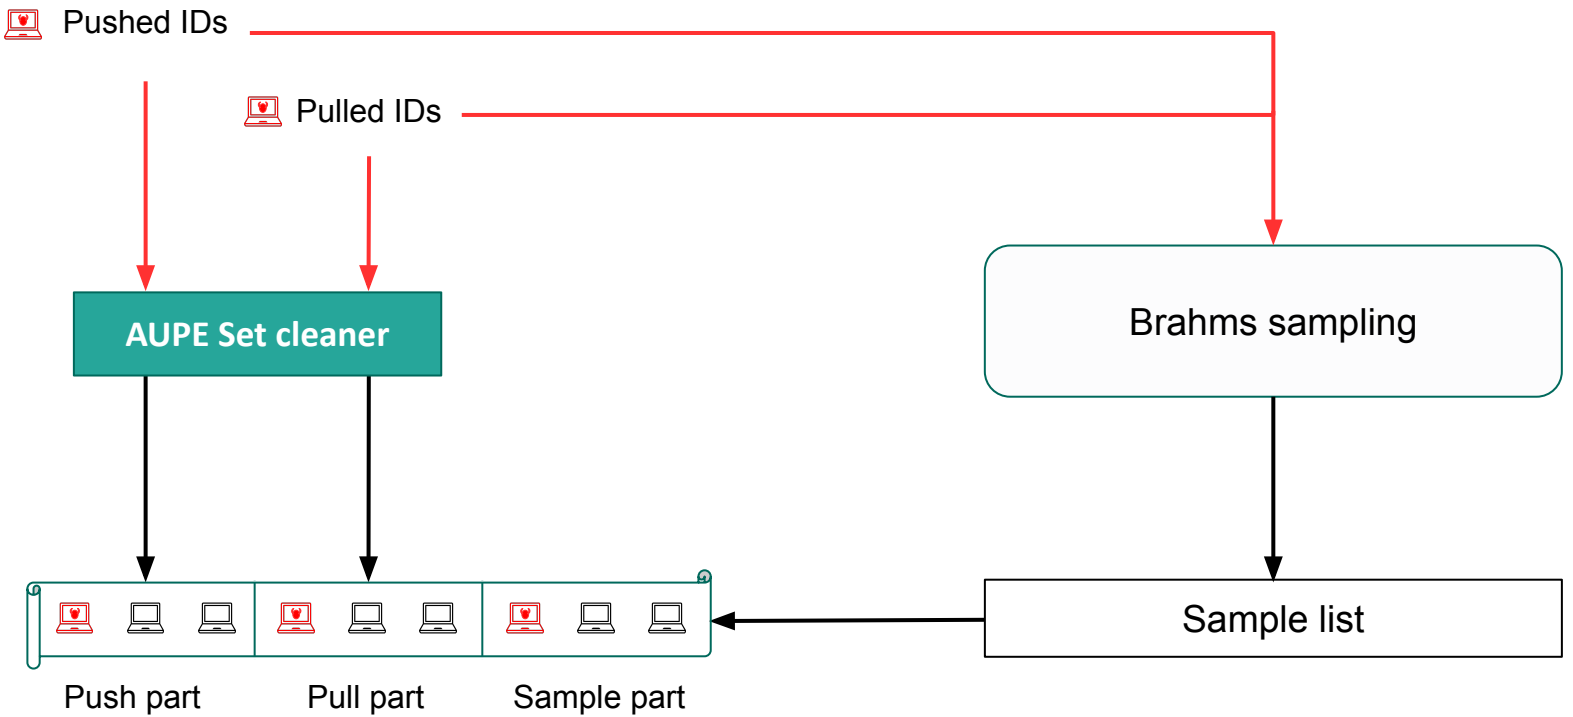

Debiasing procedure

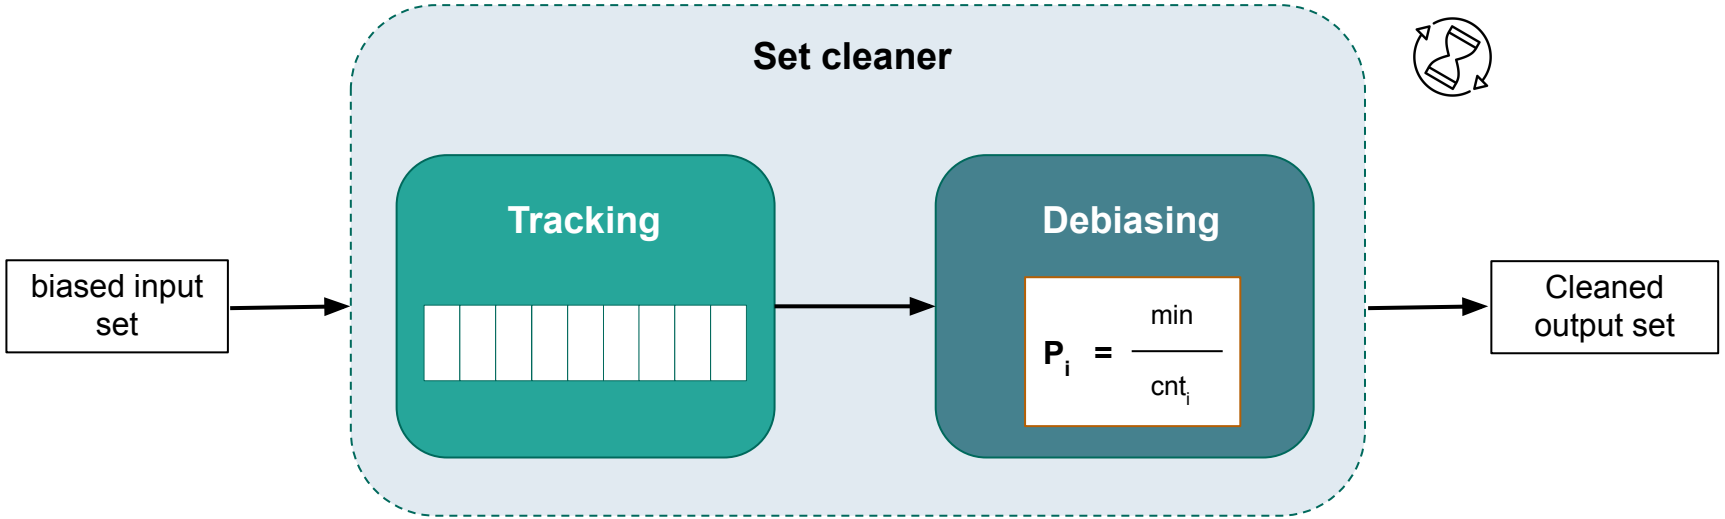

Anceaume *et al.* (2013)

Collaborative counting

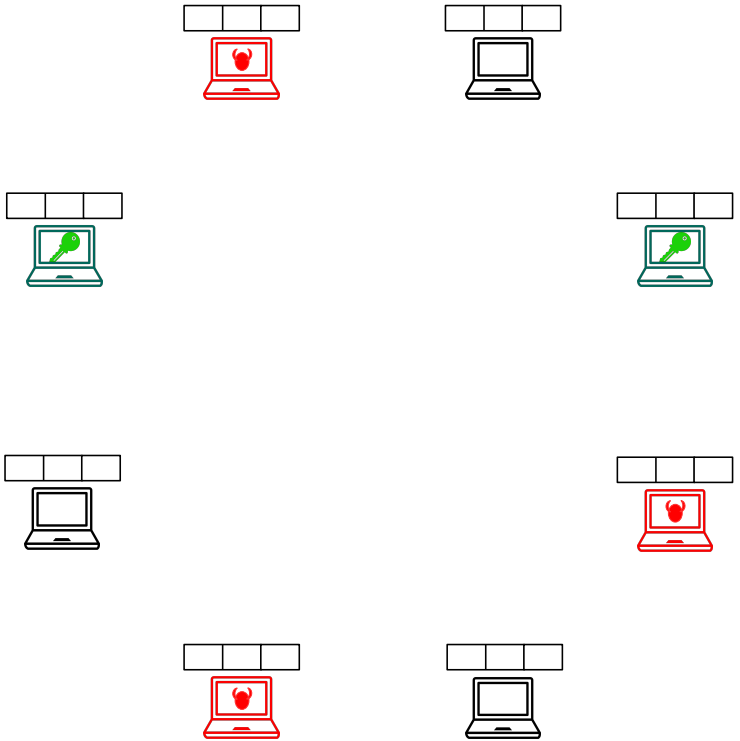

Trusted nodes

➔ mutual authentication

Collaborative counting

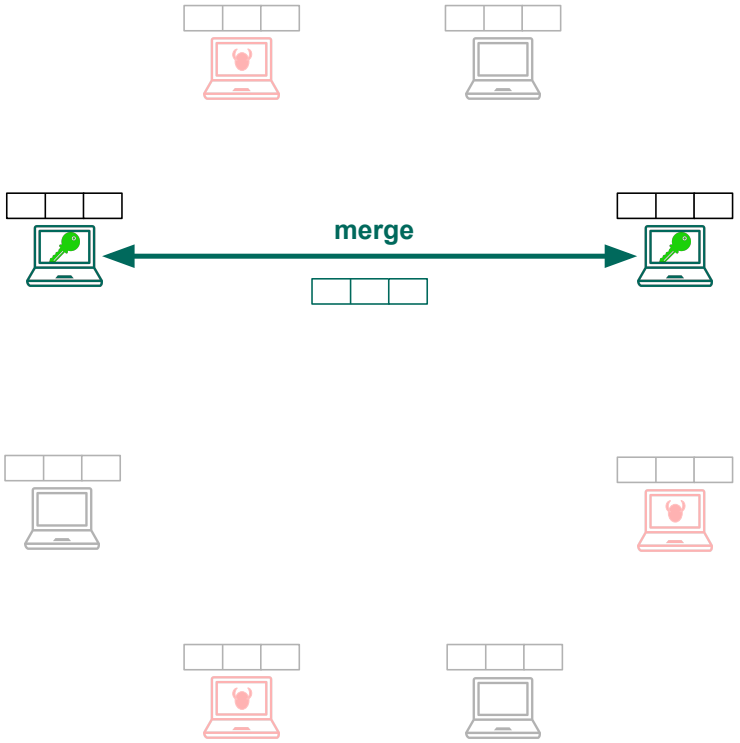

Trusted nodes

- ➔ mutual authentication
- ➔ tracking component merge

Fault tolerance evaluation

N=10K f=8-50%

AUPE Set cleaner + Collaborative debiasing

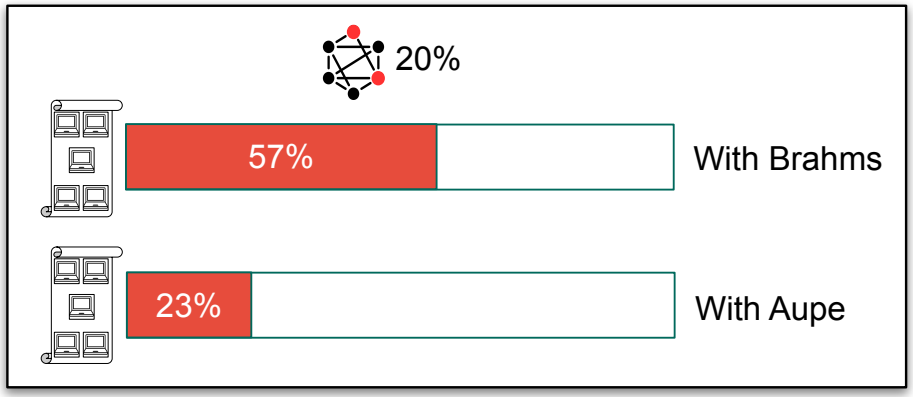

Fault tolerance evaluation

N=10K f=8-50%

Bortnikov *et al.* (2015)
Auvolat *et al.* (2023)

