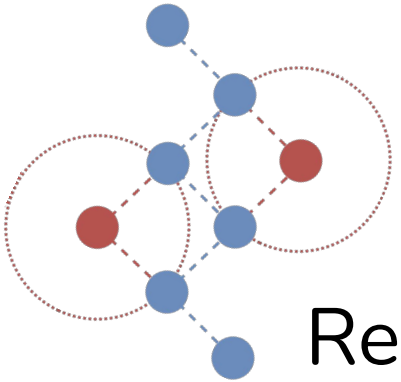

Reliable Distributed Consensus for Low-Power Multi-Hop Networks

Author
Alberto Spina

Supervisors
Prof. Julie McCann
Dr. Michael Breza

Second Marker
Dr. Anandha Gopalan

Overview

Background

Wireless Sensor Networks

Synchronous Transmissions

Agreement Protocols

Consensus

Problems

Latency

Reliability

Configurability

Replicability

Contributions

Hybrid

WISP

WIMP

Wireless Sensor Network Nodes

One **Microcontroller** (MCU)

Multiple **Sensors/Actuators**

One Wireless Communicator or **Radio**

Battery

Wireless Sensor Networks

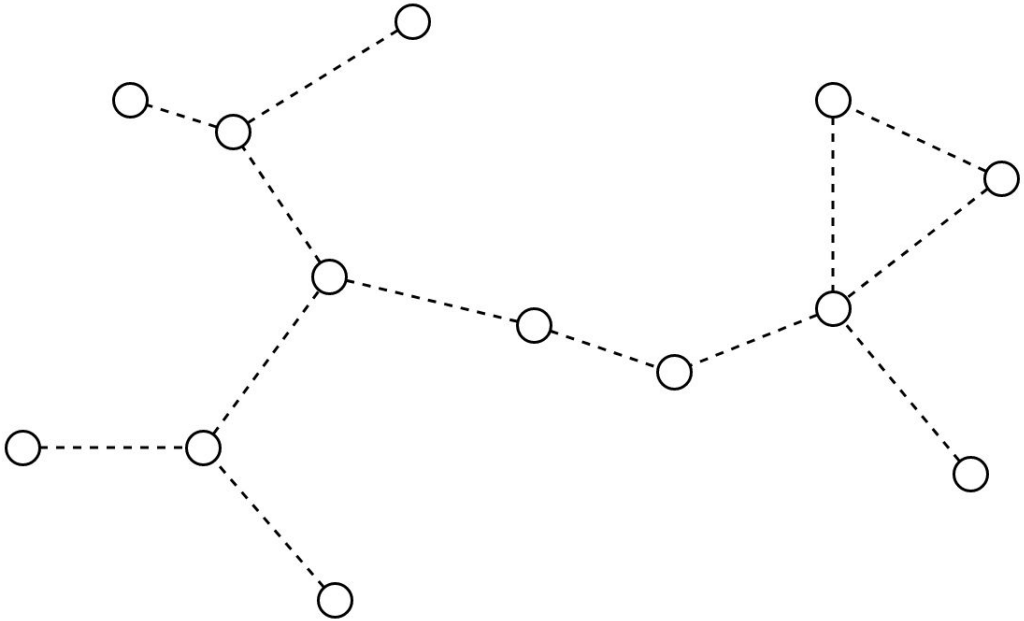

Wireless Sensor Networks

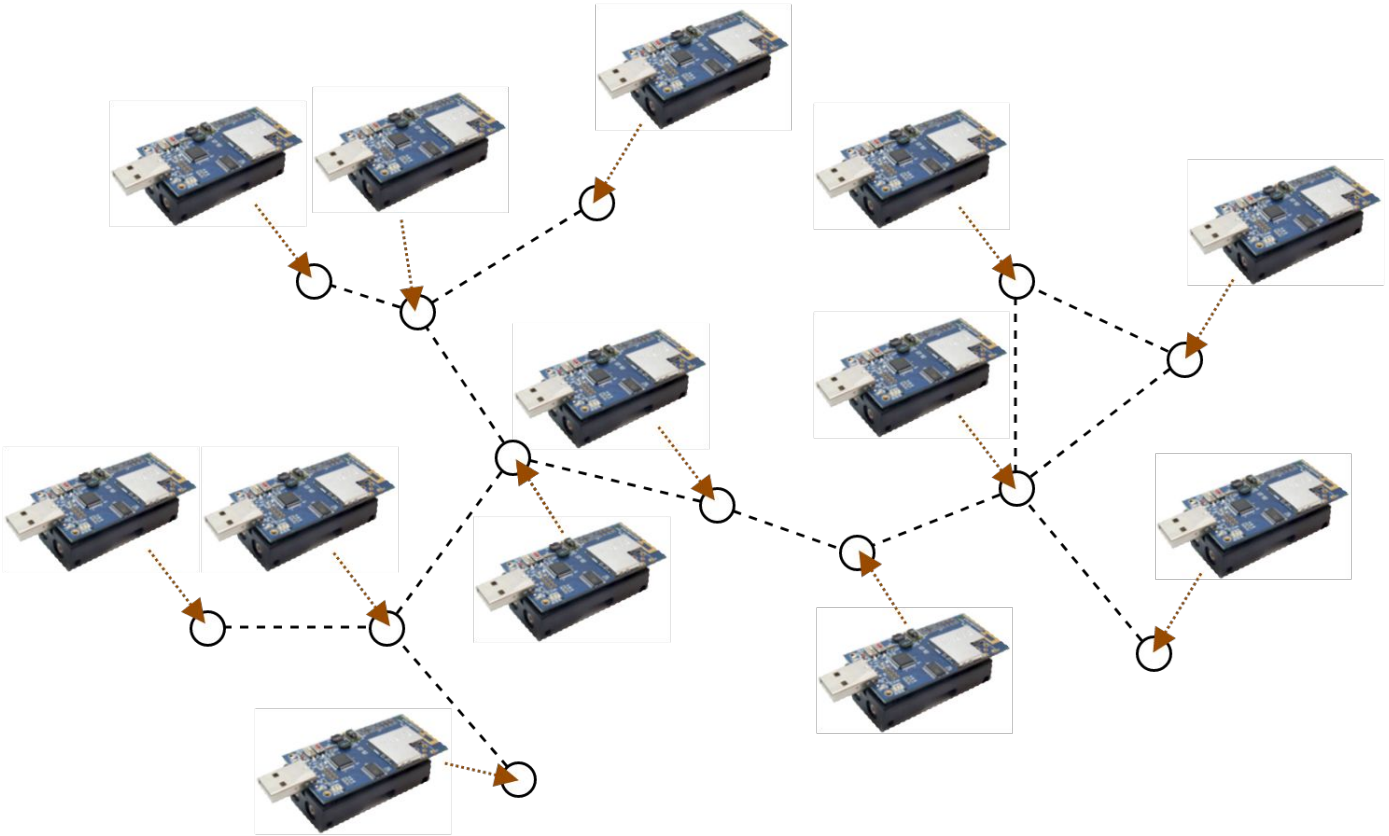

Wireless Sensor Networks

Use **low-power** radios

Environment causes **high interference**

Links are **unreliable**

FlockLab: a WSN testbed

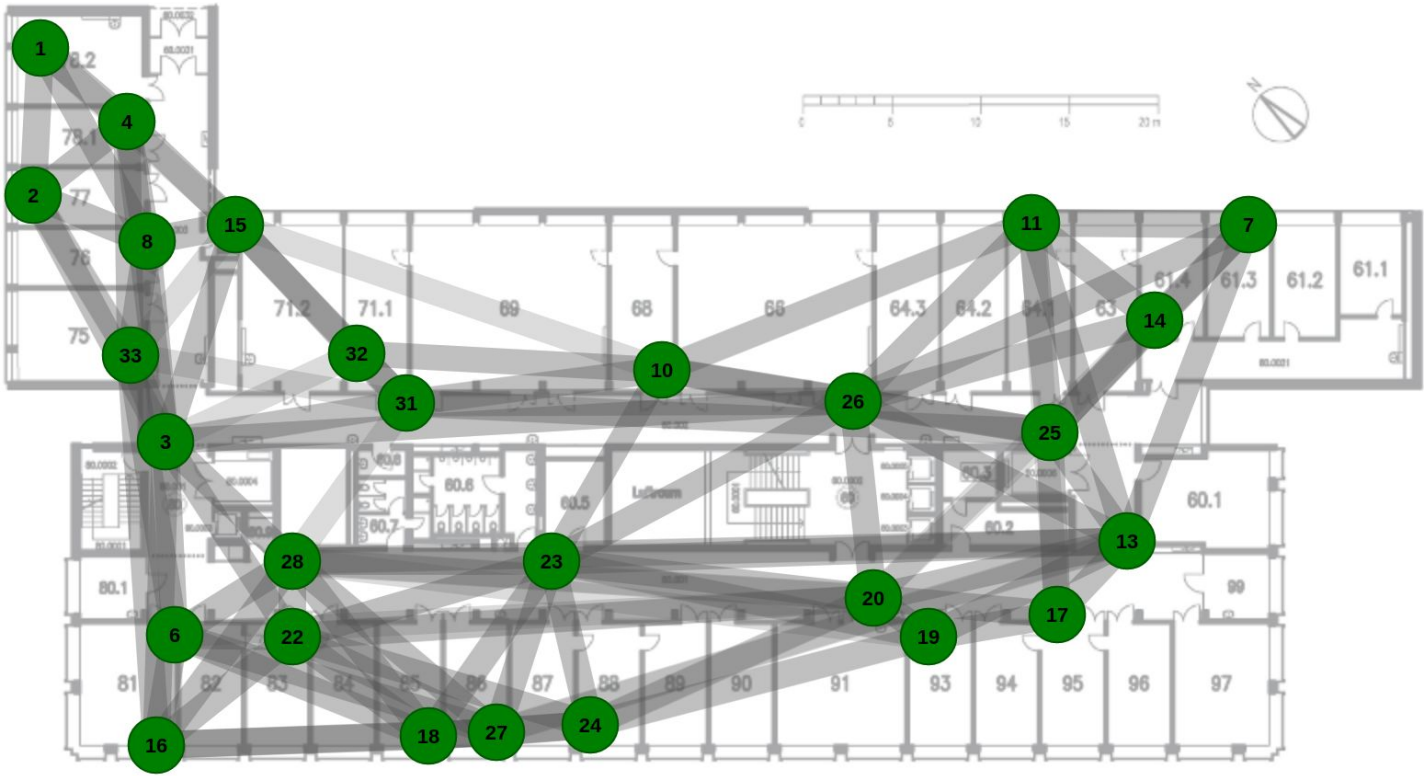

Broadcast Interference

Phase difference is perceived as a **time offset** Δt