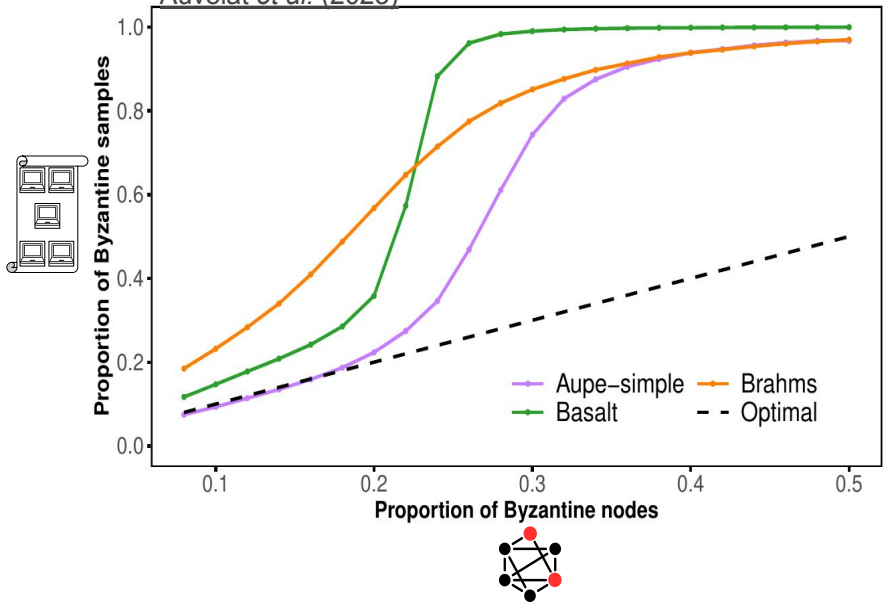

Bortnikov *et al.* (2015)
Auvolat *et al.* (2023)

Collaborative work evaluation

N=10K f=8-50% t=10-30%

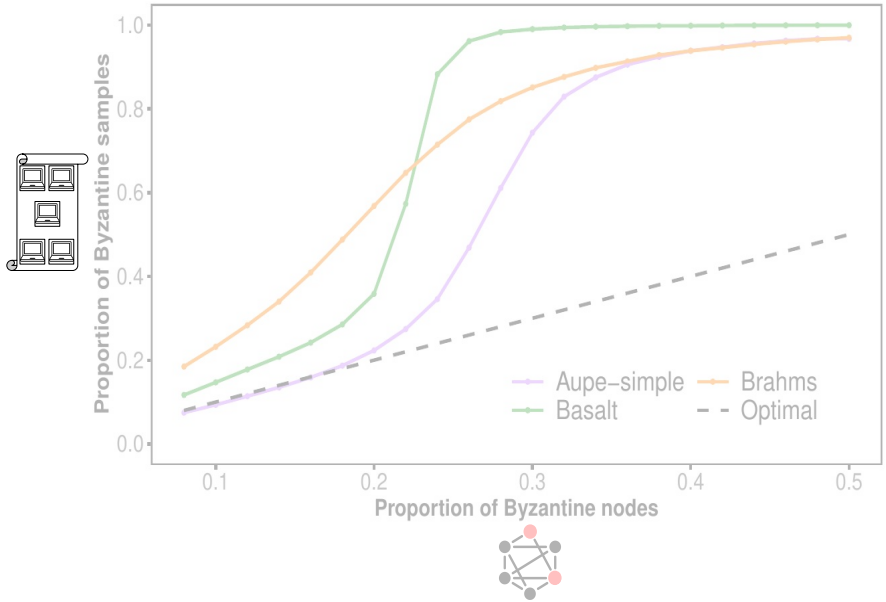

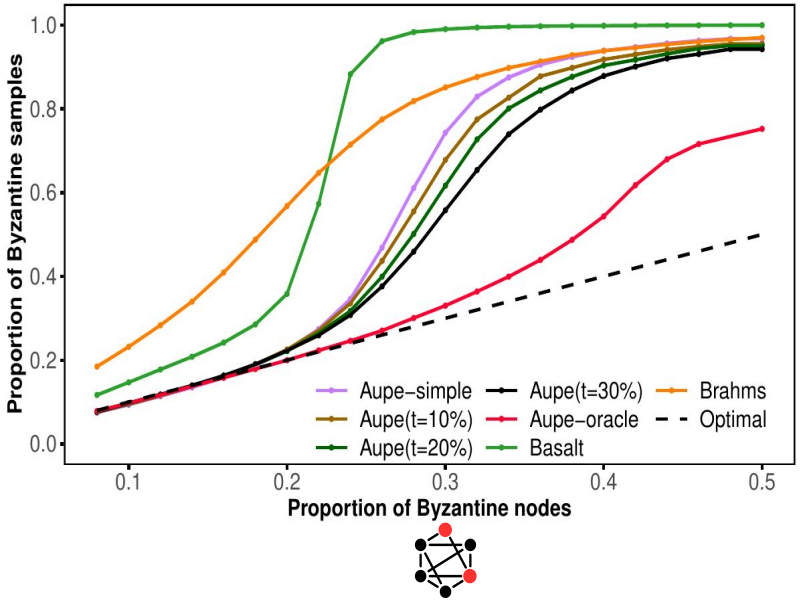

Collaborative work evaluation

N=10K f=8-50% t=10-30%

➔ **60%** better than Brahms
➔ **65%** better than Basalt

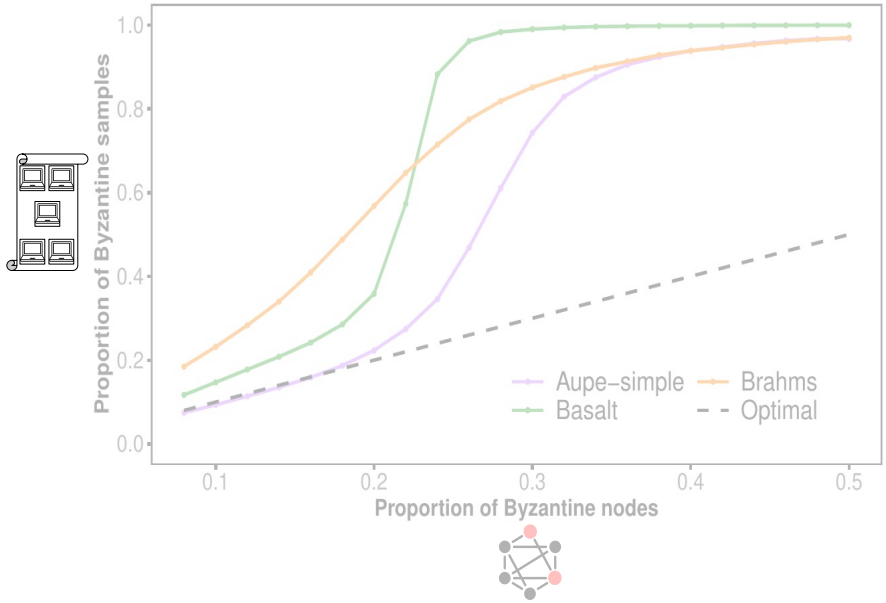

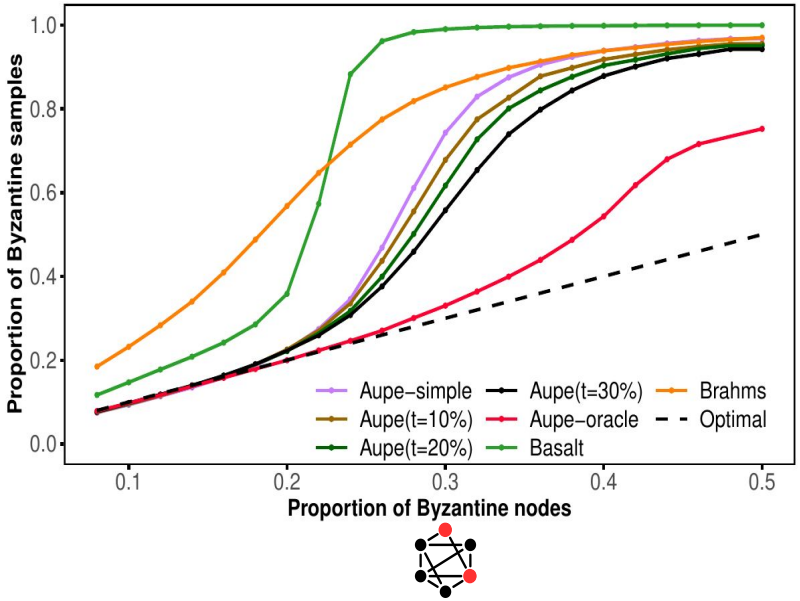

Robust frequency estimation for peer sampling in adversarial context

Sketches overview

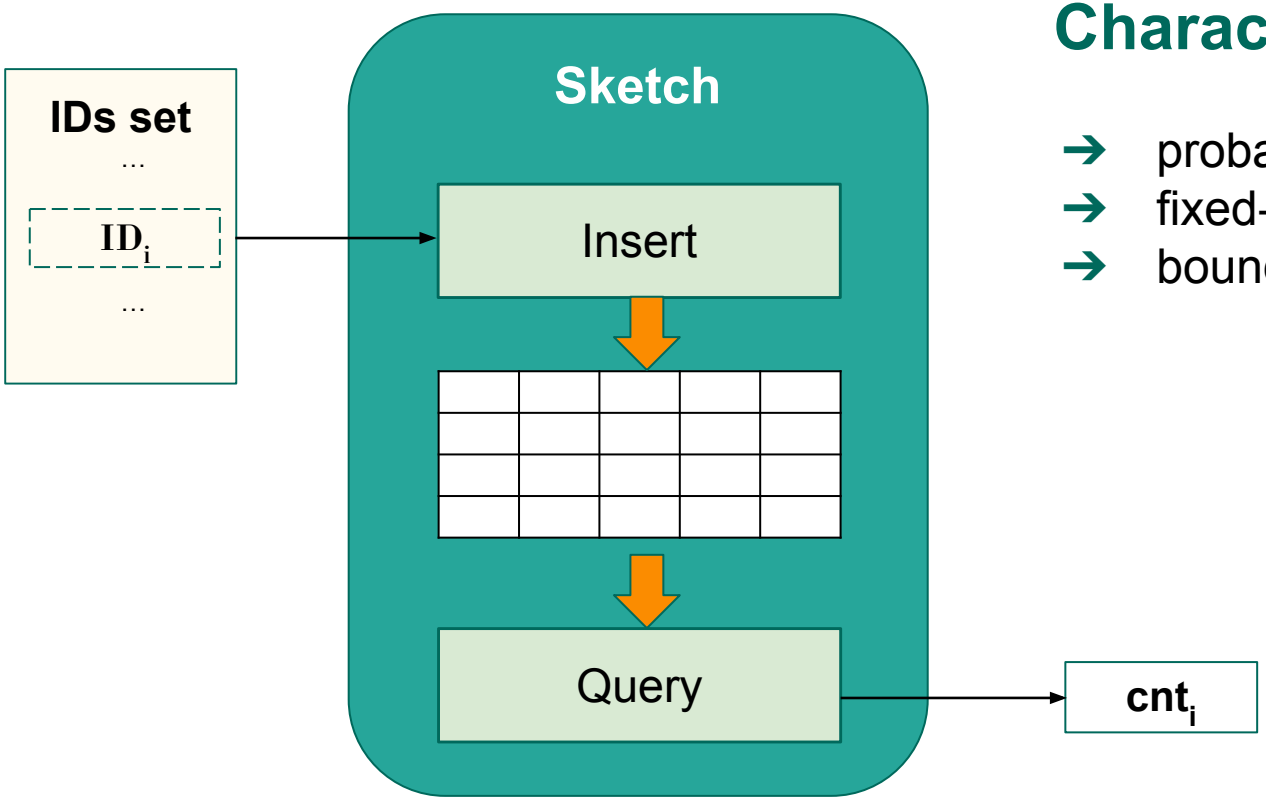

Characteristic

- ➔ probabilistic datastructure
- ➔ fixed-size
- ➔ bounded error

Count min sketch

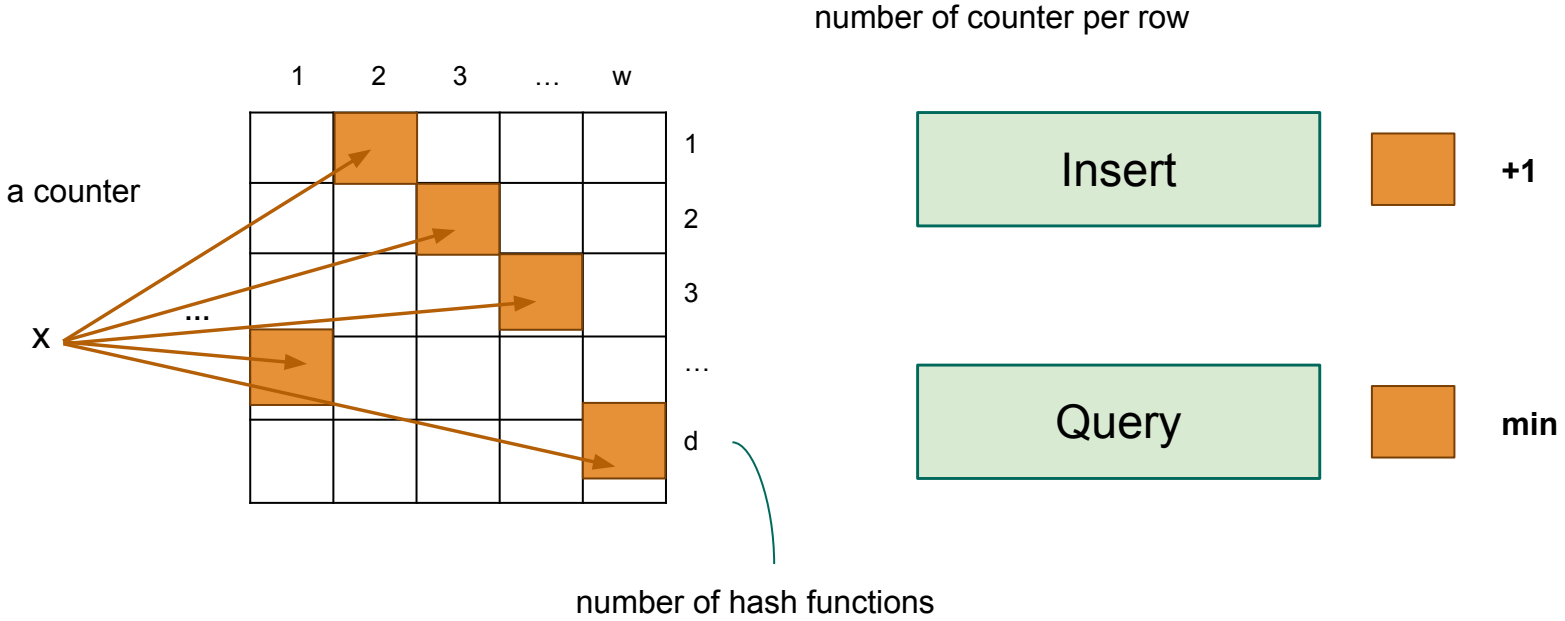

Count min sketch

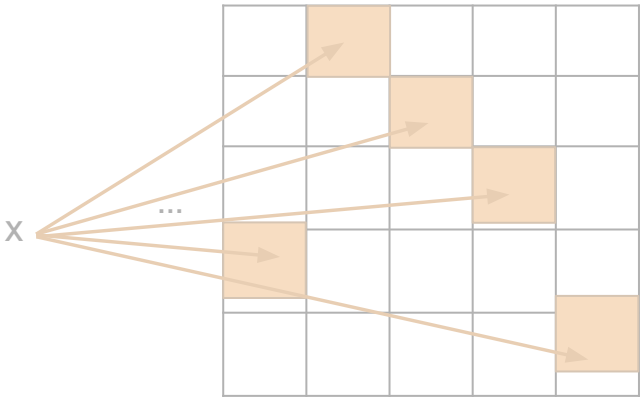