Destructive Interference for Δt odd multiple of π

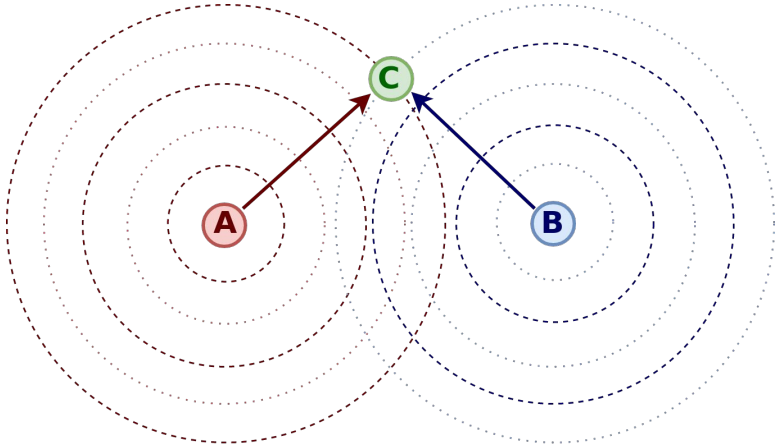

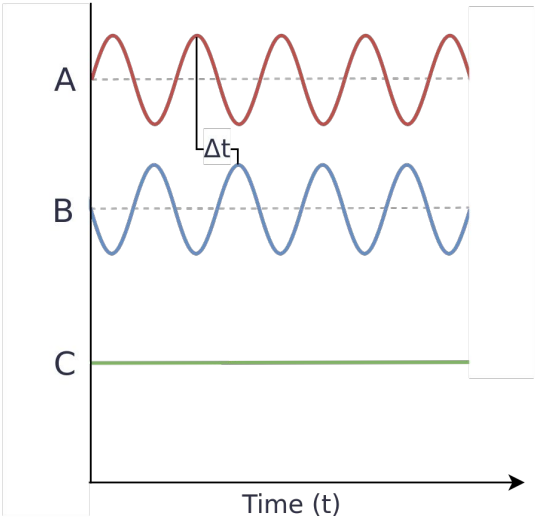

Broadcast Interference

Increase the **reliability** of transmissions

Constructive Interference for Δt multiple of 2π

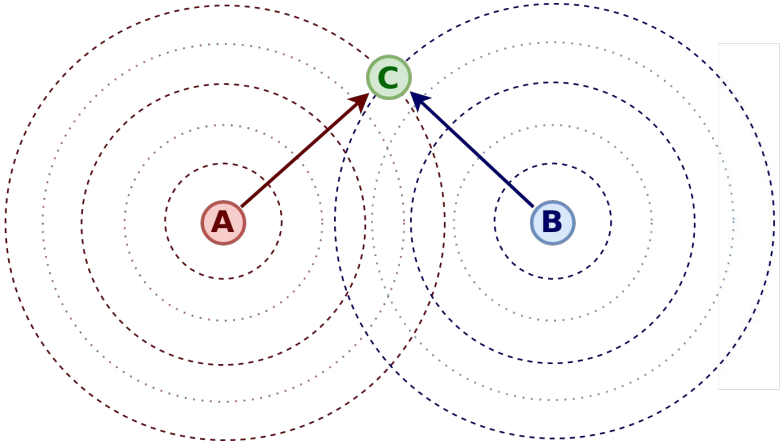

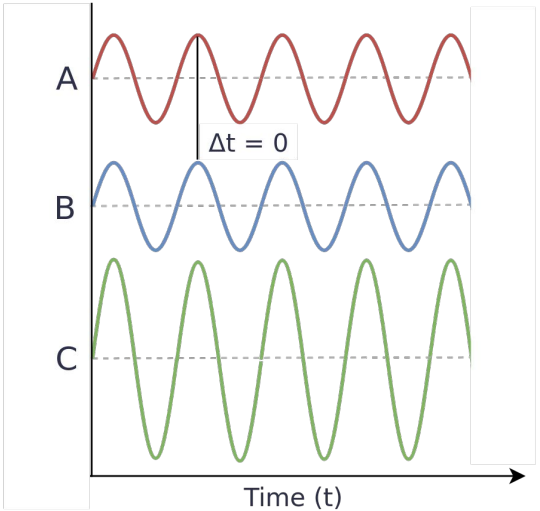

Synchronous Transmissions

Constructive Interference
+
Capture Effect

Repeatedly **flood packets** to the whole network

Enables **many-to-all communication**

Route-less packet broadcasts

Constant **winner at** EWSN **dependability competitions**

Synchronous Transmissions

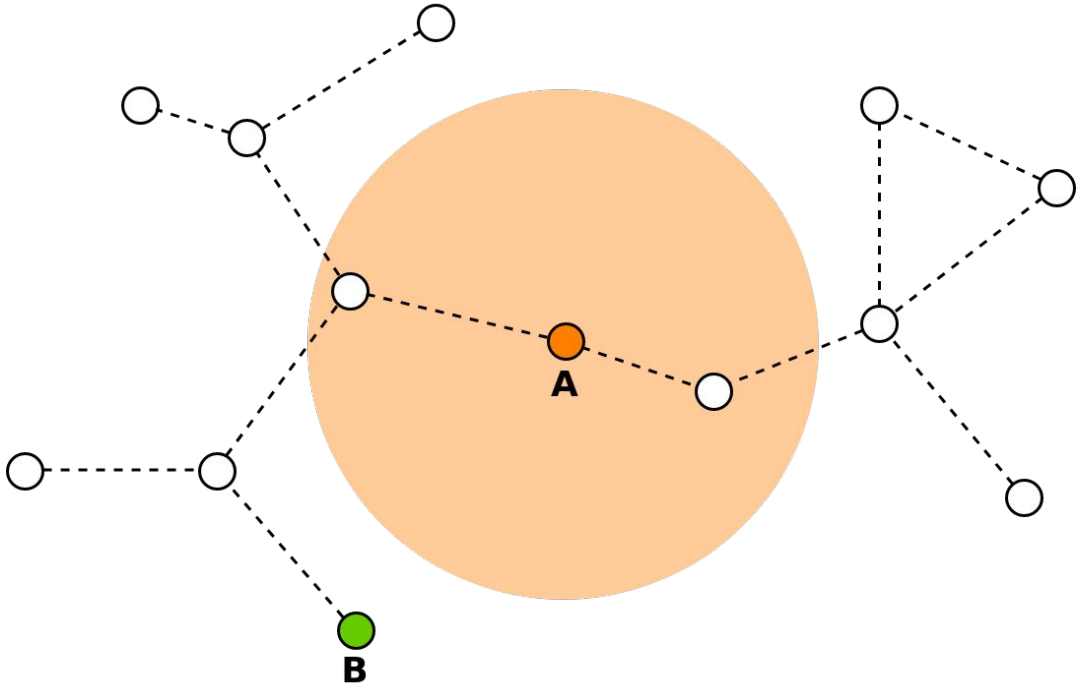

Synchronous Transmissions

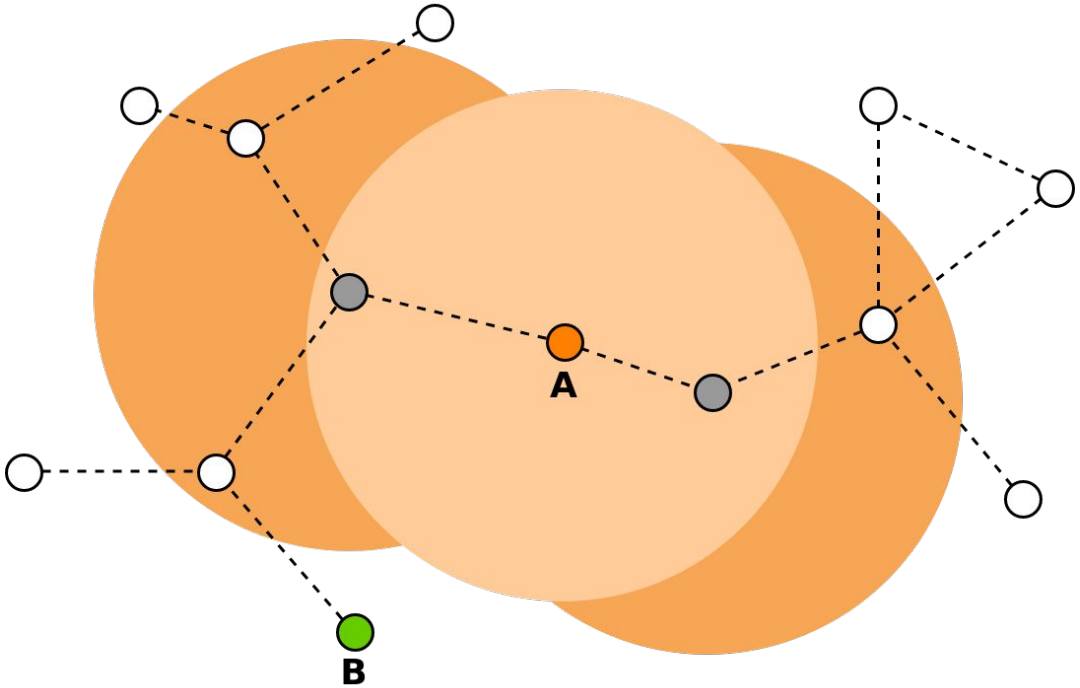

Synchronous Transmissions

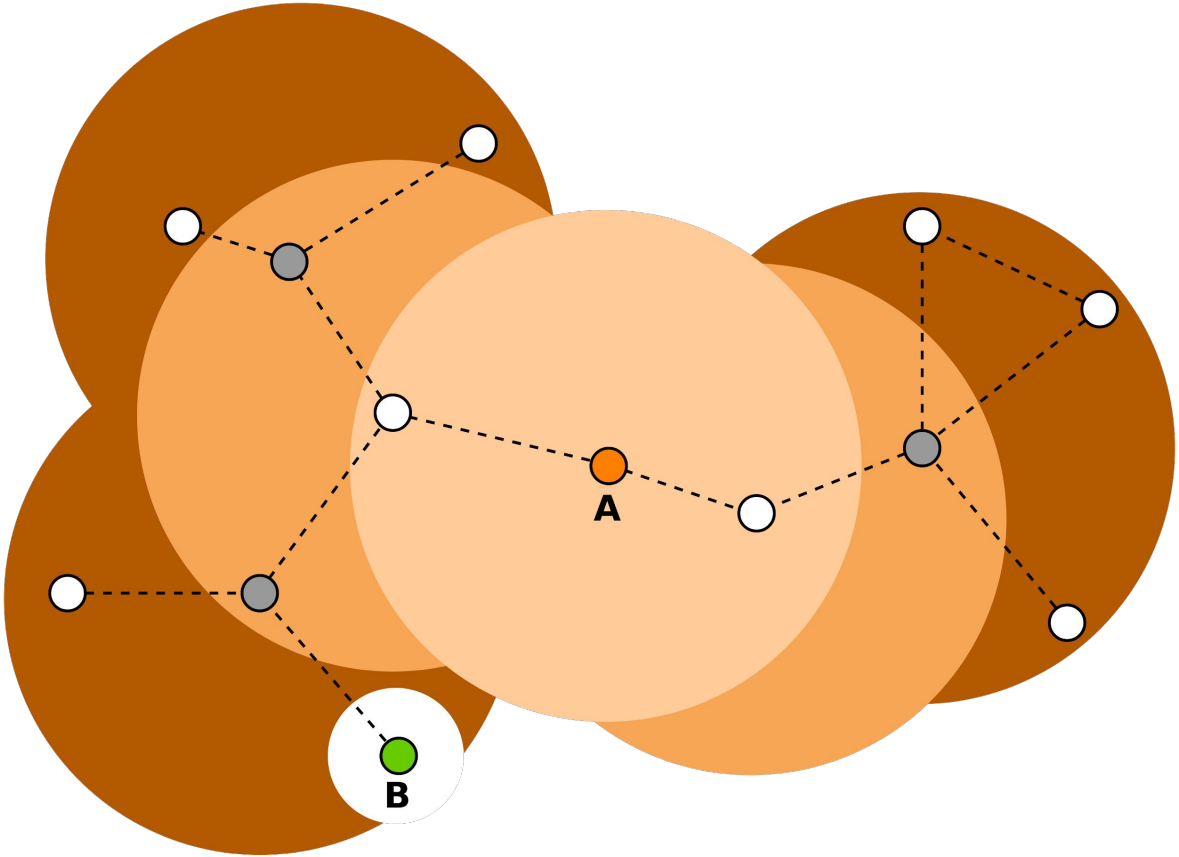

Synchronous Transmission Primitives

2011	**Glossy**
2013	**Chaos**
2013	Splash
2013	SCIF
2014	P3
2015	Pando
2016	RedFixHop
2017	LiM
2017	Robust Flooding
2018	Mixer
2018	Codecast

Glossy

F. Ferrari et al. (2011)

Synchronous Transmission (ST) **primitive**

One-to-all data dissemination

High reliability guarantees

Glossy

F. Ferrari et al. (2011)

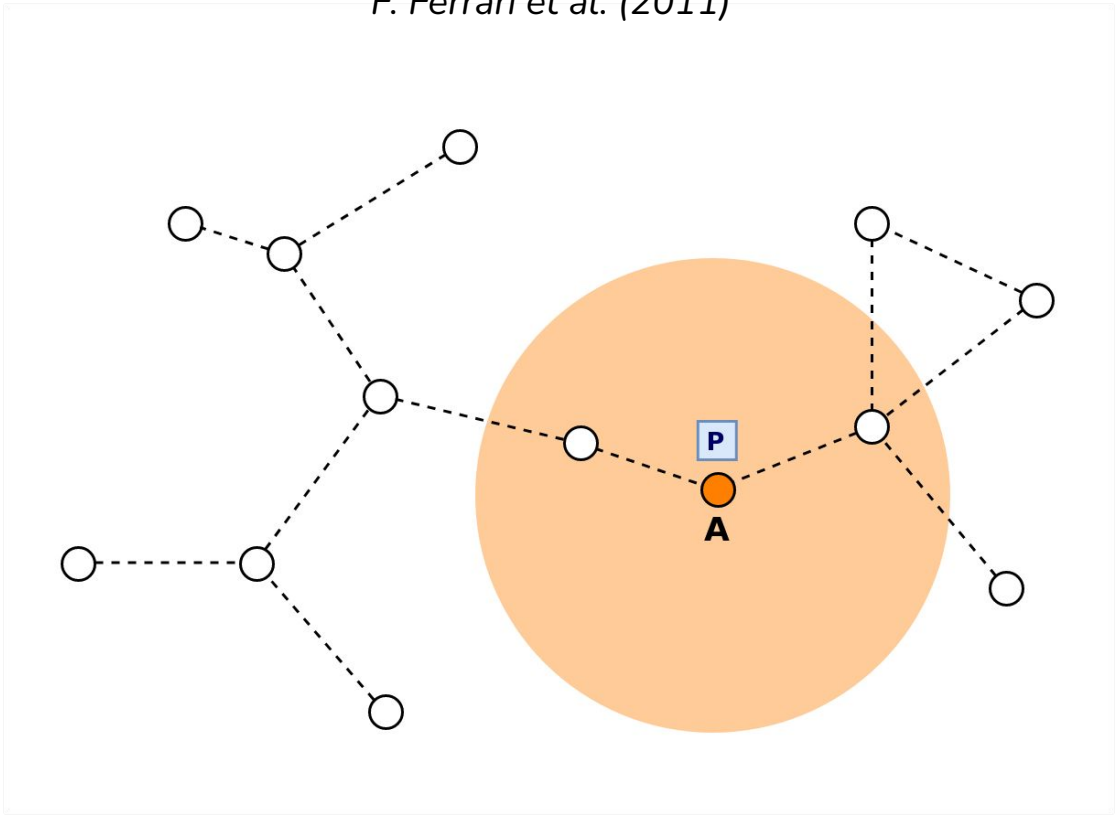

Glossy

F. Ferrari et al. (2011)

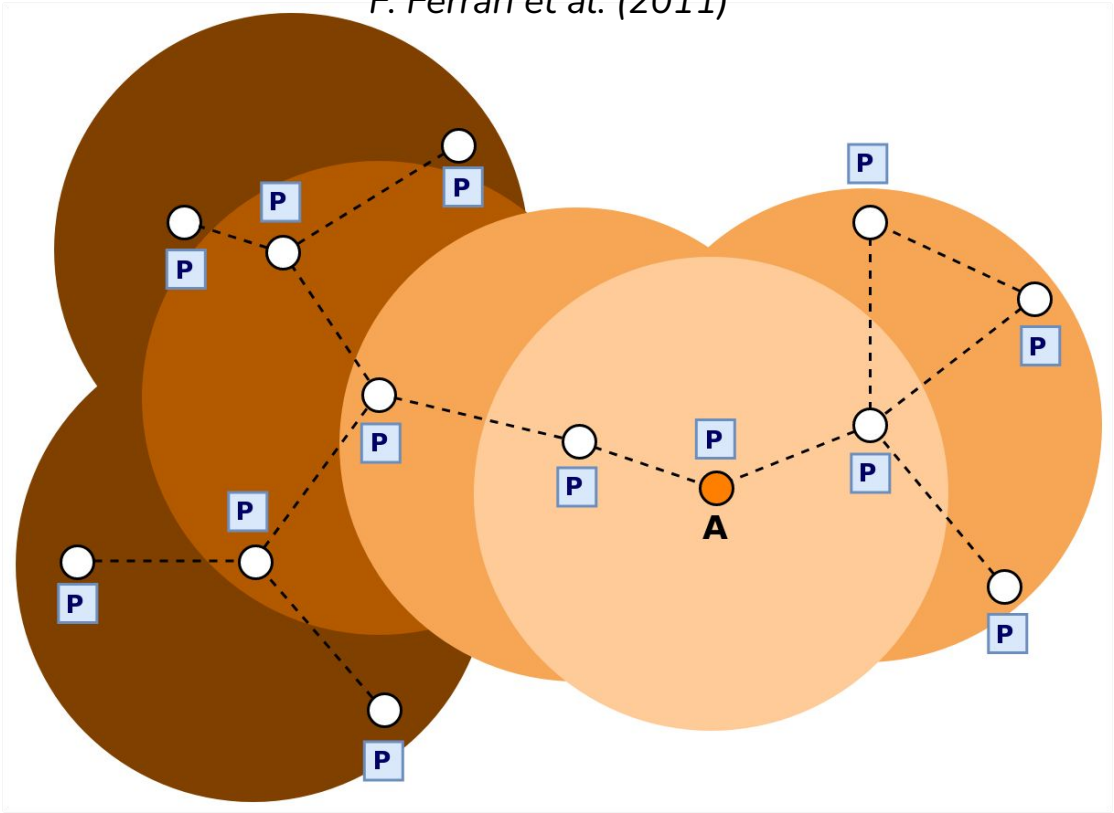

Glossy

F. Ferrari et al. (2011)

High latency for all-to-all communication

High energy cost for dissemination of individual payloads

Chaos

O. Landsiedel et al. (2013)

Many-to-all data **dissemination**

All **nodes contribute to the payload** being disseminated

Power and **time-efficient** flooding

Chaos

O. Landsiedel et al. (2013)

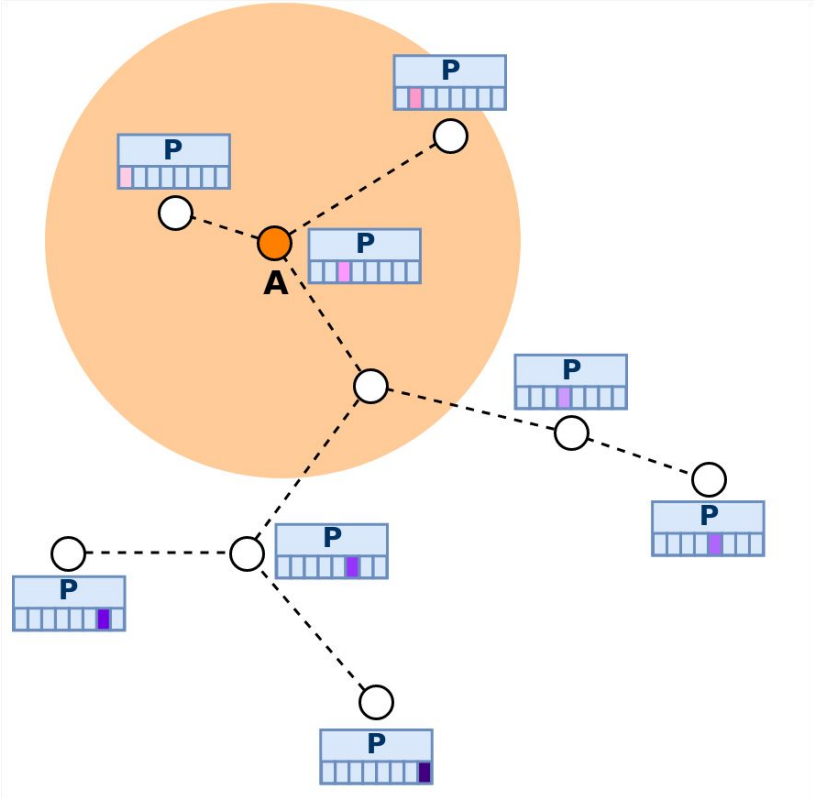

Chaos

O. Landsiedel et al. (2013)

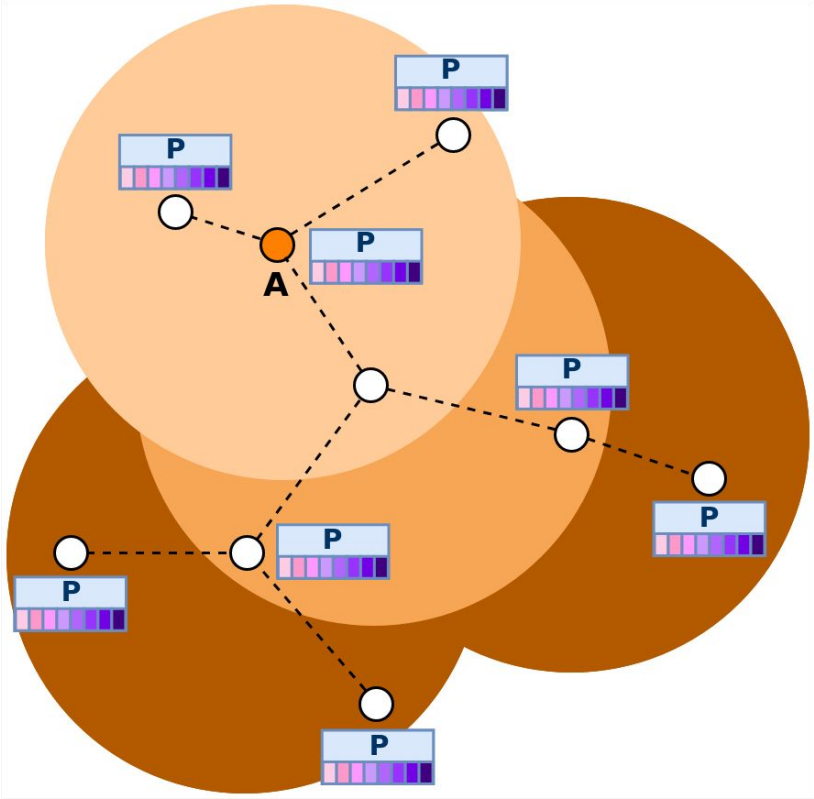

Chaos

O. Landsiedel et al. (2013)

Unreliable for one-to-all communication

Unknown termination time

Voting Protocols

A **coordinator proposes a value** to the network

A network of nodes (**cohort**) **votes** on the proposed value

Nodes conditionally commit

Atomic Commit Protocols:

2PC & 3PC

Consensus Protocols:

Paxos

Two Phase Commit

J. Gray (1978)

All nodes will **commit** the **same value**

Blocking protocol: not guaranteed to terminate

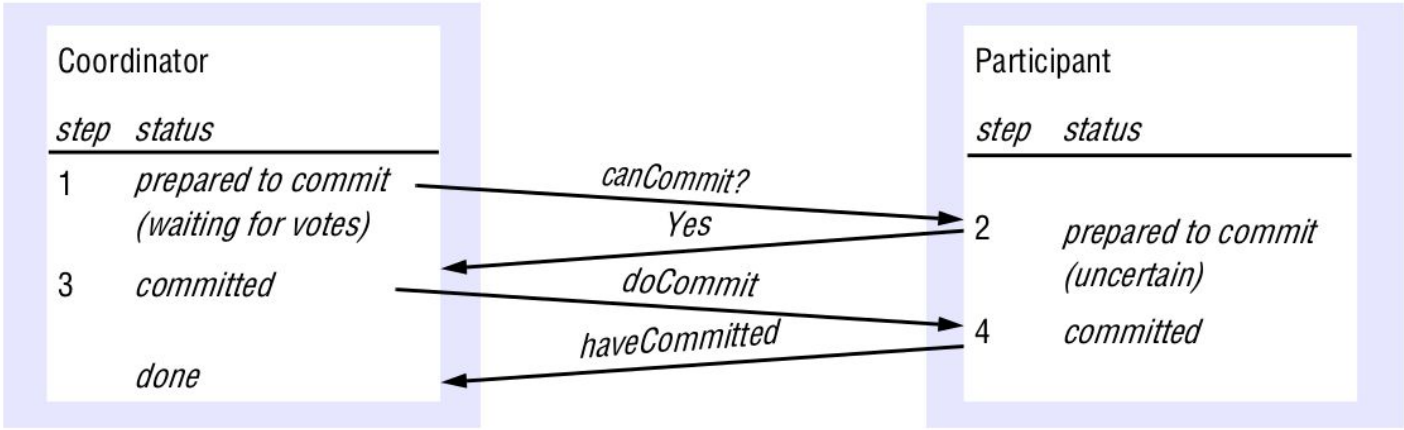

Image from: G. Coulouris et al. Distributed Systems: Concepts and Design (2011)

Three Phase Commit

D. Skeen (1981)

Network might be **inconsistent** (**safety property**)

Non-blocking: guaranteed termination (**liveliness property**)

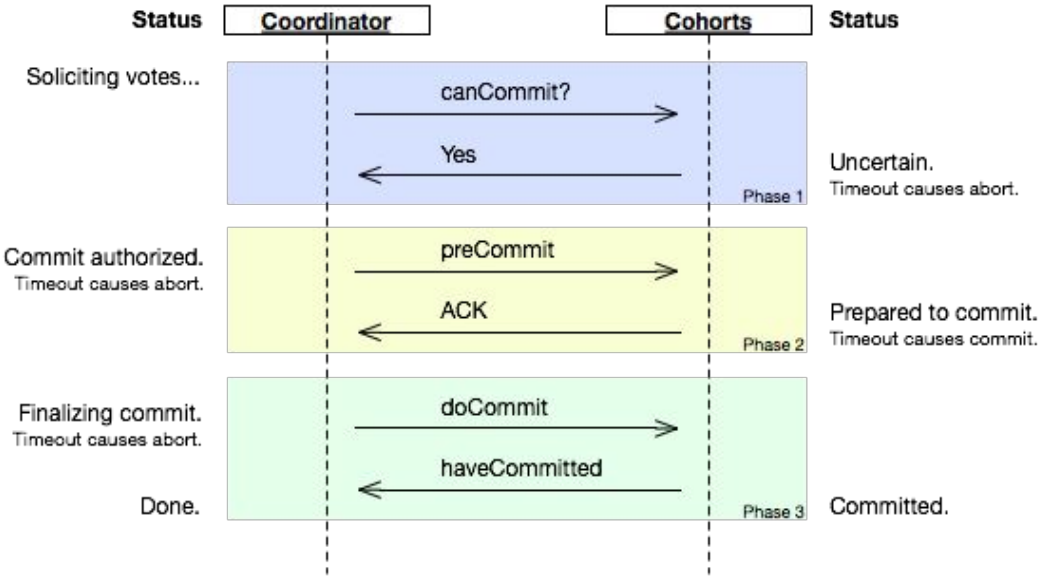

Image from: https://en.wikipedia.org/wiki/Three-phase_commit_protocol#/media/File:Three-phase_commit_diagram.png

Paxos

L. Lamport (1998)

Values are **proposed to acceptors**

Only a **quorum** of acceptors **must commit**

Committed **values** are **sent to learners**

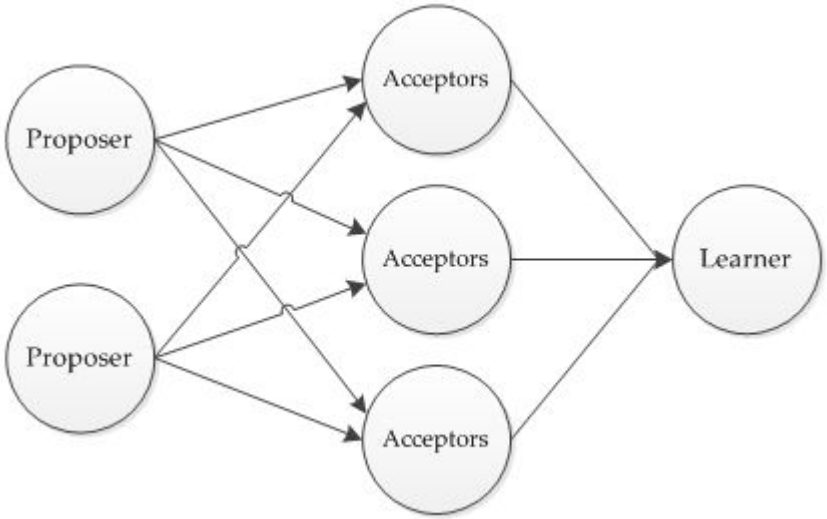

Image from: https://angus.nyc/2012/paxos-by-example/

Voting Protocols

WSNs are **distributed asynchronous systems**

Common for broadcast **packets** to be **lost**

Nodes may become **unreachable with interference**

A^2: Agreement in the Air

B. Al Nahas et al. (2017) and V. Poirot et al. (2019)

Uses **Chaos ST primitive**

Implements **2PC and 3PC** protocols (2017)