Insert

+1

Query

min

Inflated counts

Adversarial stream modeling

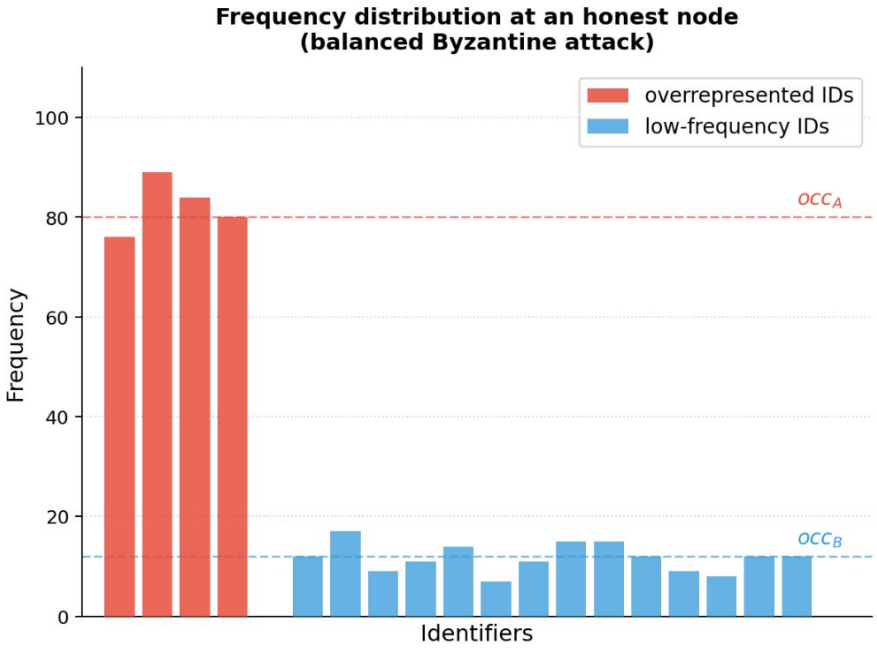

Problem with existing metrics

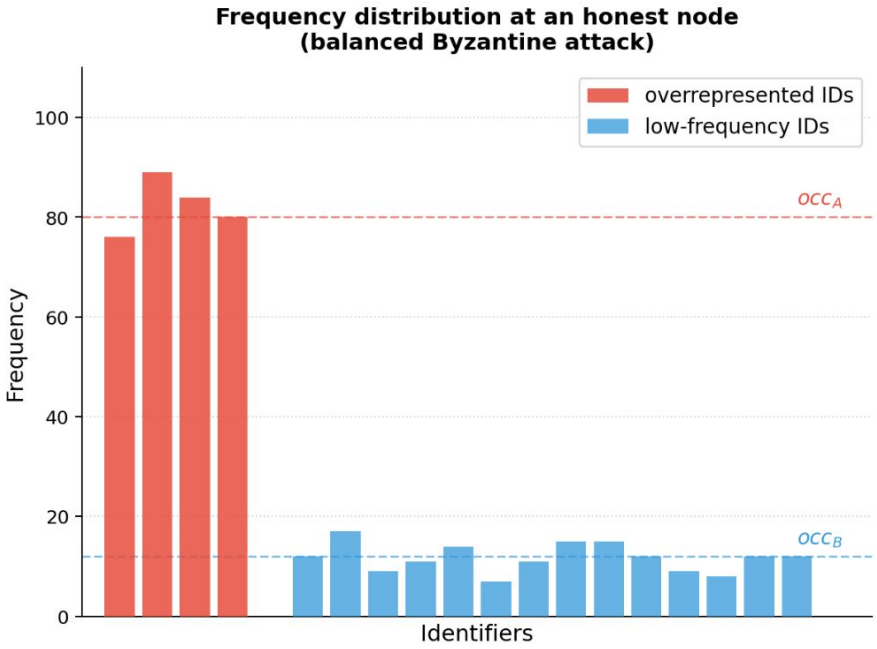

Absolute Average Error

Key metrics for adversarial context

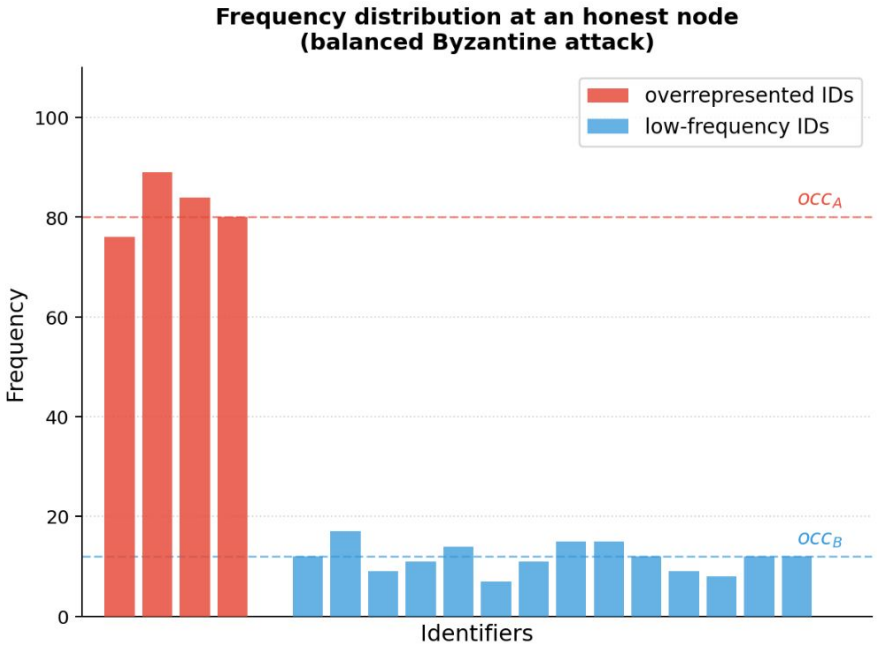

➔ preserve frequency estimation

Key metrics for adversarial context

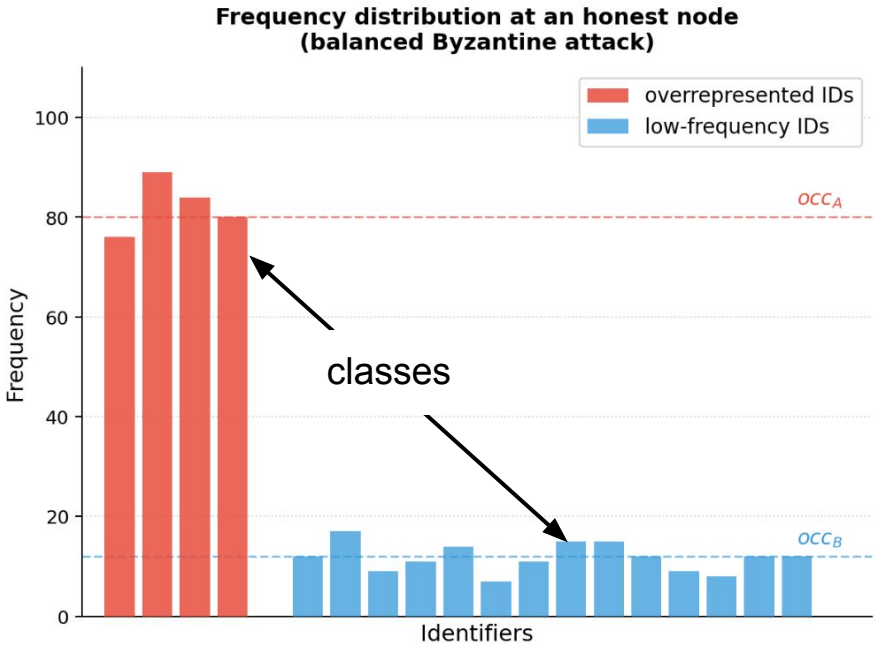

- ➔ preserve frequency estimation
- ➔ distinguish id classes

Key metrics for adversarial context

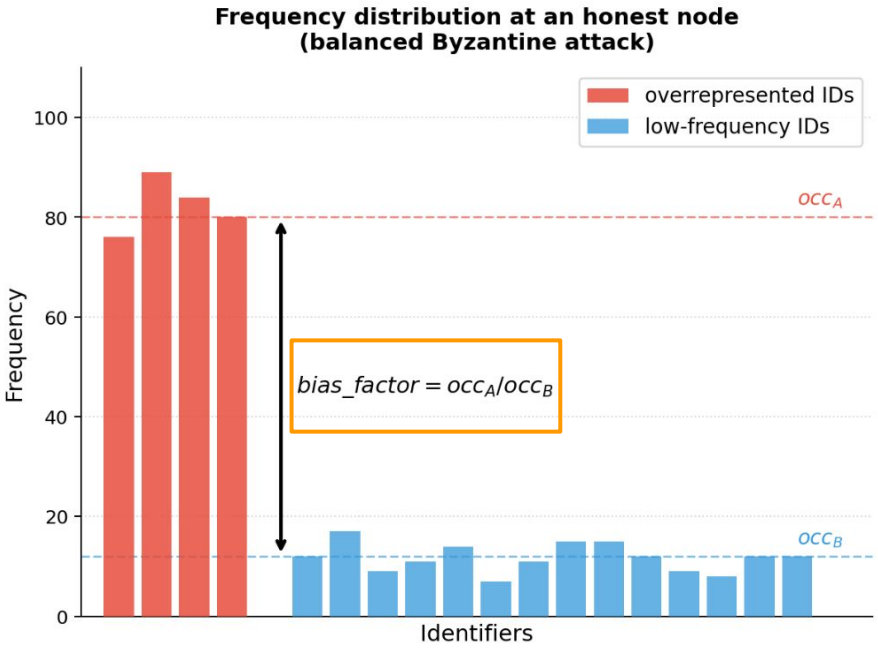

- ➔ preserve frequency estimation
- ➔ distinguish id classes
- ➔ estimate Byzantine ids overrepresentation