Implements **consensus with WPaxos** (2019)

A^2: Agreement in the Air

B. Al Nahas et al. (2017) and V. Poirot et al. (2019)

Chaos ST primitive is **unreliable**

No termination guarantees

Greatly **disrupted by interference**

Baloo

R. Jacob et al. (2019)

Proposes a **middleware** to **interact with radio** and hardware

Can be used with **multiple ST primitives**

Protocols implemented using **callback functions**

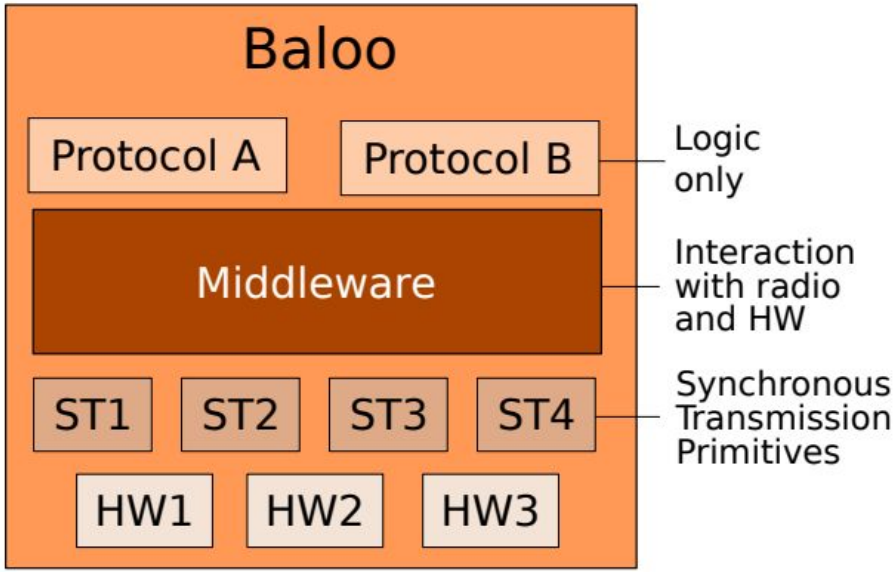

Baloo

R. Jacob et al. (2019)

Hard to use same protocol with **more than one ST primitive**

Different **primitives** must be **scheduled differently**

Protocols must be in charge of **control and timing**

Our Aim

C1. Protocols must be able to **easily switch between ST primitives**

C2. Create a **new ST primitive** which is able to provide the **robustness of Glossy** together with the **performance of Chaos**

C3. **Consensus** protocols must run **reliably on WSNs**

C4. Protocol reliability has to be **tested with replicable results**

Contribution I

XPC: A Voting Protocol Coordinator

Baloo requires a coordinator for voting protocols

Handles flood **timing and** primitive **control**

Has **intermediate data representation** for primitives

Addresses C1

Contribution I

XPC: A Voting Protocol Coordinator

Baloo requires a coordinator for voting protocols

Handles flood **timing and** primitive **control**

Has **intermediate data representation** for primitives

Addresses C1

“Protocols must be able to **easily switch between ST primitives**”

XPC: A Voting Protocol Coordinator

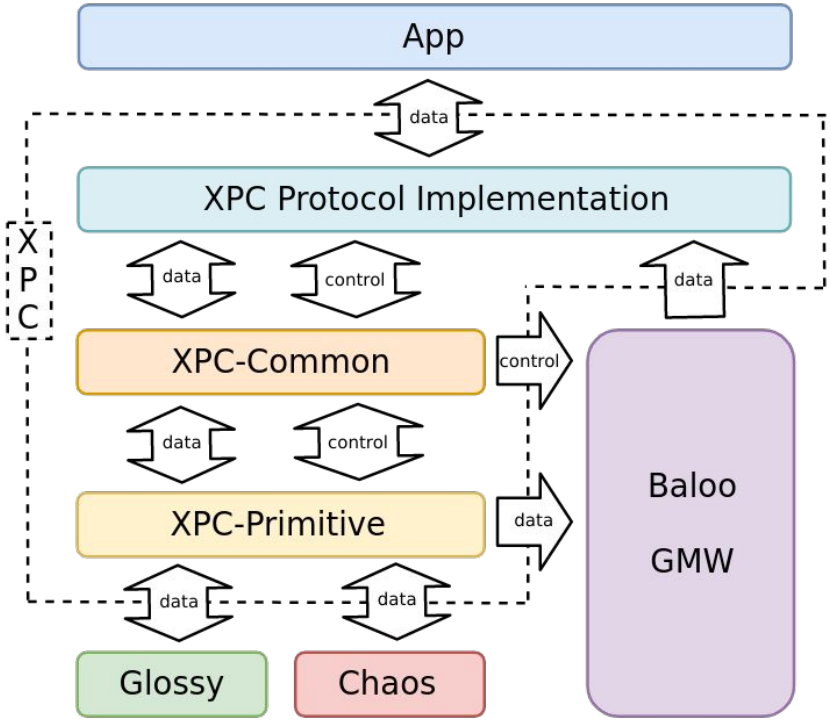

XPC Additions

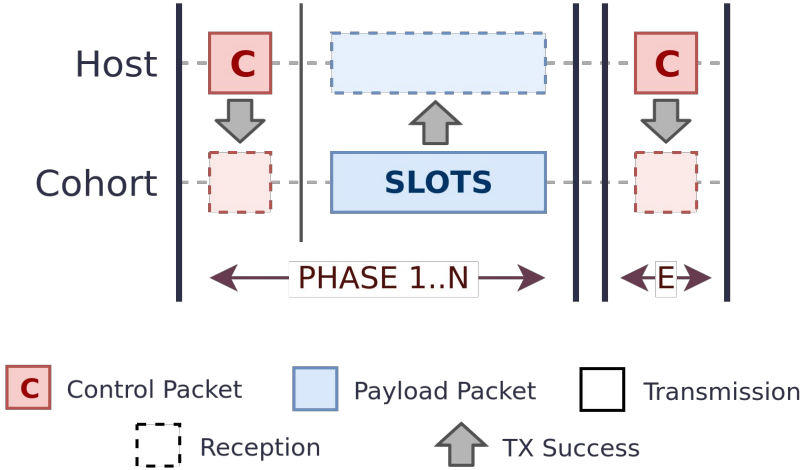

Single Initiator

Retransmissions for Reliability

Additional **Final Round** (E)

Two Phase Commit with XPC

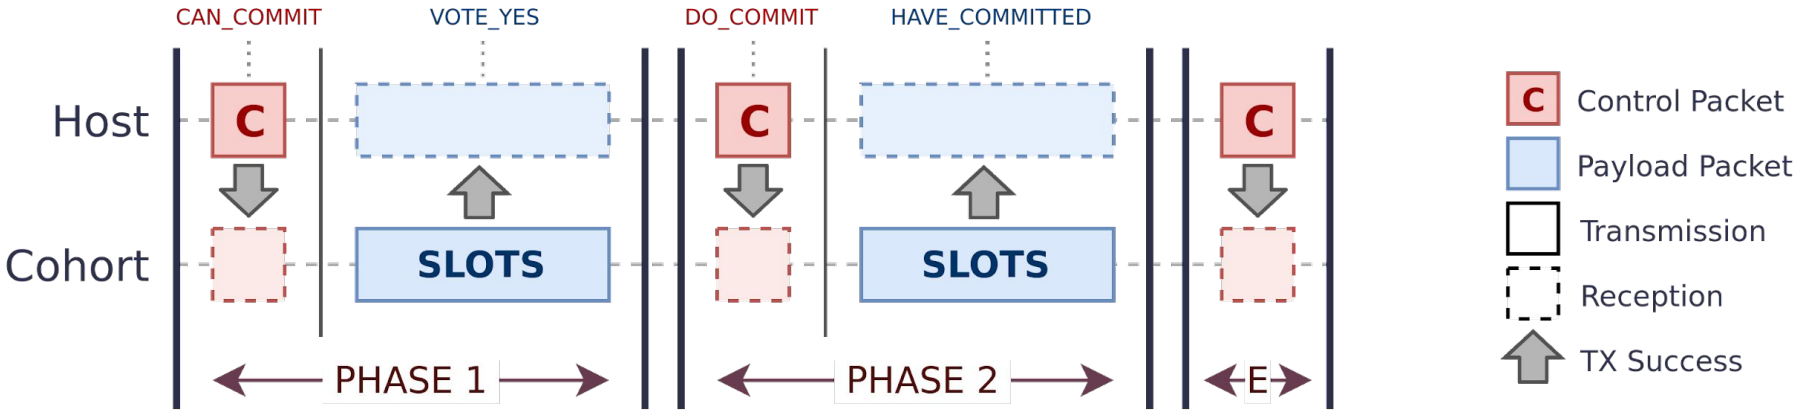

2PC-Glossy

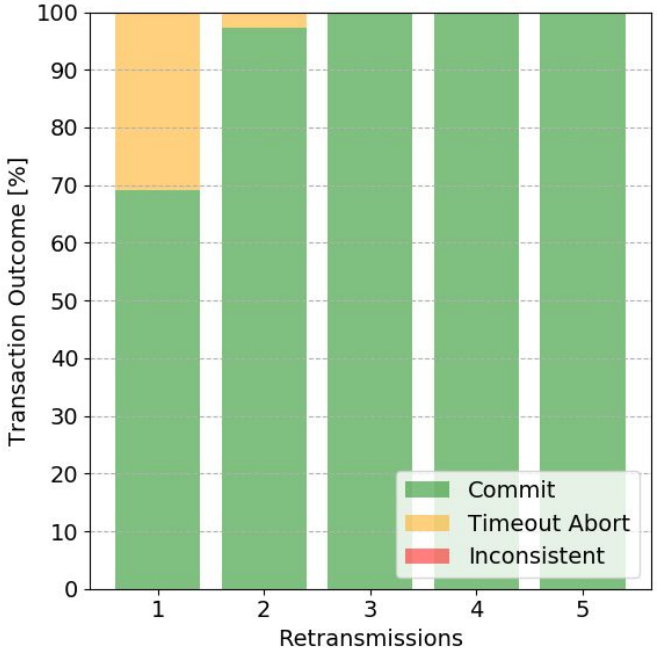
Transaction Outcome [%]
100
90
80
70
60
50
40
30
20
10
0
1
2
3
4
5
Retransmissions
Commit
Timeout Abort
Inconsistent

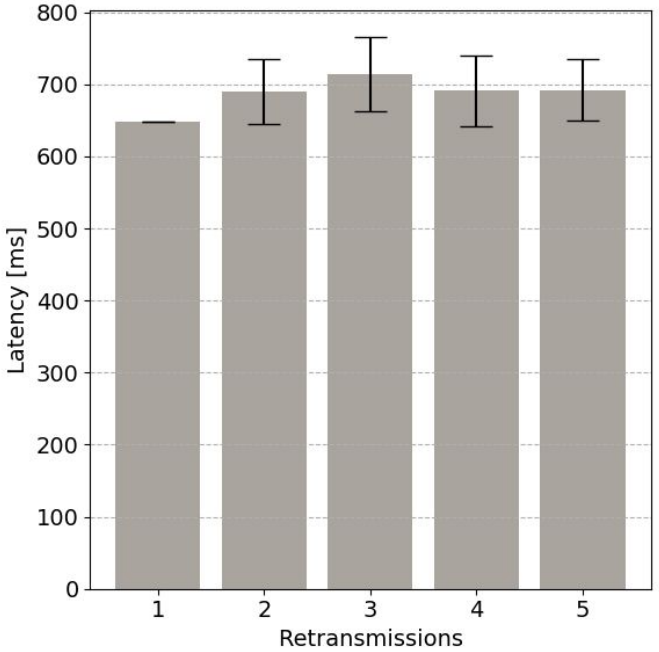
Latency [ms]
800
700
600
500
400
300
200
100
0
1
2
3
4
5
Retransmissions

Contribution I

2PC-Chaos

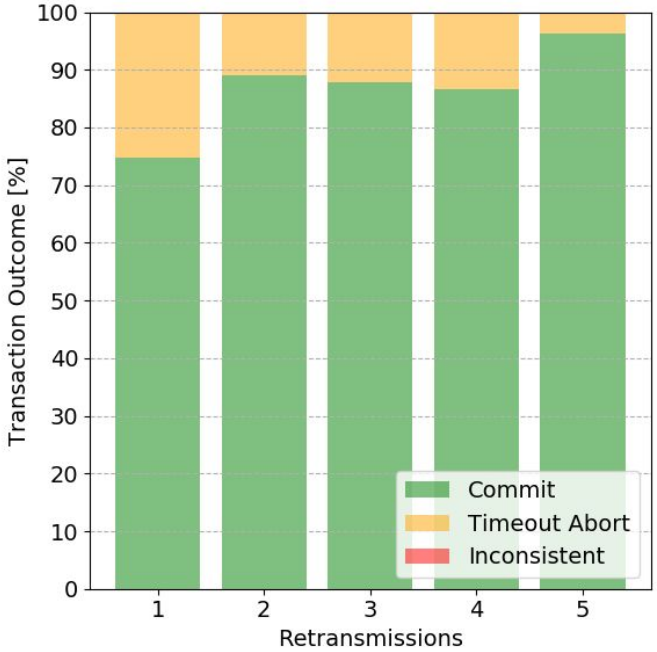

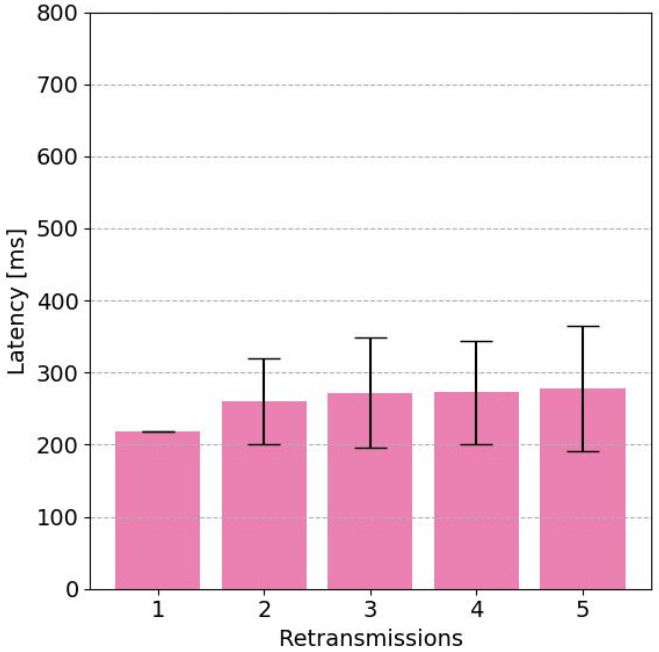

Contribution I

Hybrid ST Primitive

Leverages **XPC** to schedule **first** a **Chaos** flood **and then Glossy** rounds

Minimises latency and **maximises reliability**

Addresses C2

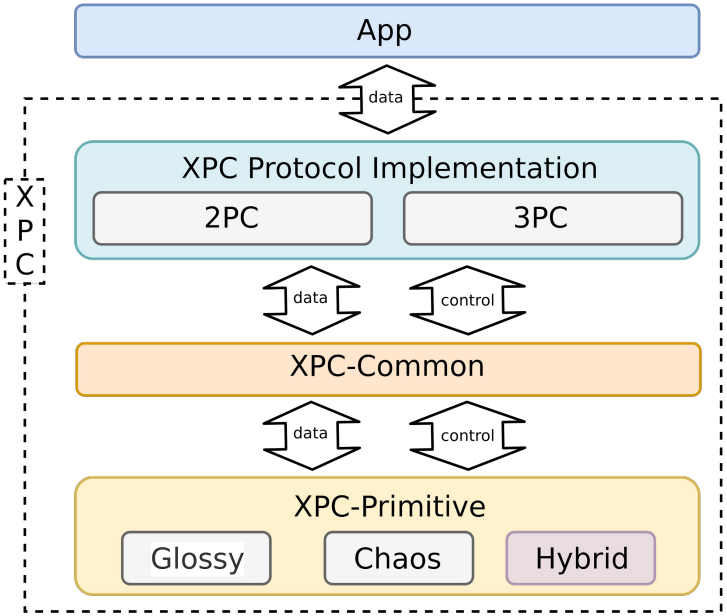

Contribution I

Hybrid ST Primitive

Leverages **XPC** to schedule **first** a **Chaos** flood **and then Glossy** rounds

Minimises latency and **maximises reliability**

Addresses C2

“Create a **new ST primitive** which is able to provide the **robustness of Glossy** together with the **performance of Chaos**”

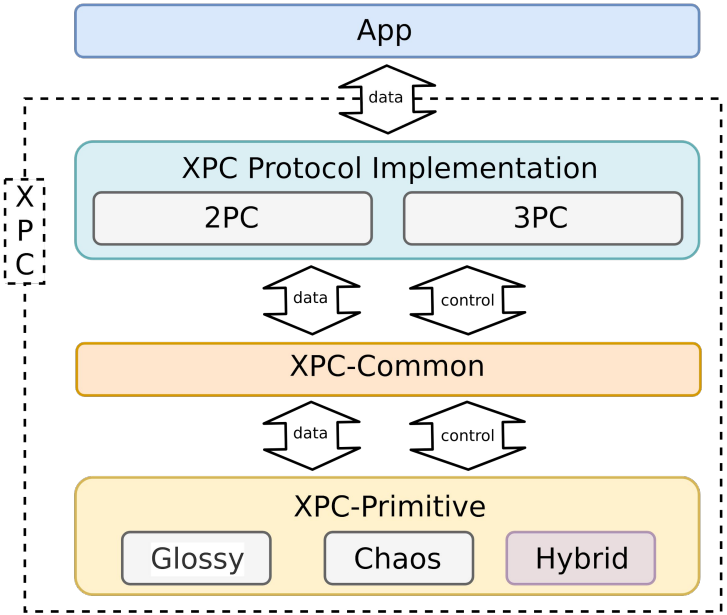

Contribution I

Hybrid ST Primitive

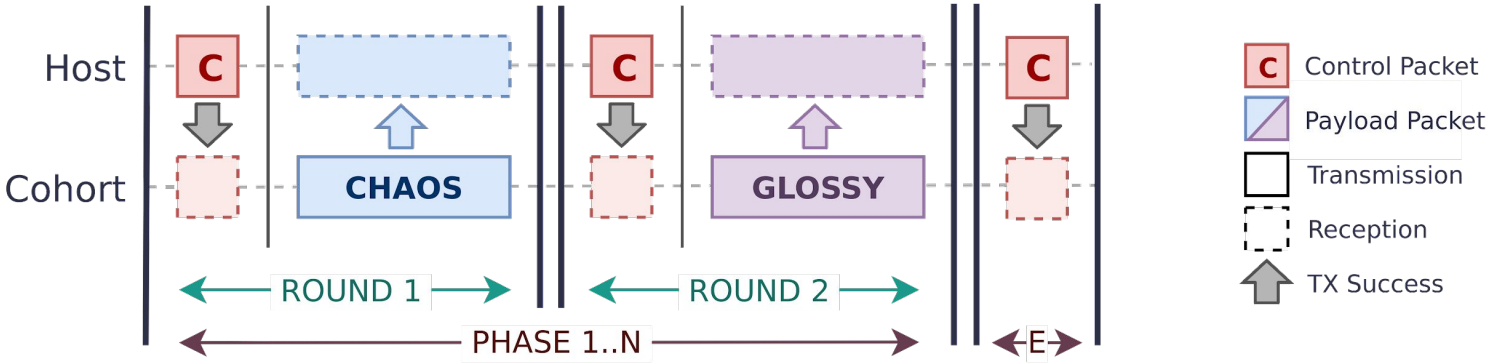

Demo Interface: Flocklab Visualiser

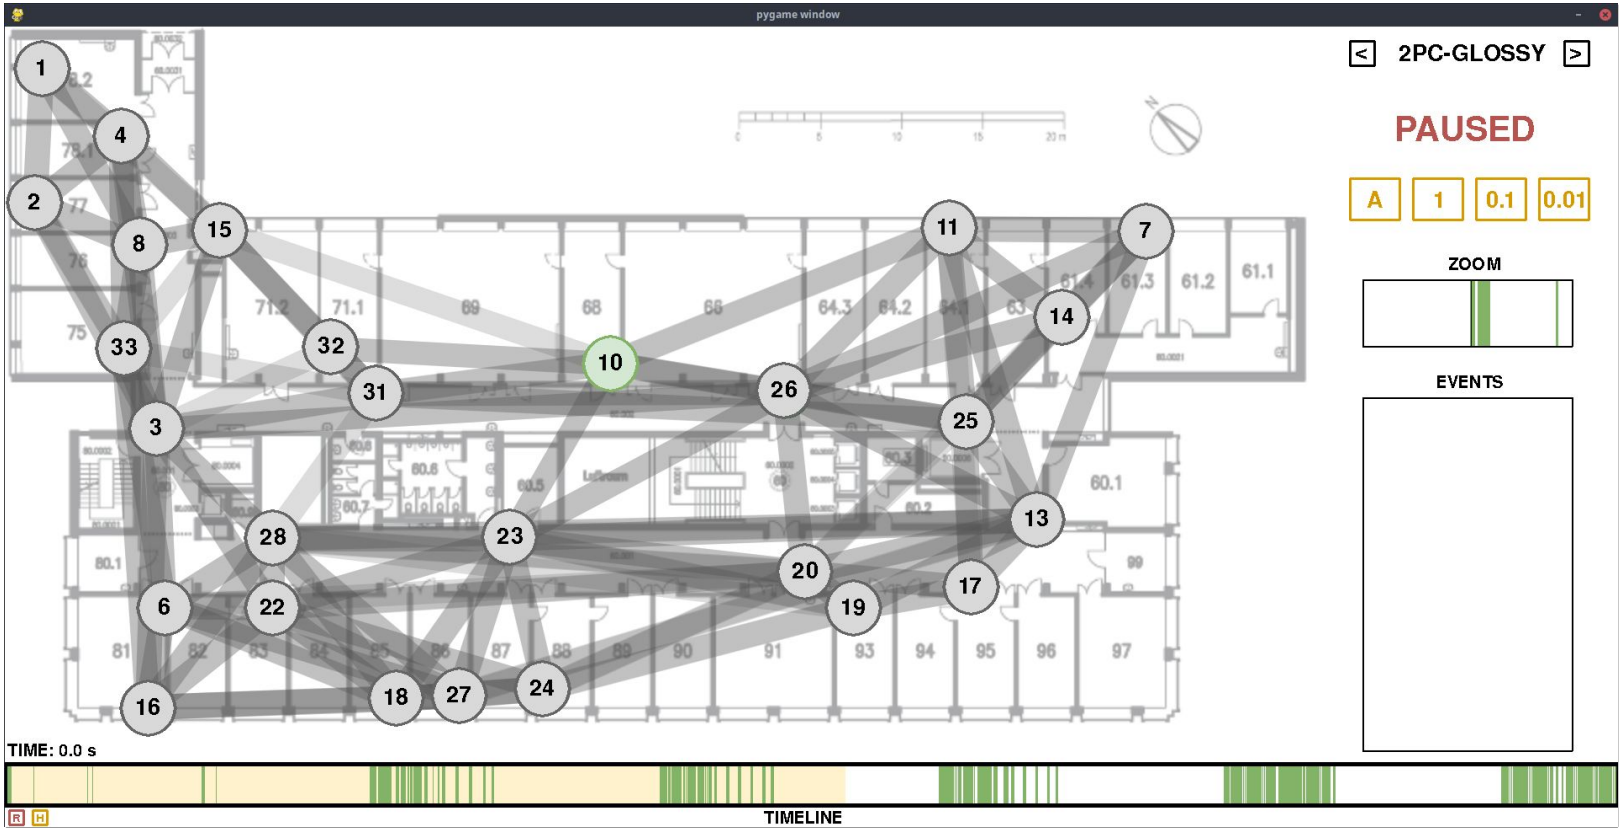


2PC-HYBRID
SPEED: x0.002
A
1
0.1
0.01
ZOOM
EVENTS
P1 (10): Control
P1 (10): Chaos
P1 (10): Control
TIME: 20.4096 s
TIMELINE