Bias factor

Experimental protocol

N=20K M=600K f=10%

20 to 80 KB

BM CMMCU LCU
CF CMSCU XY

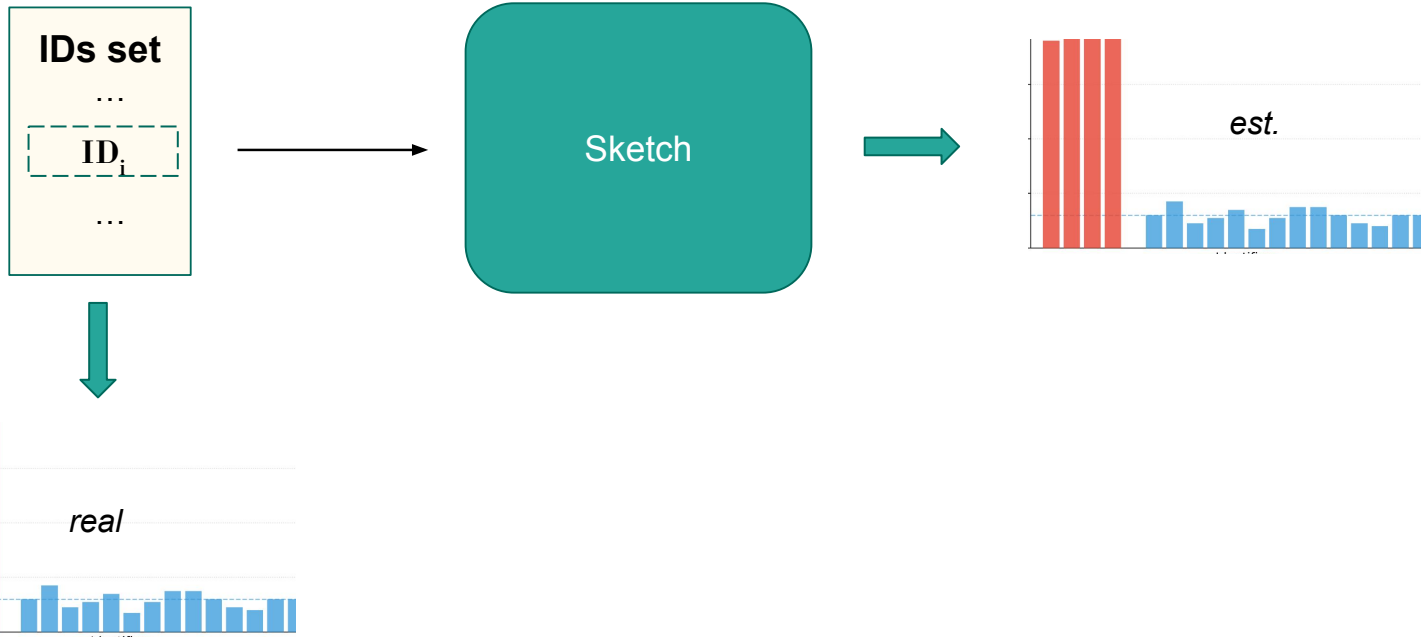

Experimental protocol

N=20K M=600K f=10%

20 to 80 KB

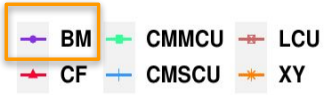

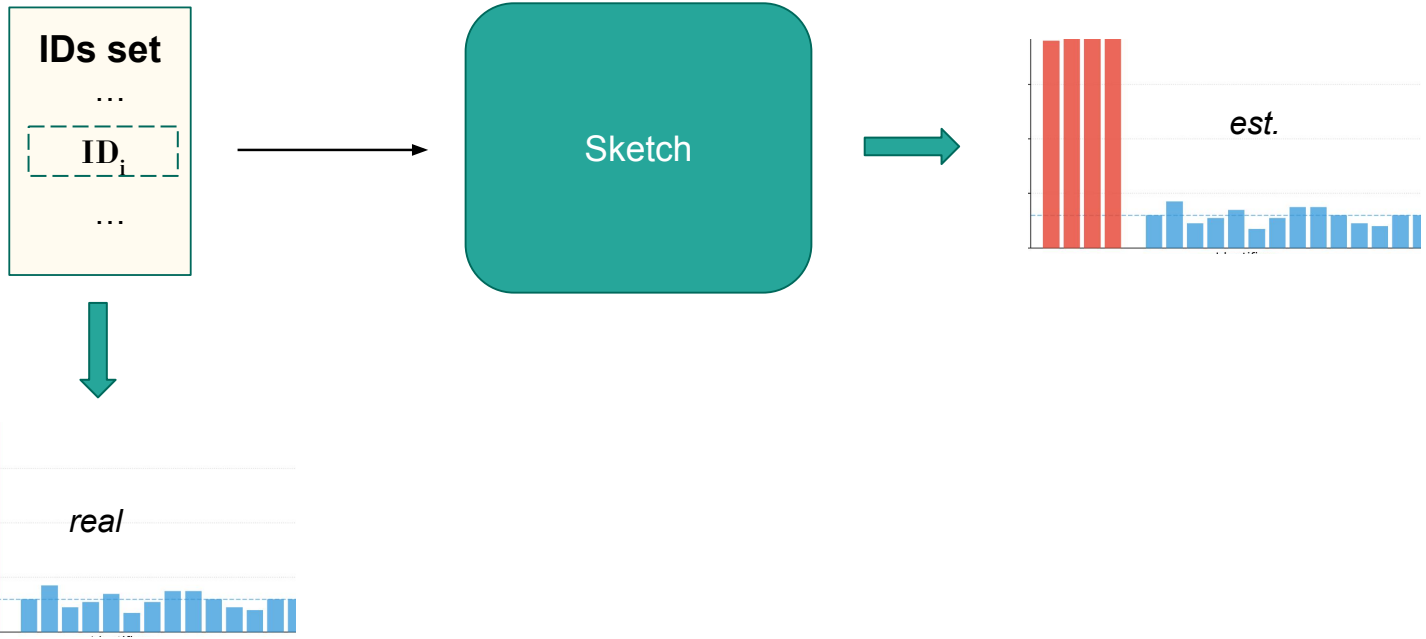

Sketches evaluation

N=20K M=600K f=10%

20 to 80 KB

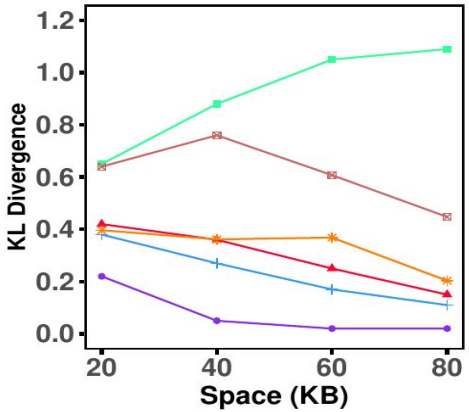

KL divergence

Sketches evaluation

N=20K M=600K f=10%

20 to 80 KB

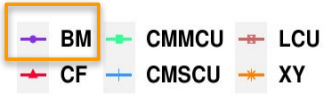

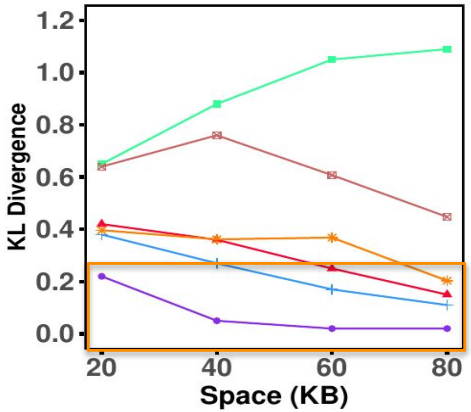

KL divergence

Sketches evaluation

N=20K M=600K f=10%

20 to 80 KB

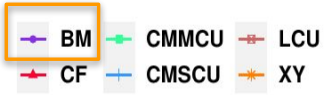

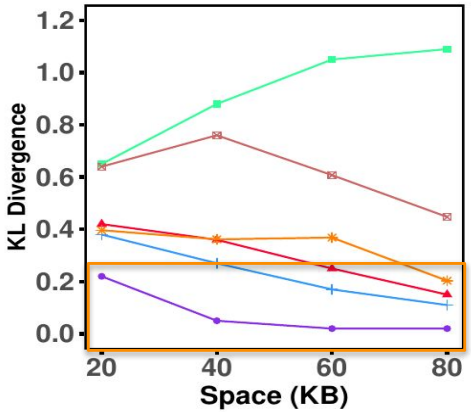

KL divergence

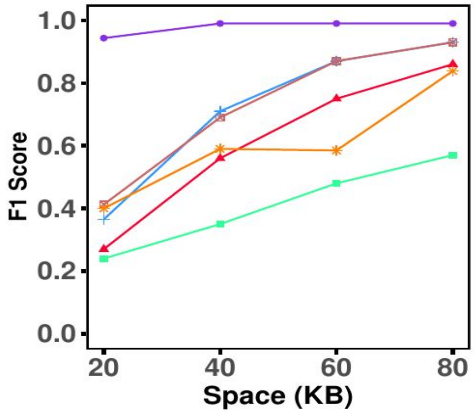

F1 score

Sketches evaluation

N=20K M=600K f=10%

20 to 80 KB

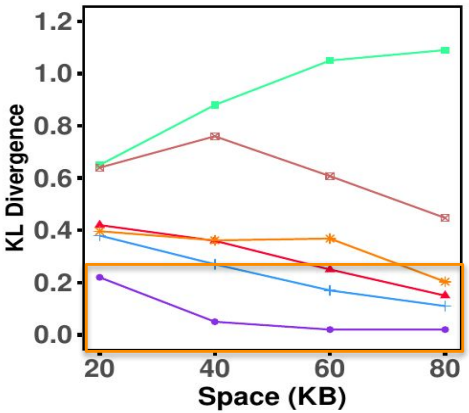

KL divergence

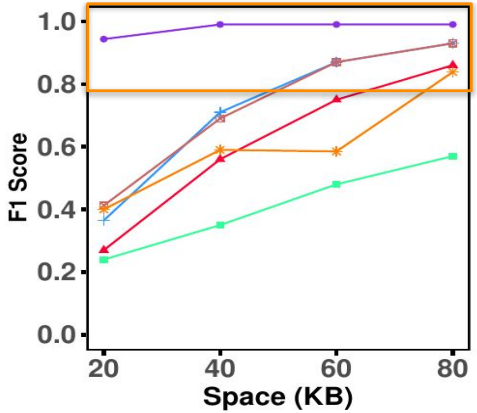

F1 score

Sketches evaluation

N=20K M=600K f=10%

20 to 80 KB

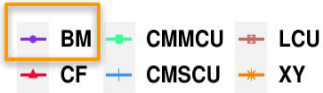

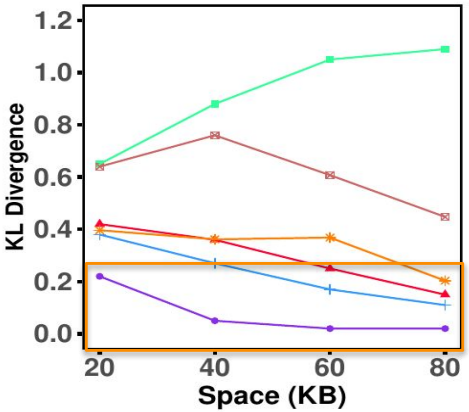

KL divergence

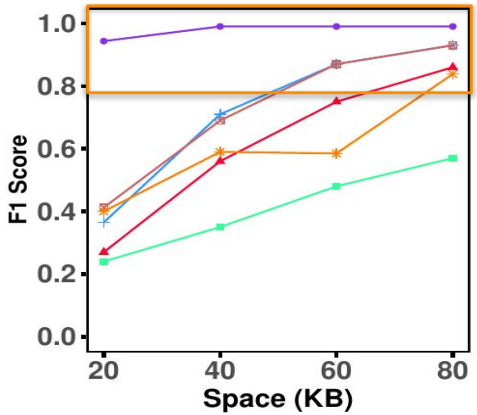

F1 score

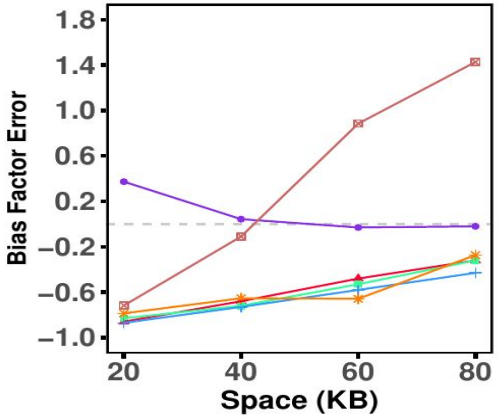

Bias factor error

Sketches evaluation

N=20K M=600K f=10%

20 to 80 KB

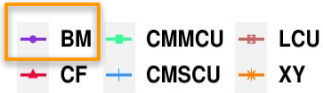

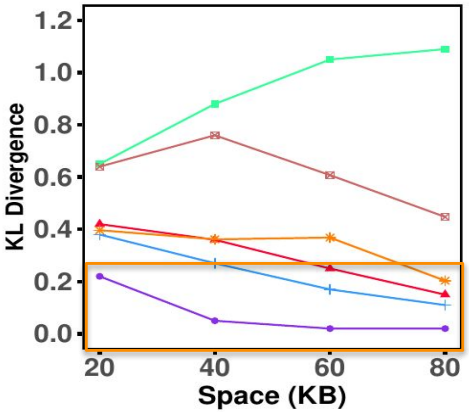

KL divergence

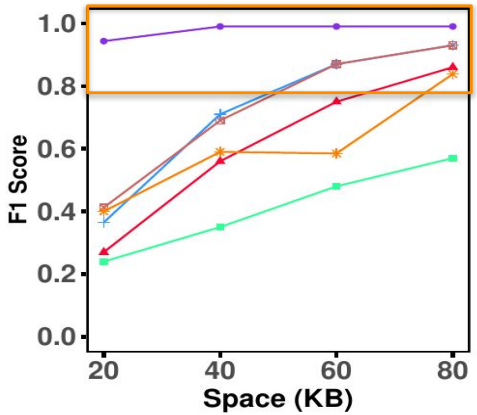

F1 score

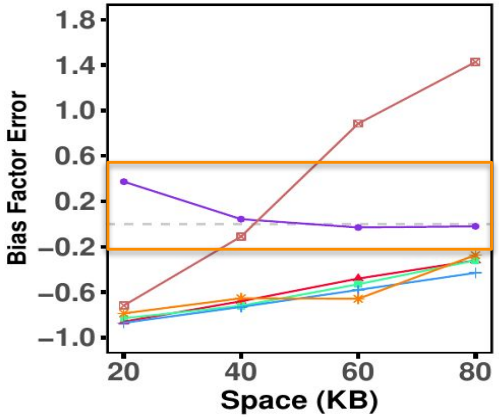

Bias factor error

Problem with the integration of sketch

How to decay the sketch?

How to aggregate two sketches?

Bitmatcher decay and merge for byzantine-resilient peer sampling

Bitmatcher overview

Characteristics

- ➔ improve memory utilization
- ➔ adjust counter sizes

Shi *et al.* (2024)