Contribution I

2PC-Hybrid

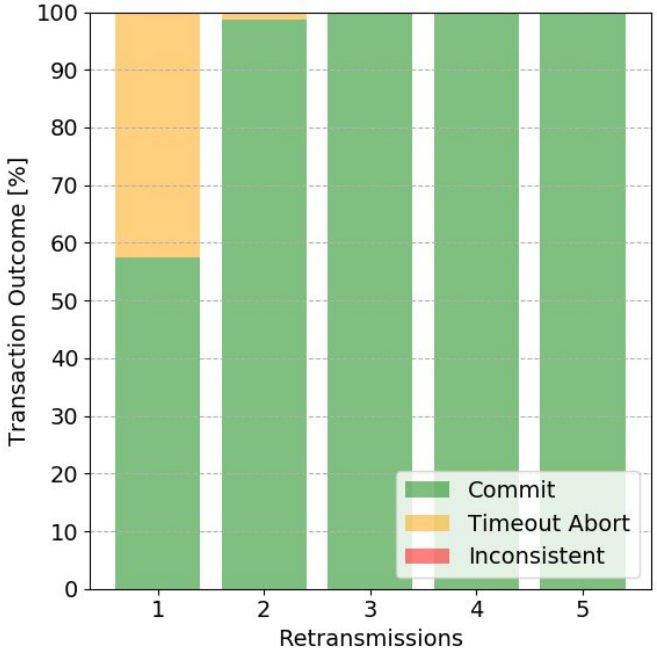

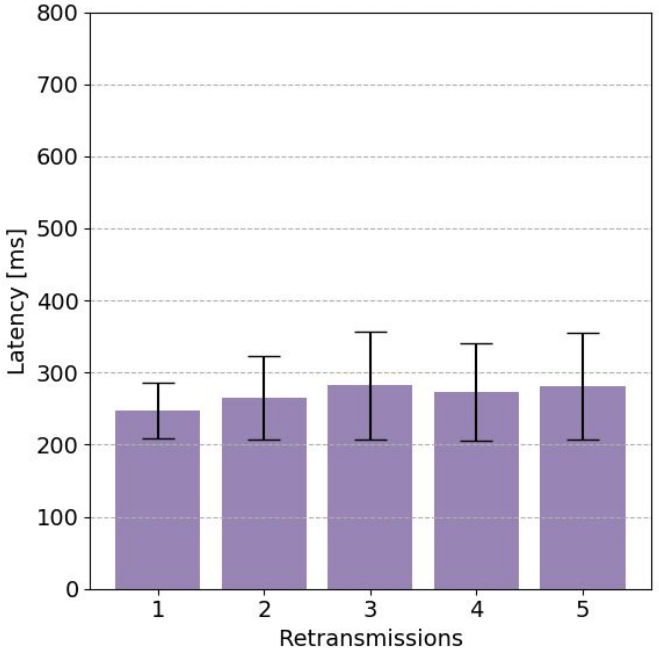

Comparison with A^2

Transaction **Commit**

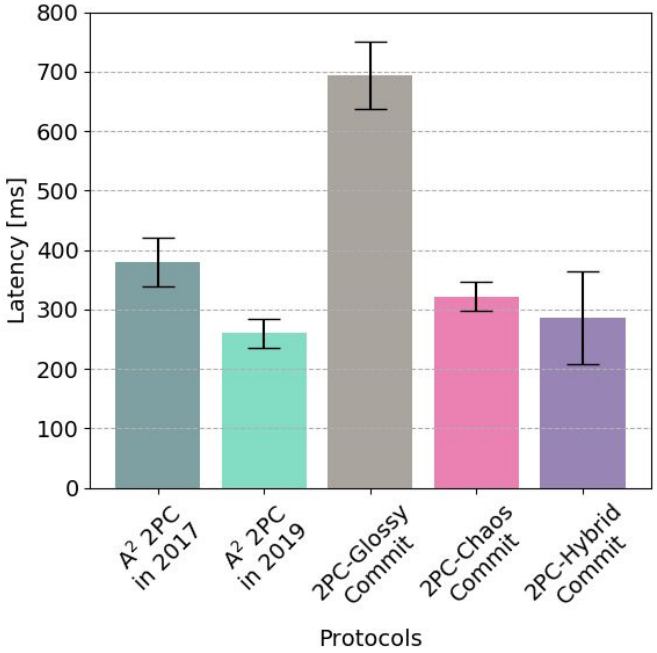

Transaction **Abort**

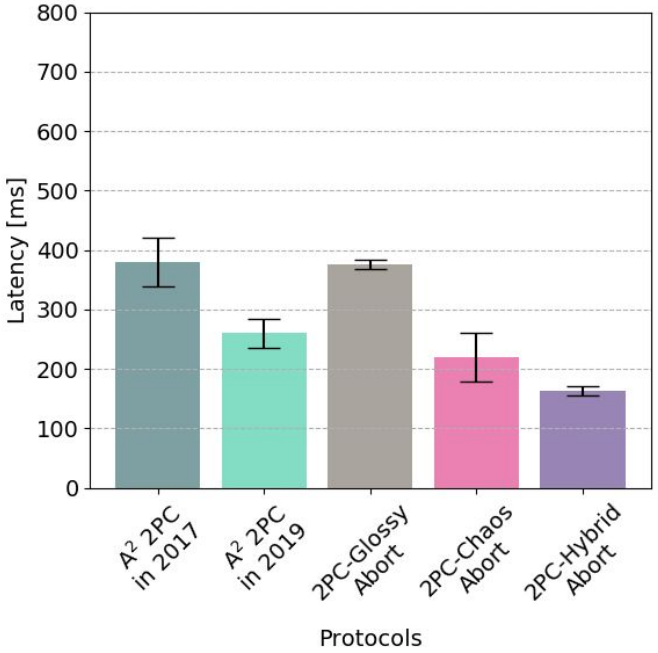

Contribution II

WiPP: Wireless Part-time Parliament

Quorum-based **majority voting**

Global dissemination of committed values

Available as an **XPC Protocol**

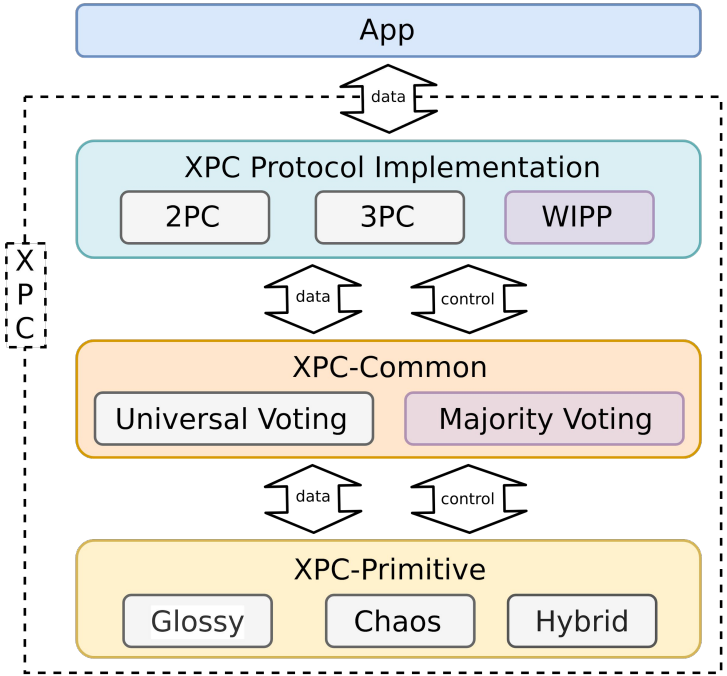

Contribution II

WISP: WiPP Simple Paxos

WiPP + **Hybrid** ST primitive

Voting phase **and** global **dissemination**

Satisfies **consensus properties**:

Validity Integrity Termination Agreement

Addresses C3

Contribution II

WISP: WiPP Simple Paxos

WiPP + **Hybrid** ST primitive

Voting phase **and** global **dissemination**

Satisfies **consensus properties**:

Validity Integrity Termination Agreement

Addresses C3

"**Consensus** protocols must run **reliably on WSNs**"

WISP: WiPP Simple Paxos

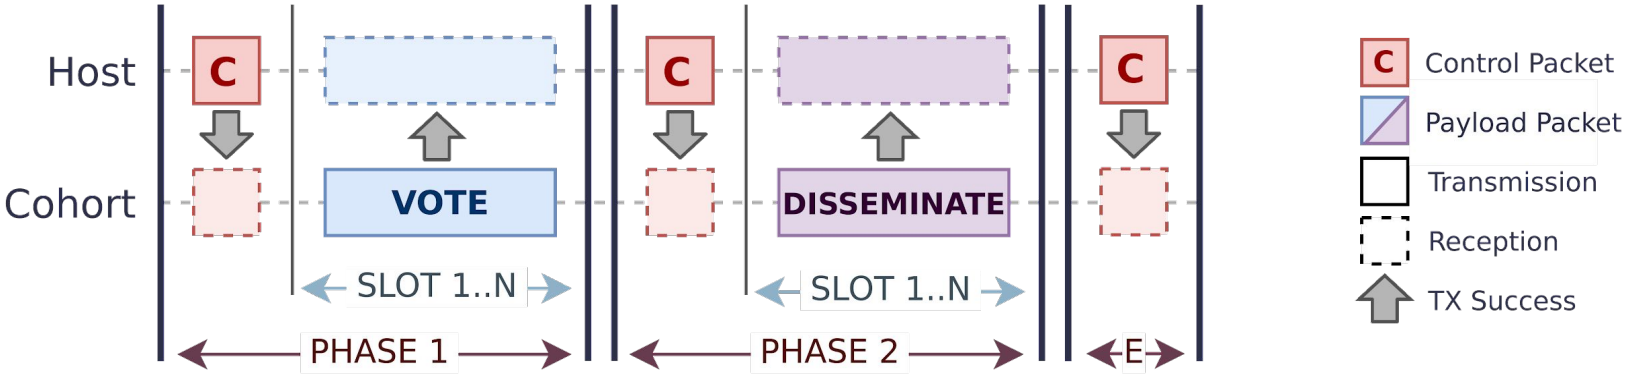


WISP
SPEED: x0.005
A
1
0.1
0.01
ZOOM
EVENTS
P1 (10): Control
P1 (10): Chaos
P2 (10): Control
P2 (10): Chaos
TIMELINE

Contribution II

WISP Applications

Configuration management

Leader election

Local node **clustering**

Failure-free commits

Contribution III

Multiple Proposers

Allows **proposals from any node**

Proposal use **contention slots**

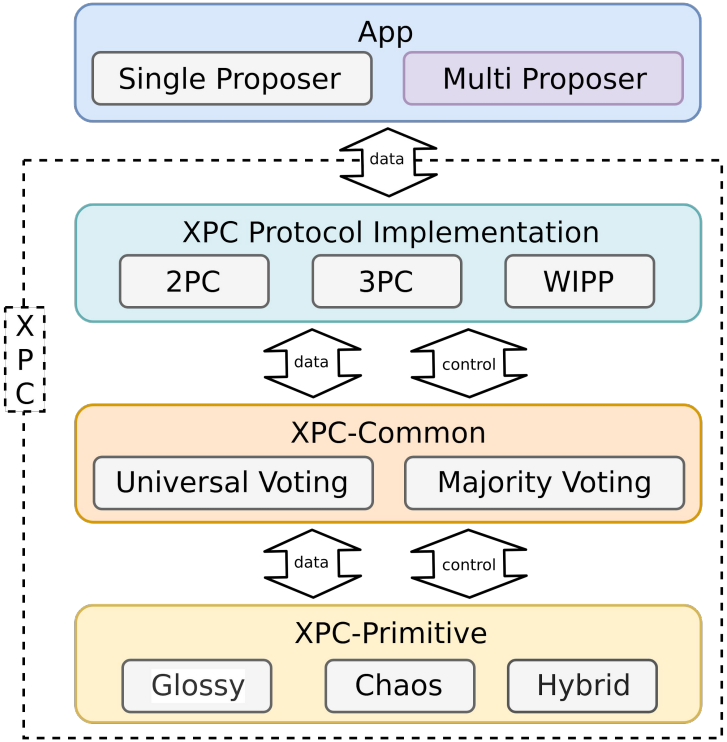

Contribution III

WIMP: WiPP Multi Paxos

Extension of WISP

Allows **proposals** from **any network node**

Uses **global leader**

Contribution III

WIMP: WiPP Multi Paxos

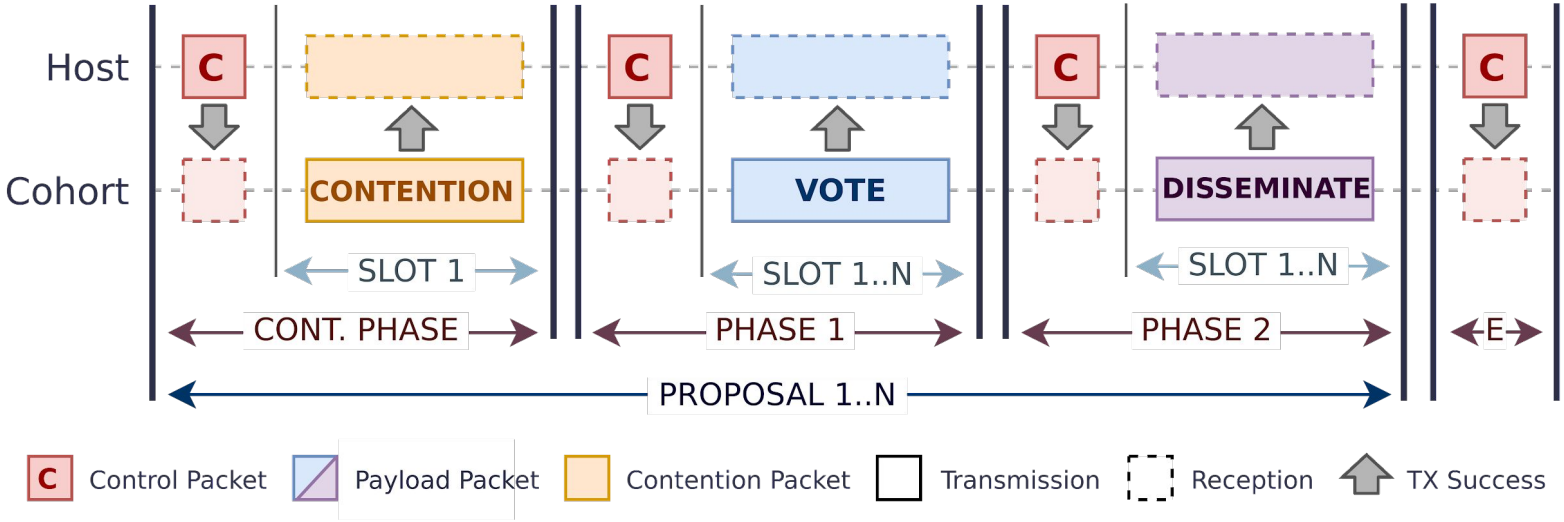


WIMP
SPEED: x0.005
A
1
0.1
0.01
ZOOM
EVENTS
P0 (10): Control
P0 (25): Content
P1 (10): Control
P1 (10): Chaos
P2 (10): Control
P2 (10): Chaos
TIMELINE

Comparison with WPaxos (A^2)

Transaction **Commit**

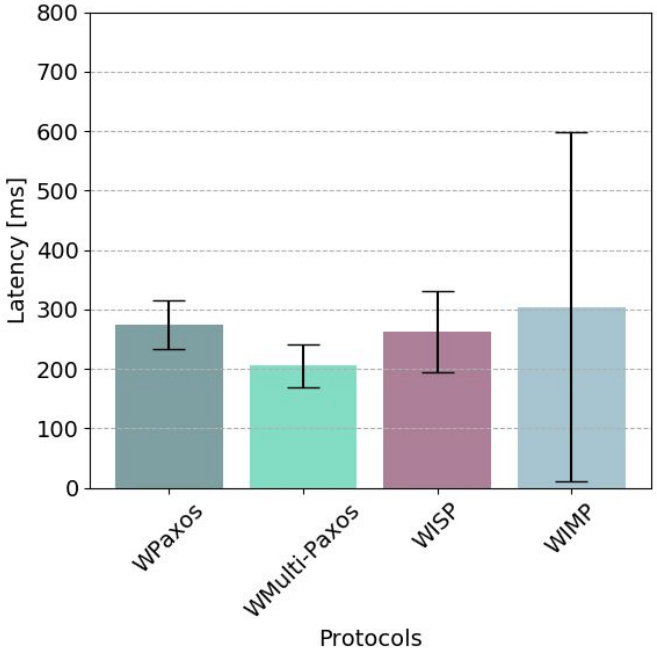

Transaction **Abort**

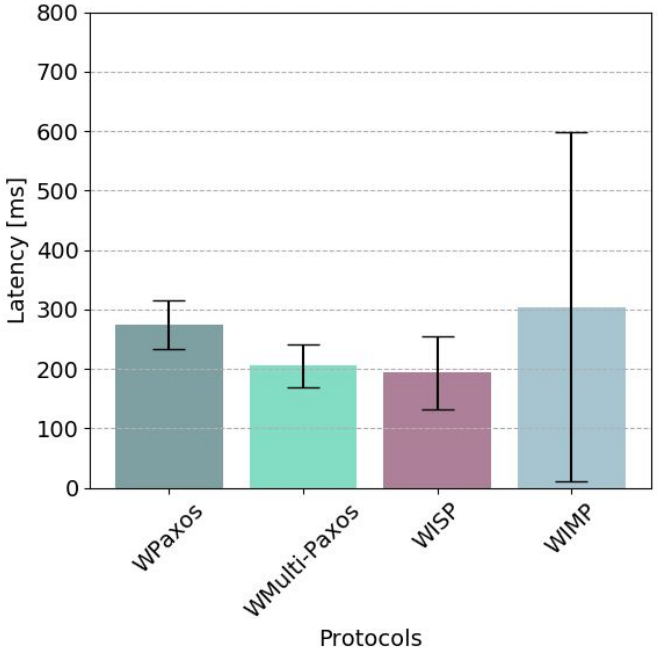

JamLab: Analysis with Interference

C. A. Boano et al. (2011)

To test **reliability** we **inject interference**

We use **JamLab**

Our protocols are **100% reliable with 1 interfering node**

Addresses C4

JamLab: Analysis with Interference

C. A. Boano et al. (2011)

To test **reliability** we **inject interference**

We use **JamLab**

Our protocols are **100% reliable with 1 interfering node**

Addresses C4

“Protocol reliability has to be **tested with replicable results**”


WISP INTERF
SPEED: x0.013
A
1
0.1
0.01
ZOOM
EVENTS
P1 (16): Glossy
P1 (17): Glossy
P1 (18): Glossy
P1 (19): Glossy
P1 (22): Glossy
P1 (23): Glossy
P1 (24): Glossy
P1 (10): Control
P1 (15): Glossy
P1 (17): Glossy
TIMELINE

Multiple Interfering Nodes

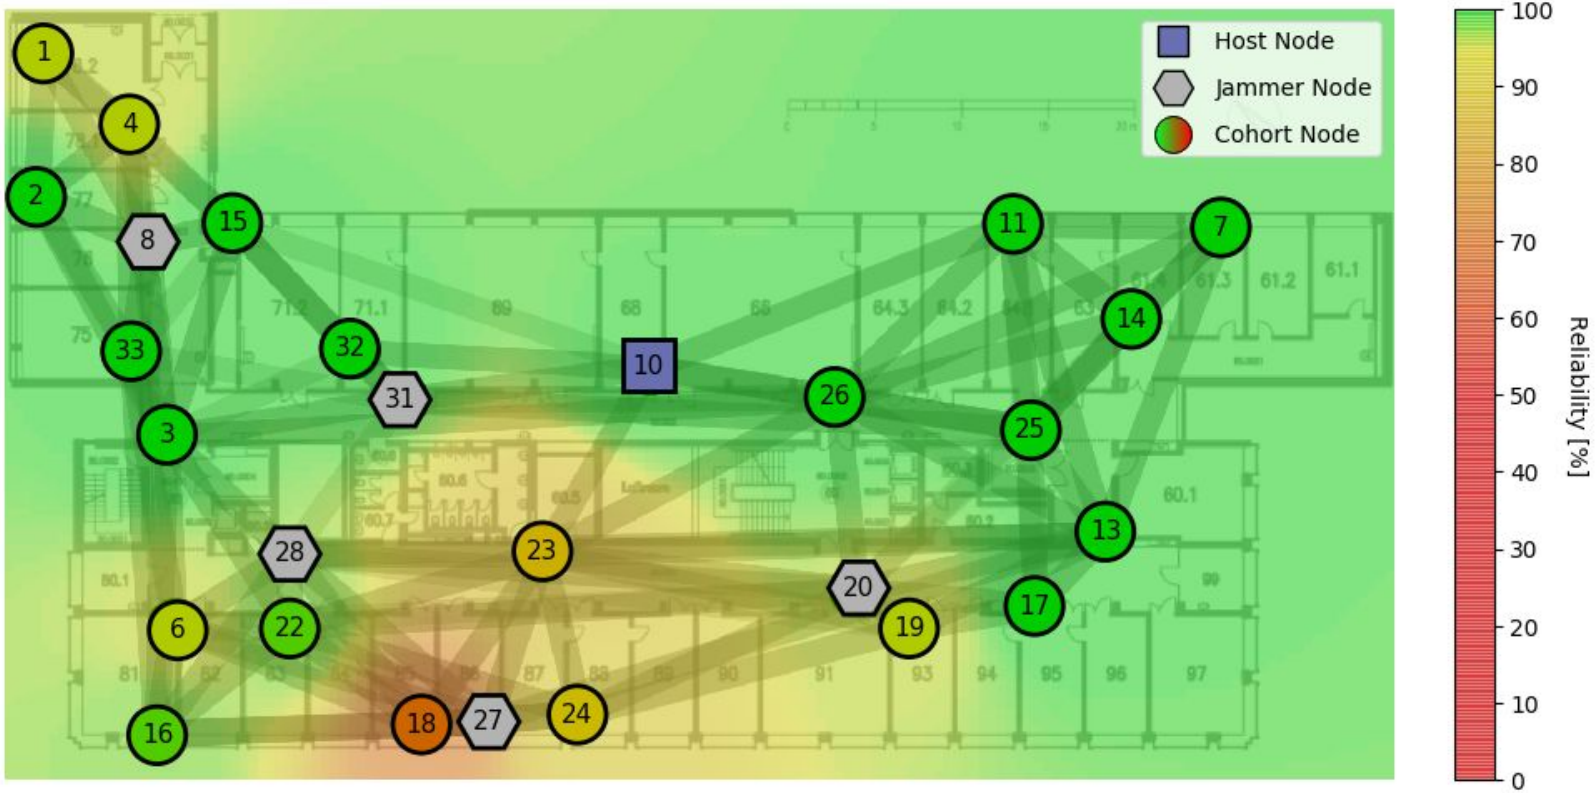

Our Contributions

Hybrid
(and XPC)

WISP
(and WiPP)

WIMP
(and Multi-Proposer)

Future Work

Group membership

XPC protocols **as a service**

Submit for **publication to IPSN**

Q & A

Additional

Security Concerns

ST primitives are **not secure**

A lot of research in WSN Security:

C. Chu et al. (**2010**). Practical ID-based Encryption for Wireless Sensor Network.

R. Gustavo et al. (**2012**). Asymmetric Encryption in Wireless Sensor Networks.

D. Shubhangi et al. (**2015**). Security in Wireless Sensor Network Using Cryptographic Techniques.

M. Elhoseny et al. (**2016**). An energy efficient encryption method for secure dynamic WSN. Security and Communication Networks.