Bitmatcher overview

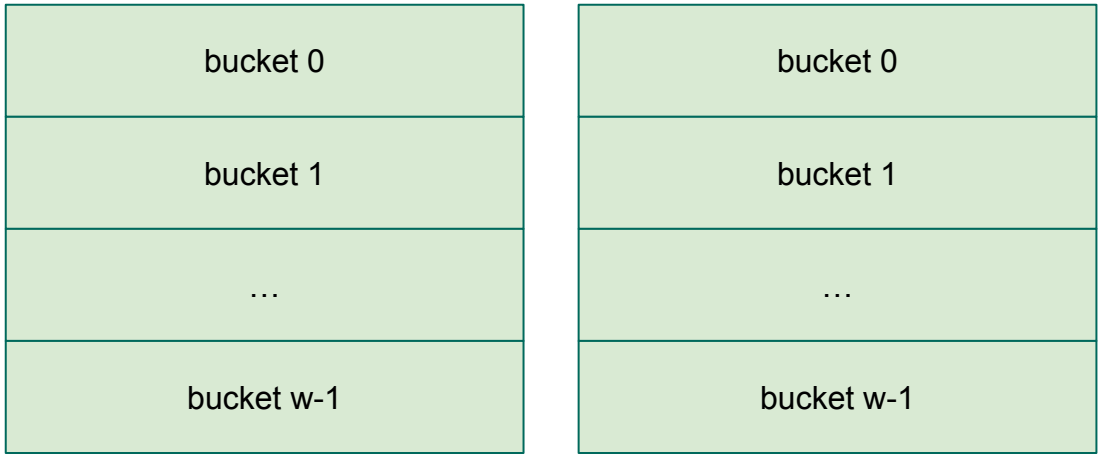

Bitmatcher overview

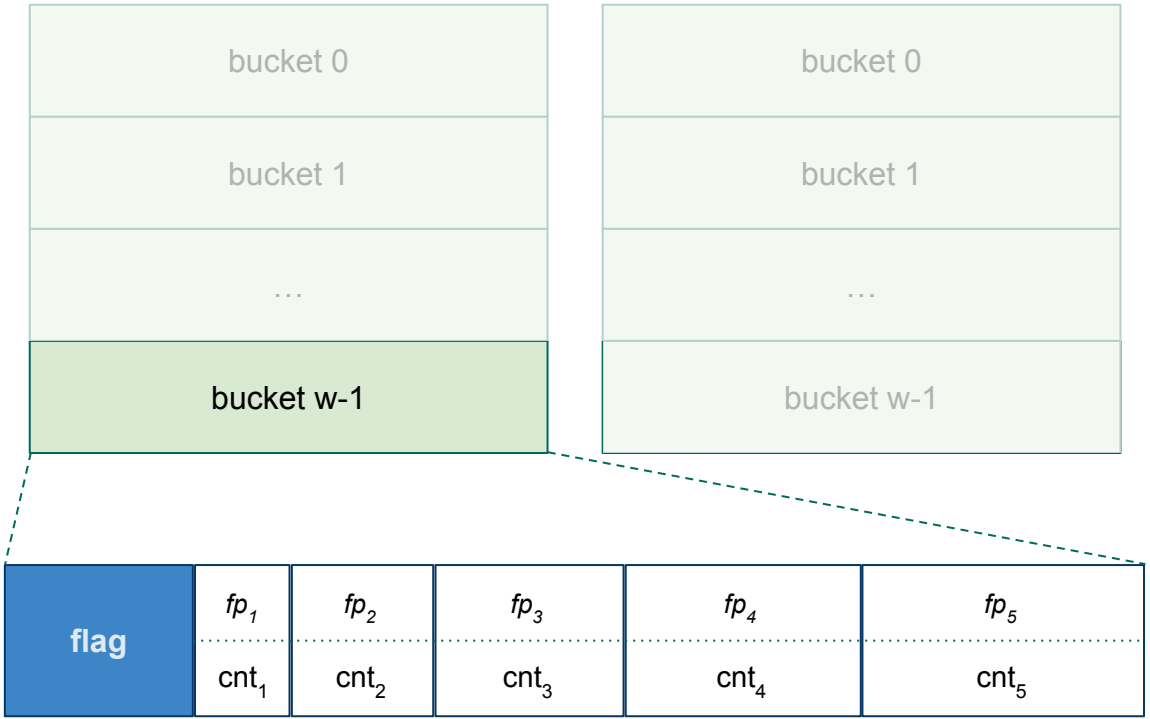

Bitmatcher overview

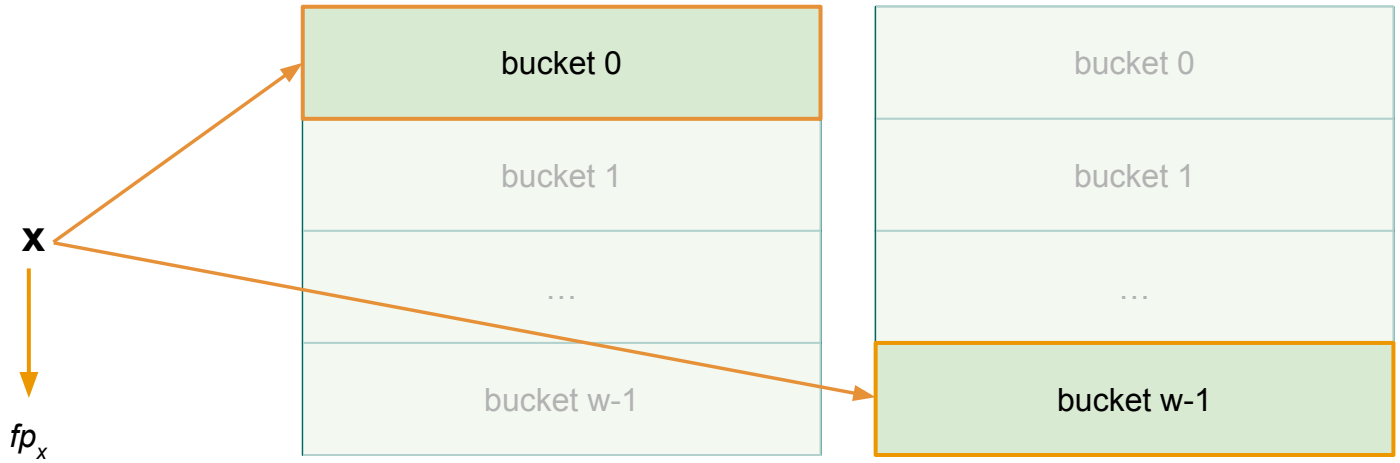

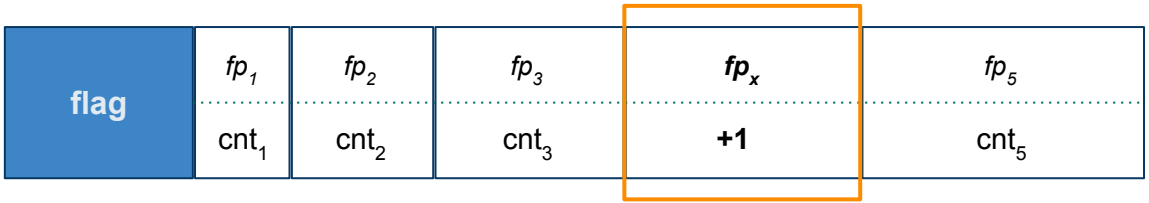

Bitmatcher overview

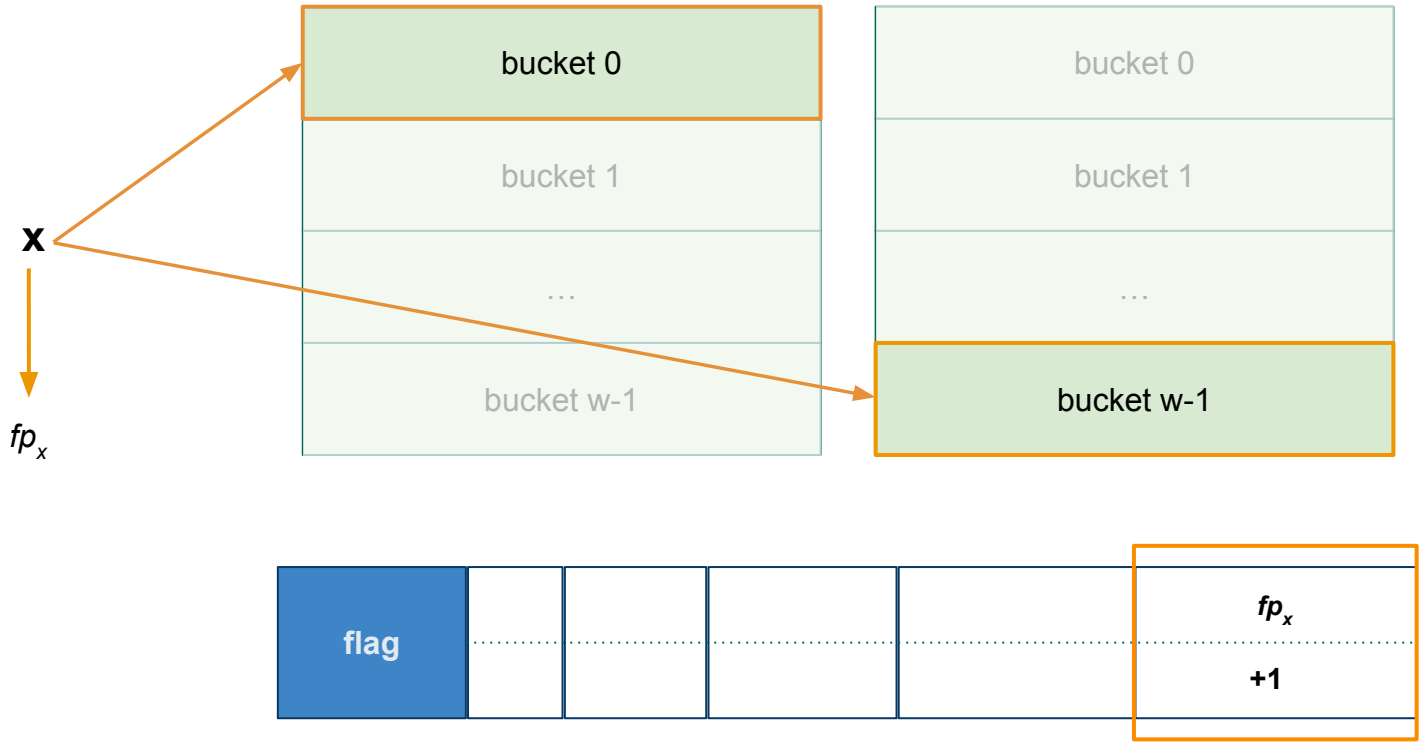

Space management

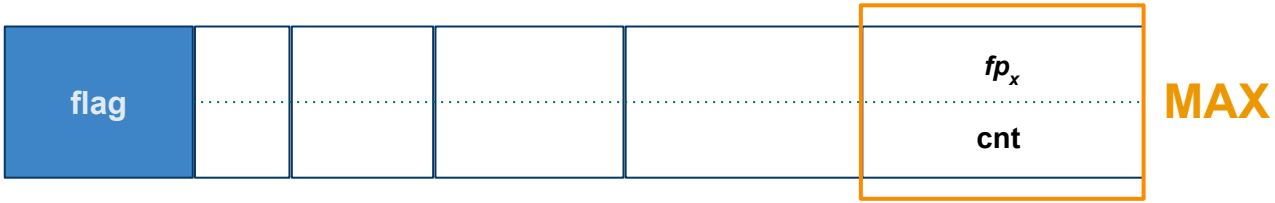

Space management

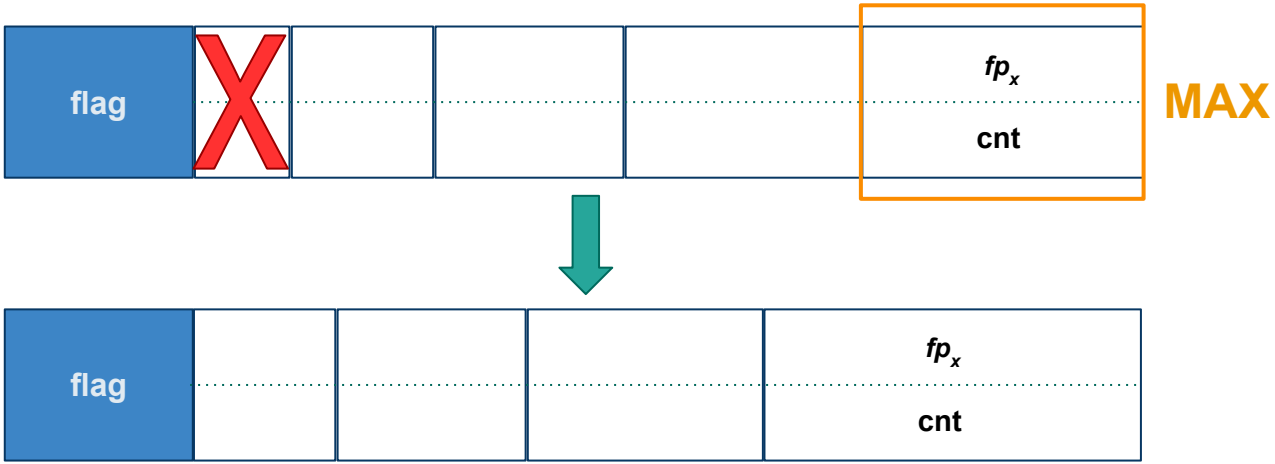

Space management

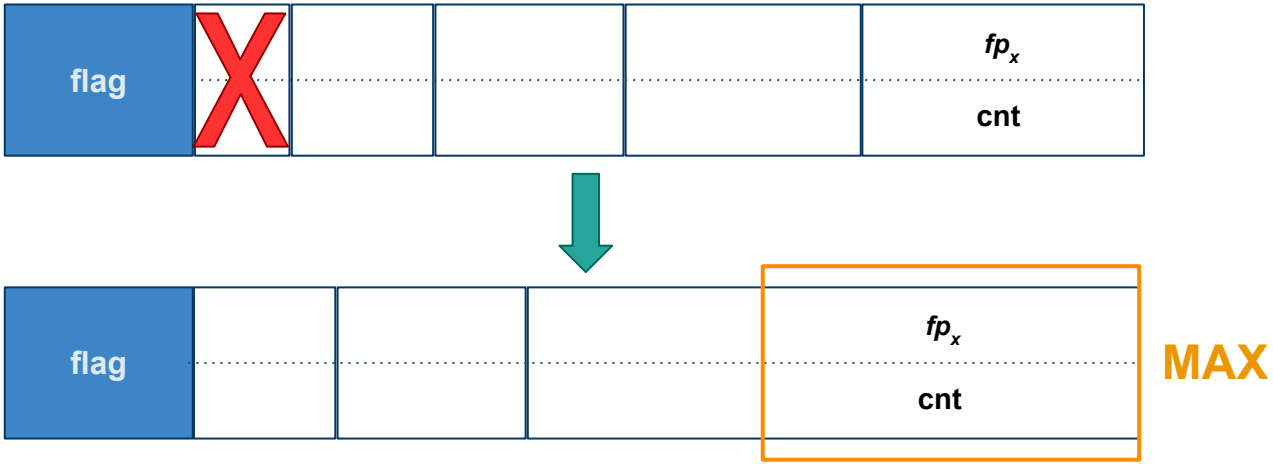

Space management

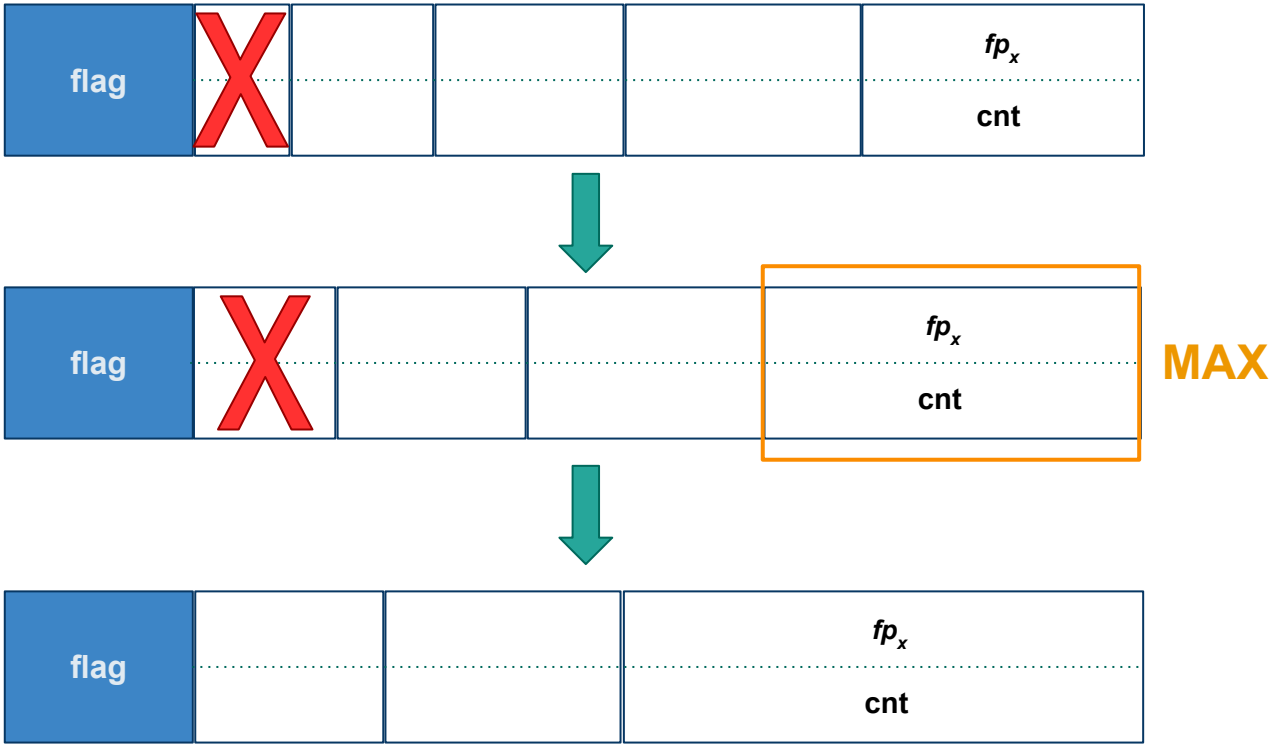

Decay challenges

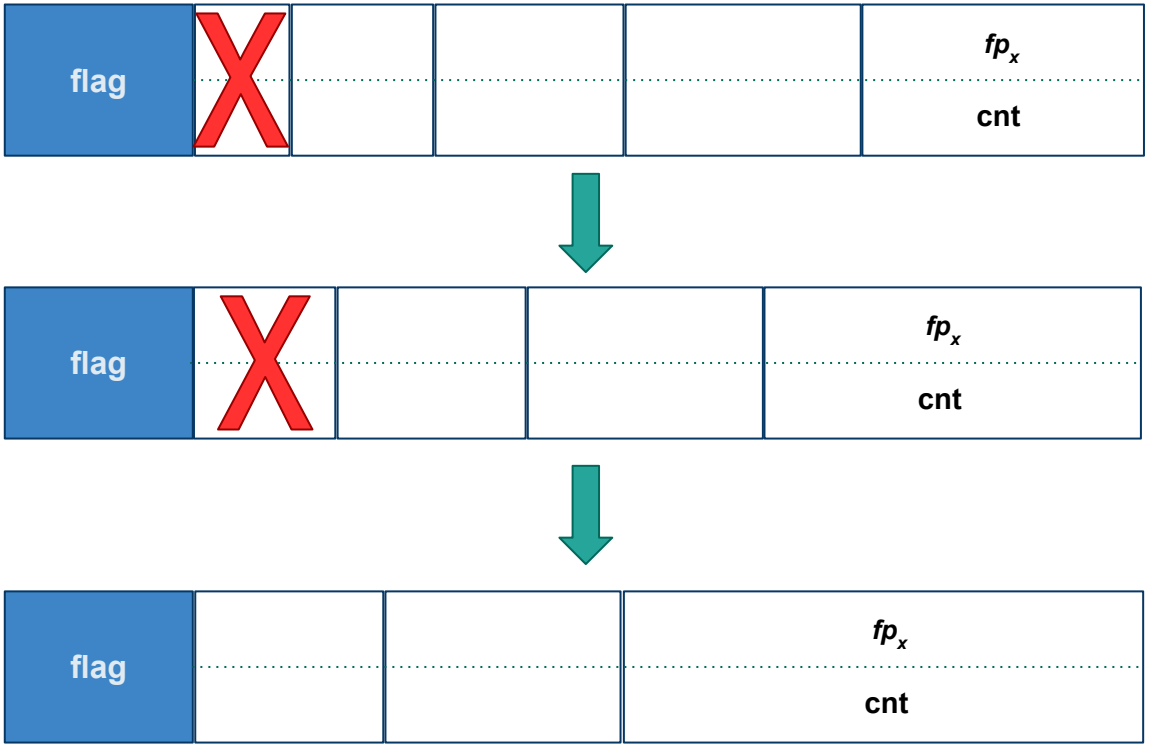

capacity = 5 entries

4 entries

3 entries

BMDecay solution

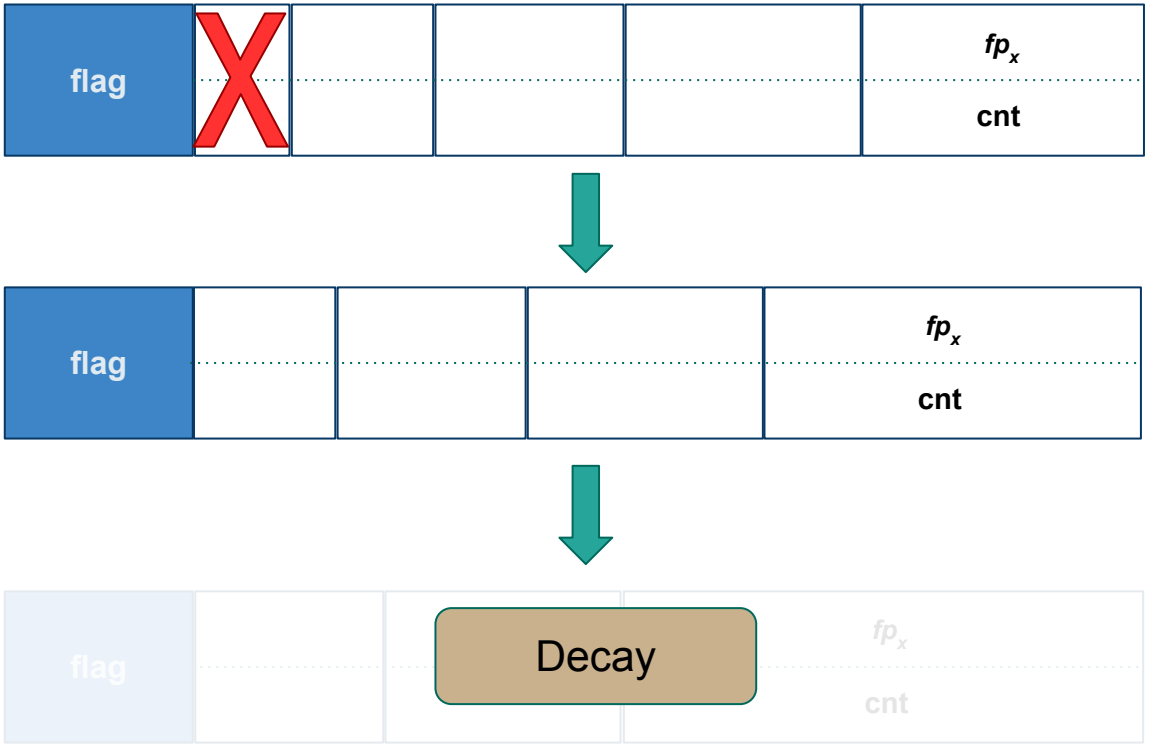

capacity = 5 entries

4 entries

3 entries

BMDecay solution

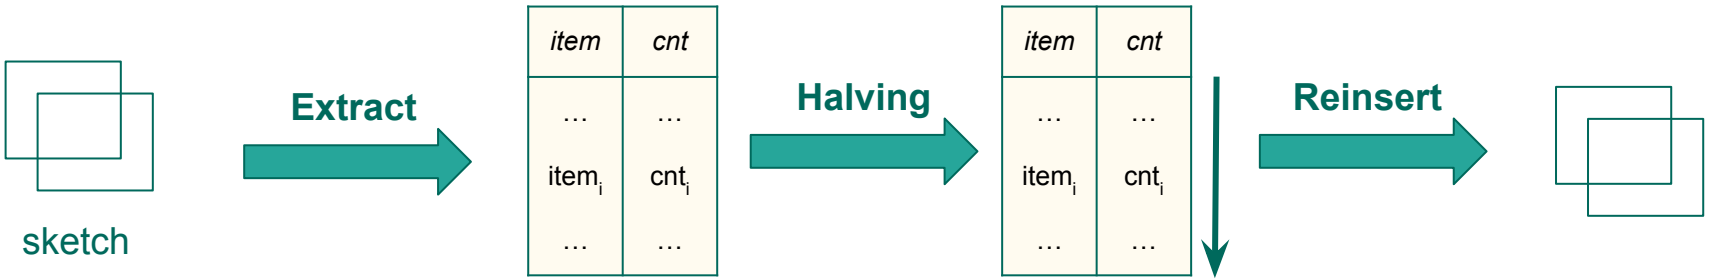

BMDecay merge

flag	a	b	c	d	e

? ? ?

flag	f	g	h	i	j

BMDecay merge

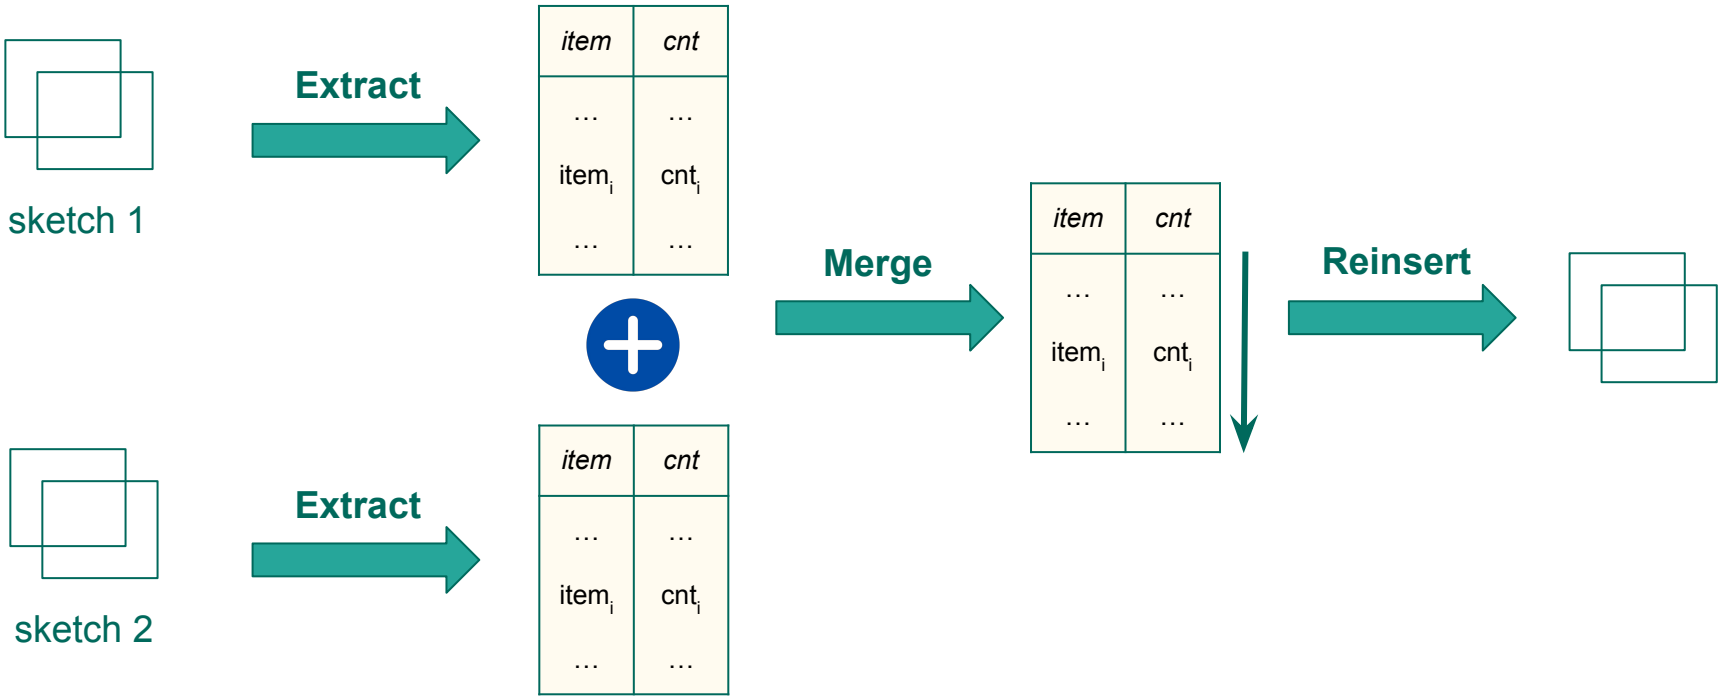

Experimental protocol

N=1K M=10K-10M f=10%

Bitmatcher BMDecay

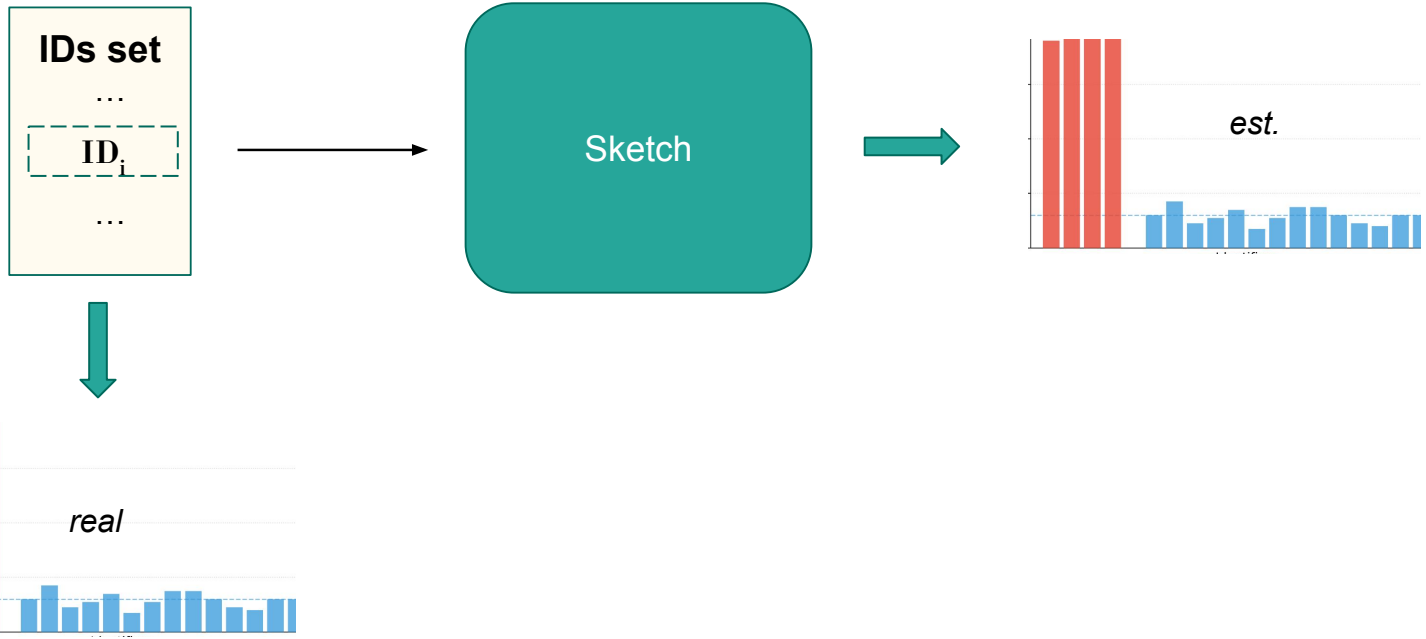

Sketches evaluation

N=1K M=10K-10M f=10%

Bitmatcher BMDecay

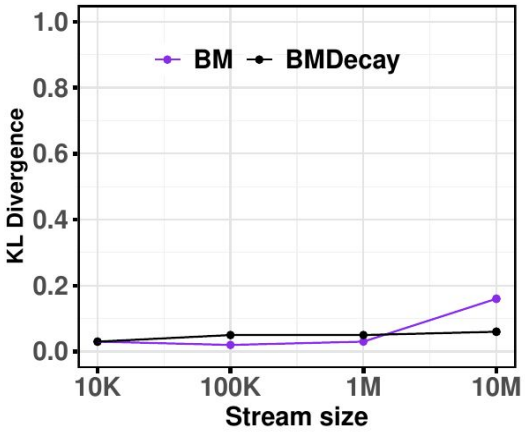

KL divergence

Sketches evaluation

N=1K M=10K-10M f=10%

Bitmatcher BMDecay

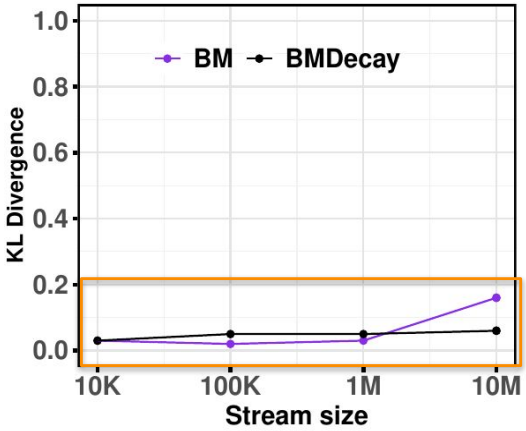

KL divergence

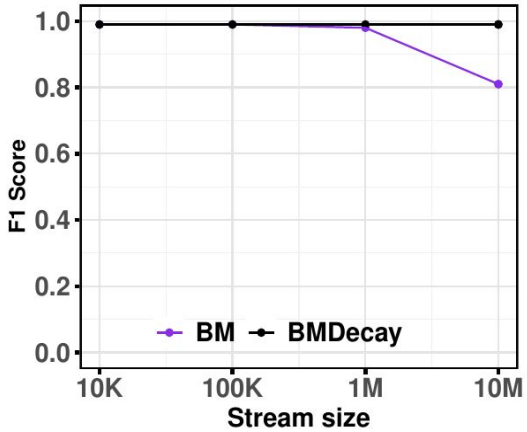

F1 score

Sketches evaluation

N=1K M=10K-10M f=10%

Bitmatcher BMDecay

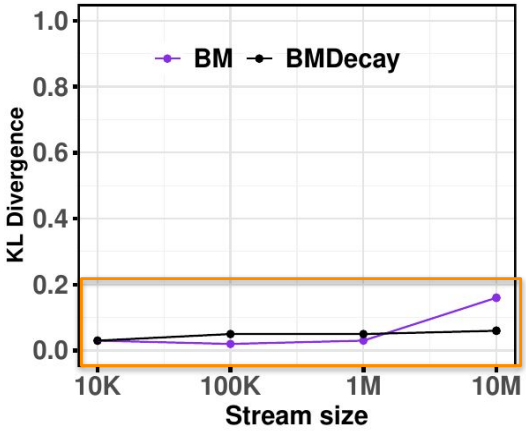

KL divergence

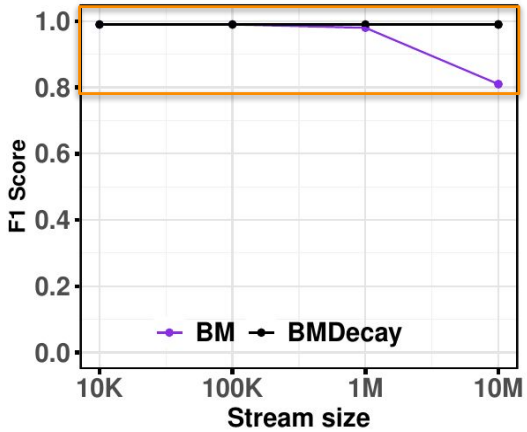

F1 score

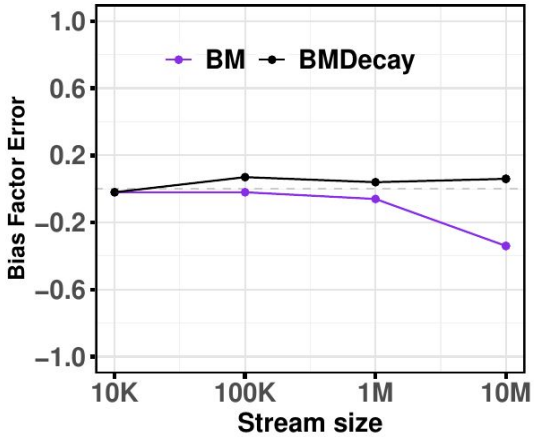

Bias factor error

Sketches evaluation

N=1K M=10K-10M f=10%

Bitmatcher BMDecay

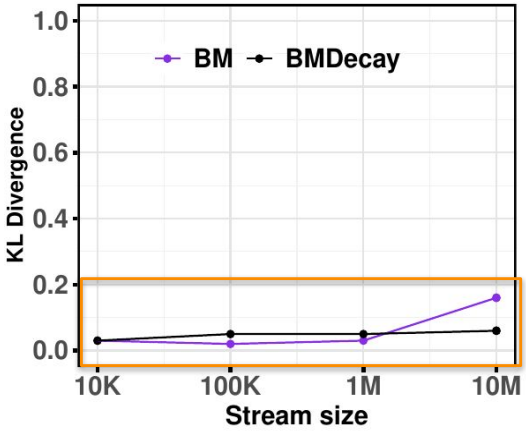

KL divergence

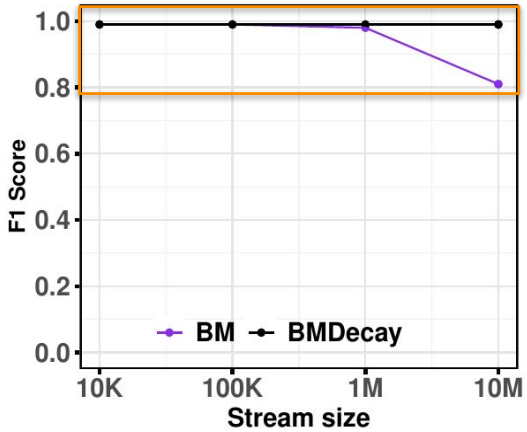

F1 score

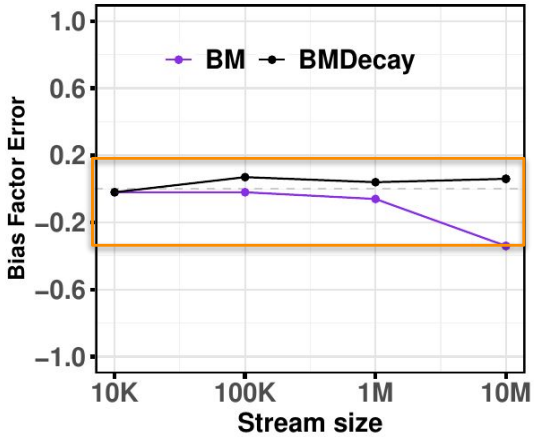

Bias factor error

Fault tolerance evaluation

N=1K f=10-30%

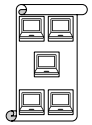

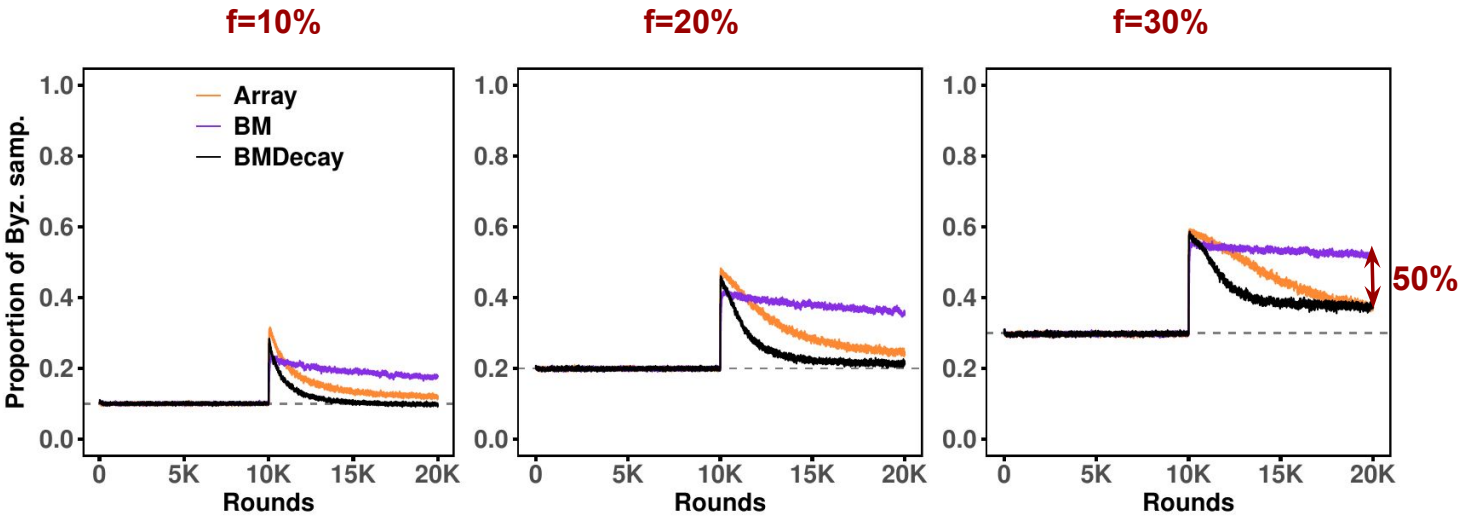

Collaborative work evaluation

N=1K f=10-30% t=10-20%

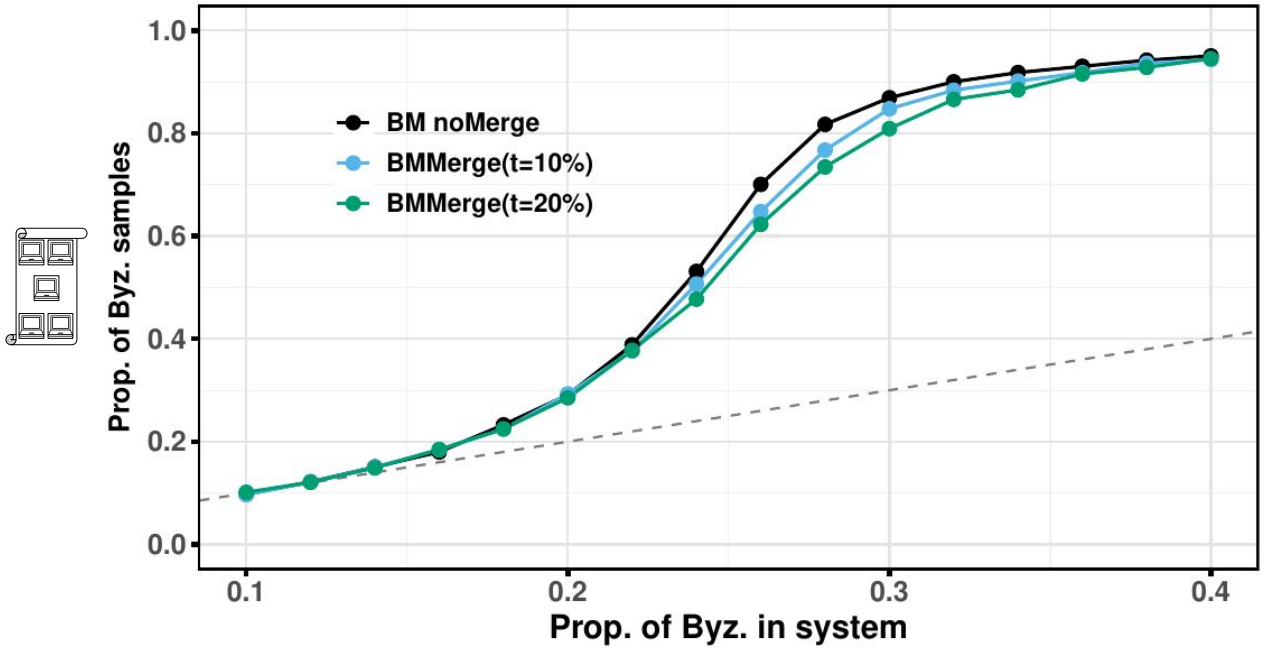

Conclusion

1. Improve tolerance to Byzantine attacks

AUPE
Counting
Gains: 60% on Brahms

Mukam *et al.* (2024)

Conclusion

1. Improve tolerance to Byzantine attacks

AUPE
Counting
Gains: 60% on Brahms

Mukam *et al.* (2024)

2. Count estimation for scalability

Comparative study of various sketches
Bias factor
Bitmatcher

Conclusion

1. Improve tolerance to Byzantine attacks

AUPE
Counting
Gains: 60% on Brahms

Mukam *et al.* (2024)

2. Count estimation for scalability

Comparative study of various sketches
Bias factor
Bitmatcher

3. Unbounded stream count estimation

Bimatcher decay and merge
Integration in Aupe
Comparable to that of exact counting

Mukam *et al.* (in preparation)

Perspectives

1. Additional analysis of AUPE

   ➔ Impact on convergence

Perspectives

1. Additional analysis of AUPE

   ➔ Impact on convergence

2. Deployment and real-world evaluation

   ➔ churn

Perspectives

1. Additional analysis of AUPE

   ➔ Impact on convergence

2. Deployment and real-world evaluation

   ➔ churn

3. Other sampling methods

LABORATOIRE BORDELAIS DE RECHERCHE EN INFORMATIQUE

LaBRI

cnrs

université de BORDEAUX

Byzantine-Resilient Peer Sampling for Large-Scale Distributed Systems

Thesis defended on April 24, 2026, by **Augusta Mukam**

Before a committee consisting of :

Mme Sonia Ben Mokhtar	Research Director, CNRS	Reviewer
M. Yérom-David Bromberg	Professor, University of Rennes	Reviewer
Mme Patricia Thébault	Professor, University of Bordeaux	Examiner
M. Léo Mendiboure	Associate Professor, University of Pau and the Pays de l'Adour	Examiner
M. Laurent Réveillère	**Full Professor, University of Bordeaux**	**Thesis Advisor**
M. Joachim Bruneau-Queyreix	**Associate Professor, Bordeaux INP**	**Co-Advisor**