K. C. Hewage et al. (**2017**). Protecting Glossy-Based Wireless Networks from Packet Injection Attacks.

K. Tsai et al. (**2018**). A Light Weight Data Encryption Method for WSN Communication.

Additional

Multi-Node Interference

Termination

Eventually each correct process decides a value

Agreement

All correct processes decide on the same value

Integrity

A process decides at most on one value

Validity

If a process decides on a value, then it must have been proposed by some process

Additional

Multiple-Proposers in Action

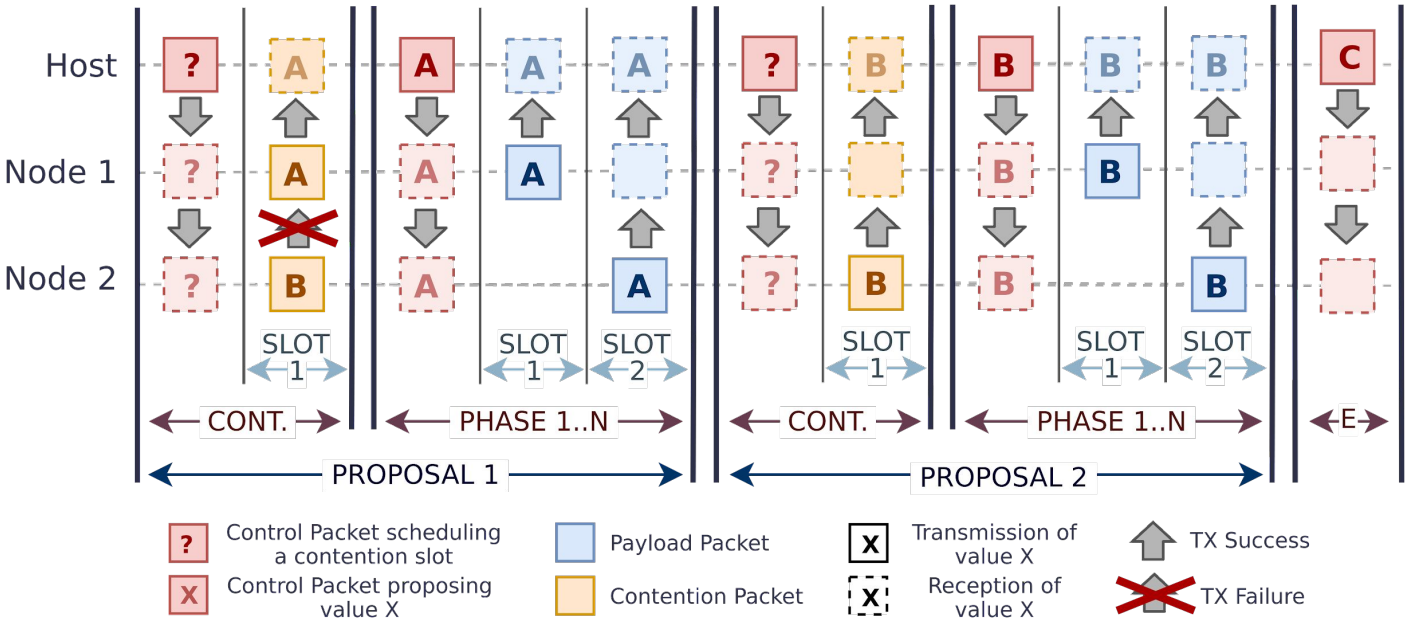

Additional

Safety in Multiple-Proposers

PPB: Pending Proposal Bit

Used when Nodes wish to propose

PVB: Proposed Value Bit

Used for proposal acknowledgement

Additional

2PC-Chaos Overview

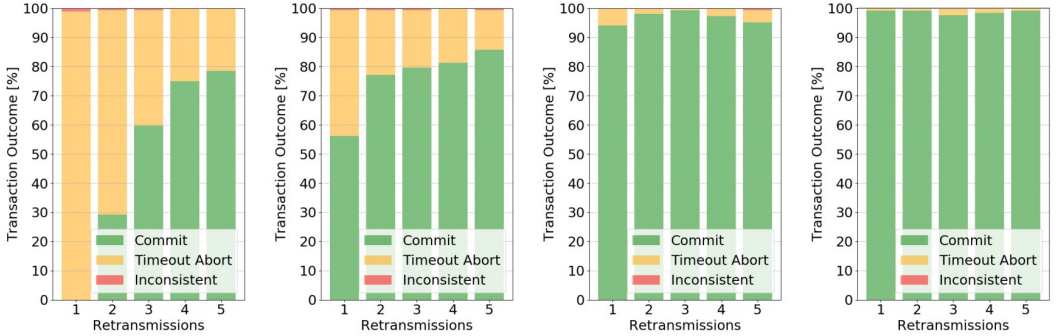

(a) 25ms slots (b) 50ms slots (c) 100ms slots (d) 200ms slots

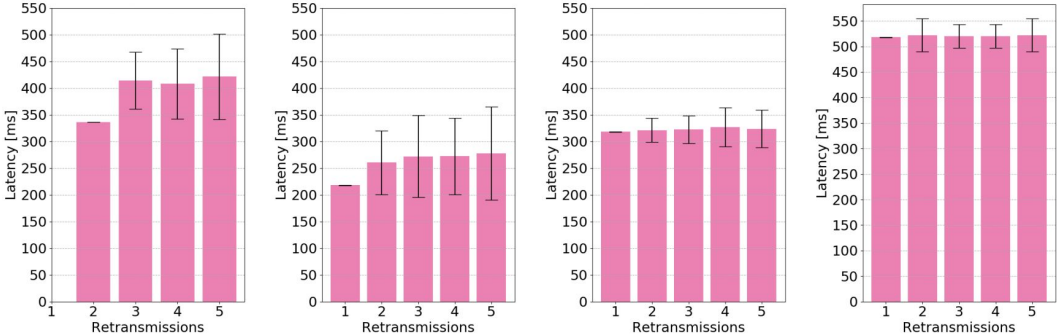

(a) 25ms slots (b) 50ms slots (c) 100ms slots (d) 200ms slots

Additional

Majority Voting Delta-Q

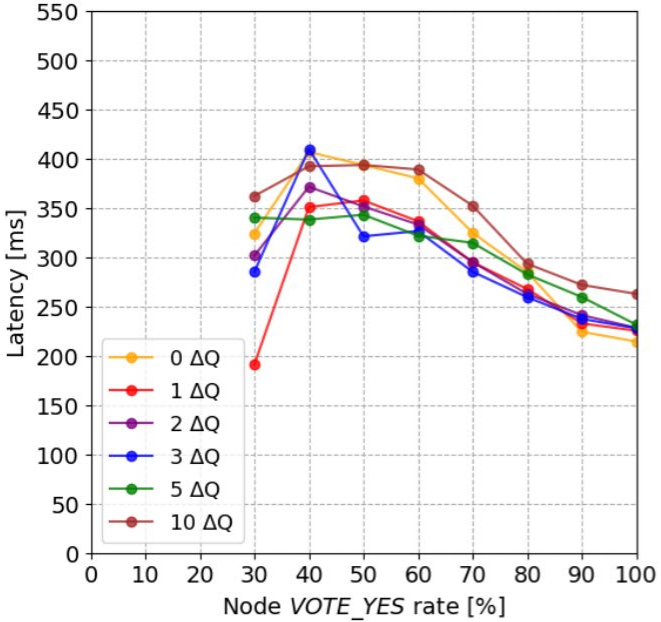

(a) Average commit latency

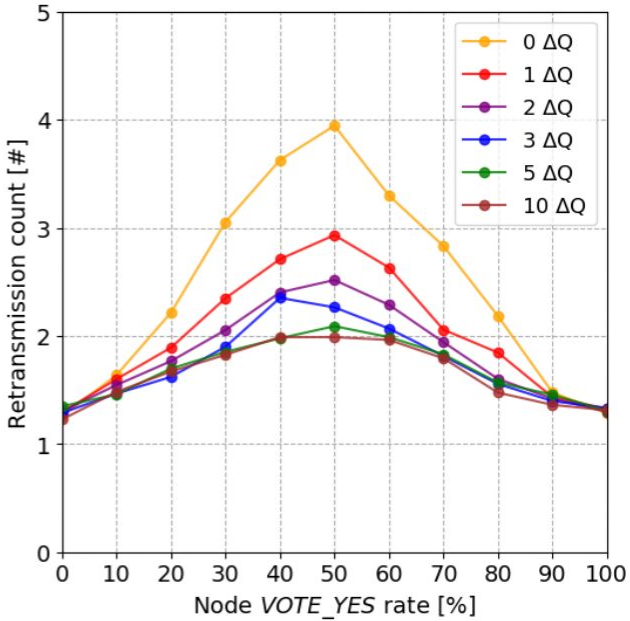

(b) Average number of retransmissions

Additional

Majority Voting M-Slots

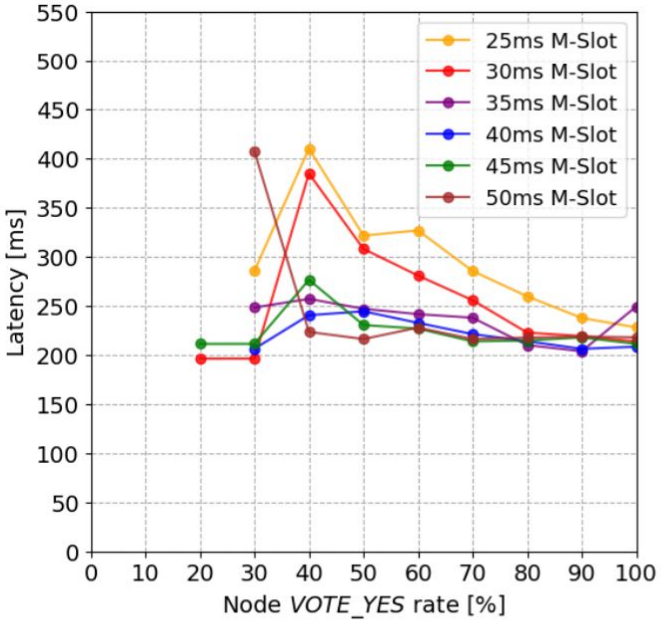

(a) Average commit latency

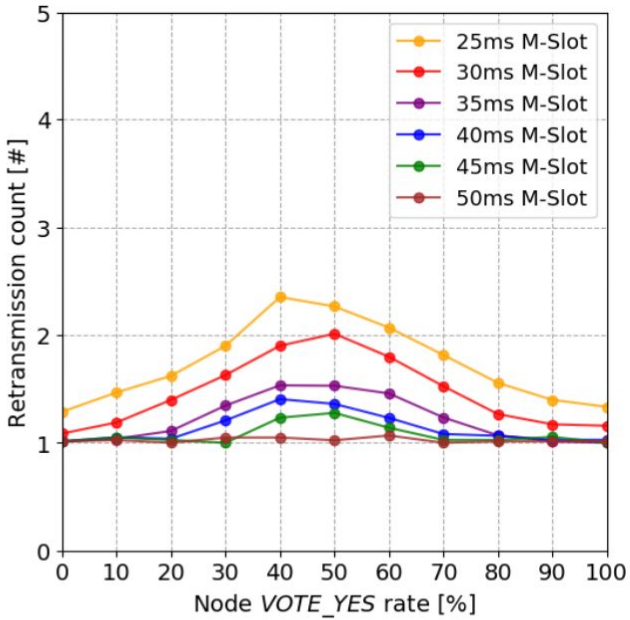

(b) Average number of retransmissions

Additional

Interference Overview

Wifi Interference	Reliability (%)	Latency (ms)	Chaos Coverage (%)	Avg. Retr.
2PC-Hybrid	100.00	383.07	**P1**: 92.19 **P2**: 90.54	**P1**: 2.29 **P2**: 2.19
3PC-Hybrid	100.00	606.20	**P1**: 91.32 **P2**: 92.92 **P3**: 89.58	**P1**: 2.26 **P2**: 1.96 **P3**: 2.31
WISP	100.00	270.93	**MP**: 87.79 **DP**: 91.98	**MP**: 1.05 **DP**: 1.44
WIMP	100.00	360.88	**MP**: 86.09 **DP**: 91.98	**MP**: 1.04 **DP**: 1.70
Microwave	**Reliability (%)**	**Latency (ms)**	**Chaos Coverage (%)**	**Avg. Retr.**
2PC-Hybrid	100.00	396.06	**P1**: 92.97 **P2**: 91.82	**P1**: 2.07 **P2**: 2.17
3PC-Hybrid	100.00	586.67	**P1**: 92.09 **P2**: 91.58 **P3**: 92.30	**P1**: 1.95 **P2**: 2.12 **P3**: 1.95
WISP	100.00	342.53	**MP**: 78.90 **DP**: 91.14	**MP**: 1.15 **DP**: 2.23
WIMP	100.00	377.31	**MP**: 86.22 **DP**: 91.69	**MP**: 1.01 **DP**: 1.66

Additional

Multi-Node Interference

Microwave (2 Nodes)	Reliability (%)	Latency (ms)	Chaos Coverage (%)	Avg. Retr.
2PC-Hybrid	100.00	1027.18	**P1**: 82.91 **P2**: 81.54	**P1**: 4.64 **P2**: 5.04
WISP	98.33	625.37	**MP**: 56.49 **DP**: 82.85	**MP**: 1.57 **DP**: 4.52
Microwave (3 Nodes)	**Reliability (%)**	**Latency (ms)**	**Chaos Coverage (%)**	**Avg. Retr.**
2PC-Hybrid	95.92	1025.86	**P1**: 80.26 **P2**: 81.43	**P1**: 4.95 **P2**: 5.11
WISP	92.86	756.73	**MP**: 47.00 **DP**: 81.37	**MP**: 1.97 **DP**: 5.22
Microwave (4 Nodes)	**Reliability (%)**	**Latency (ms)**	**Chaos Coverage (%)**	**Avg. Retr.**
2PC-Hybrid	57.14	1221.30	**P1**: 79.59 **P2**: 79.92	**P1**: 6.72 **P2**: 6.11
WISP	59.18	1024.20	**MP**: 42.06 **DP**: 78.83	**MP**: 2.27 **DP**: 6.91
Microwave (5 Nodes)	**Reliability (%)**	**Latency (ms)**	**Chaos Coverage (%)**	**Avg. Retr.**
2PC-Hybrid	47.22	1291.68	**P1**: 75.51 **P2**: 79.37	**P1**: 7.85 **P2**: 7.50
WISP	59.09	1178.84	**MP**: 41.67 **DP**: 73.88	**MP**: 2.83 **DP**: 7.67

Additional

XPC with Glossy

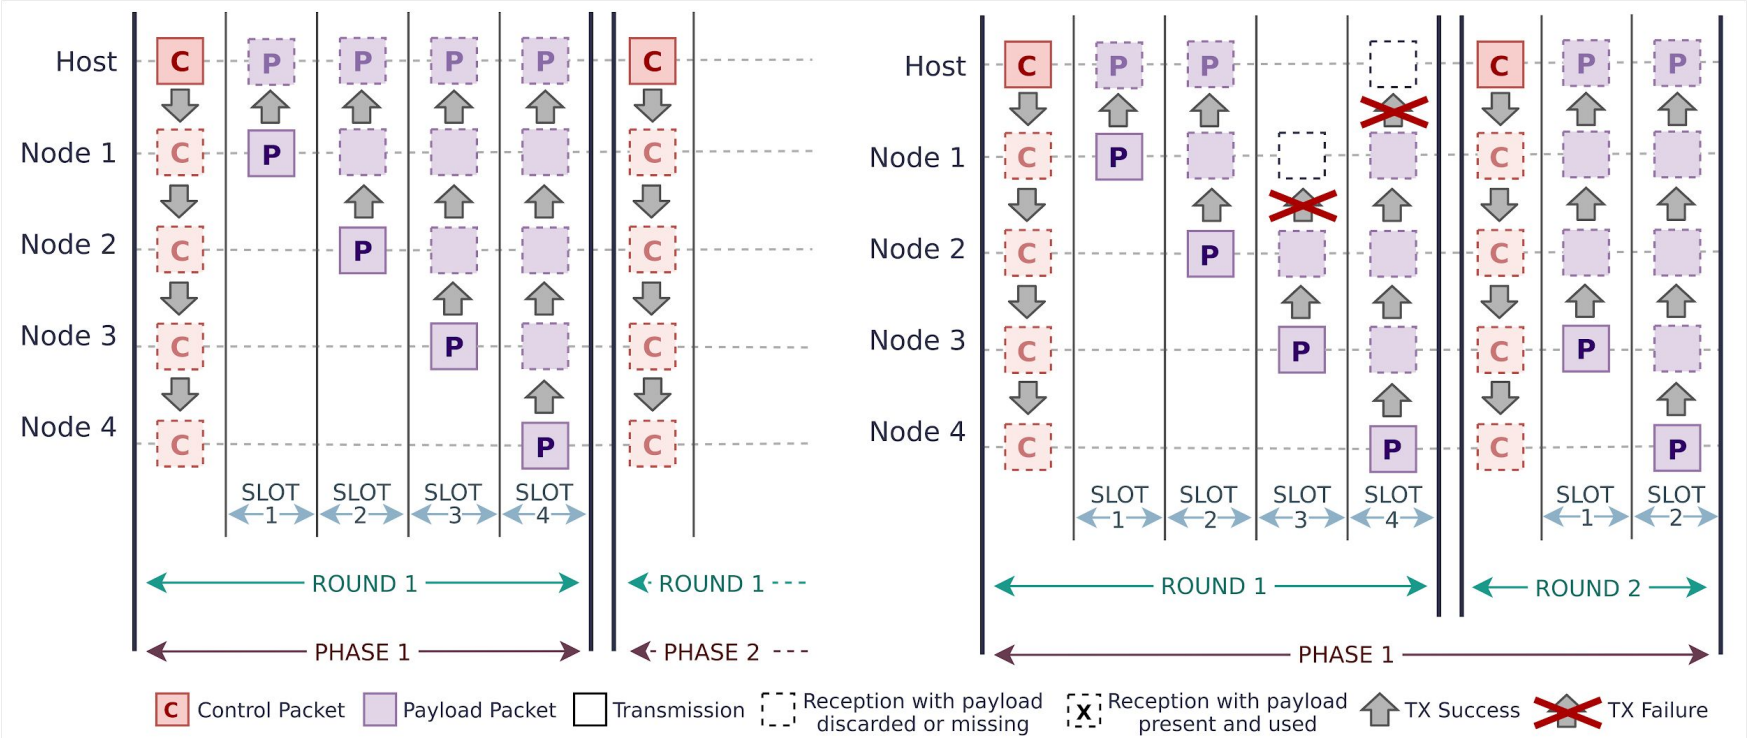

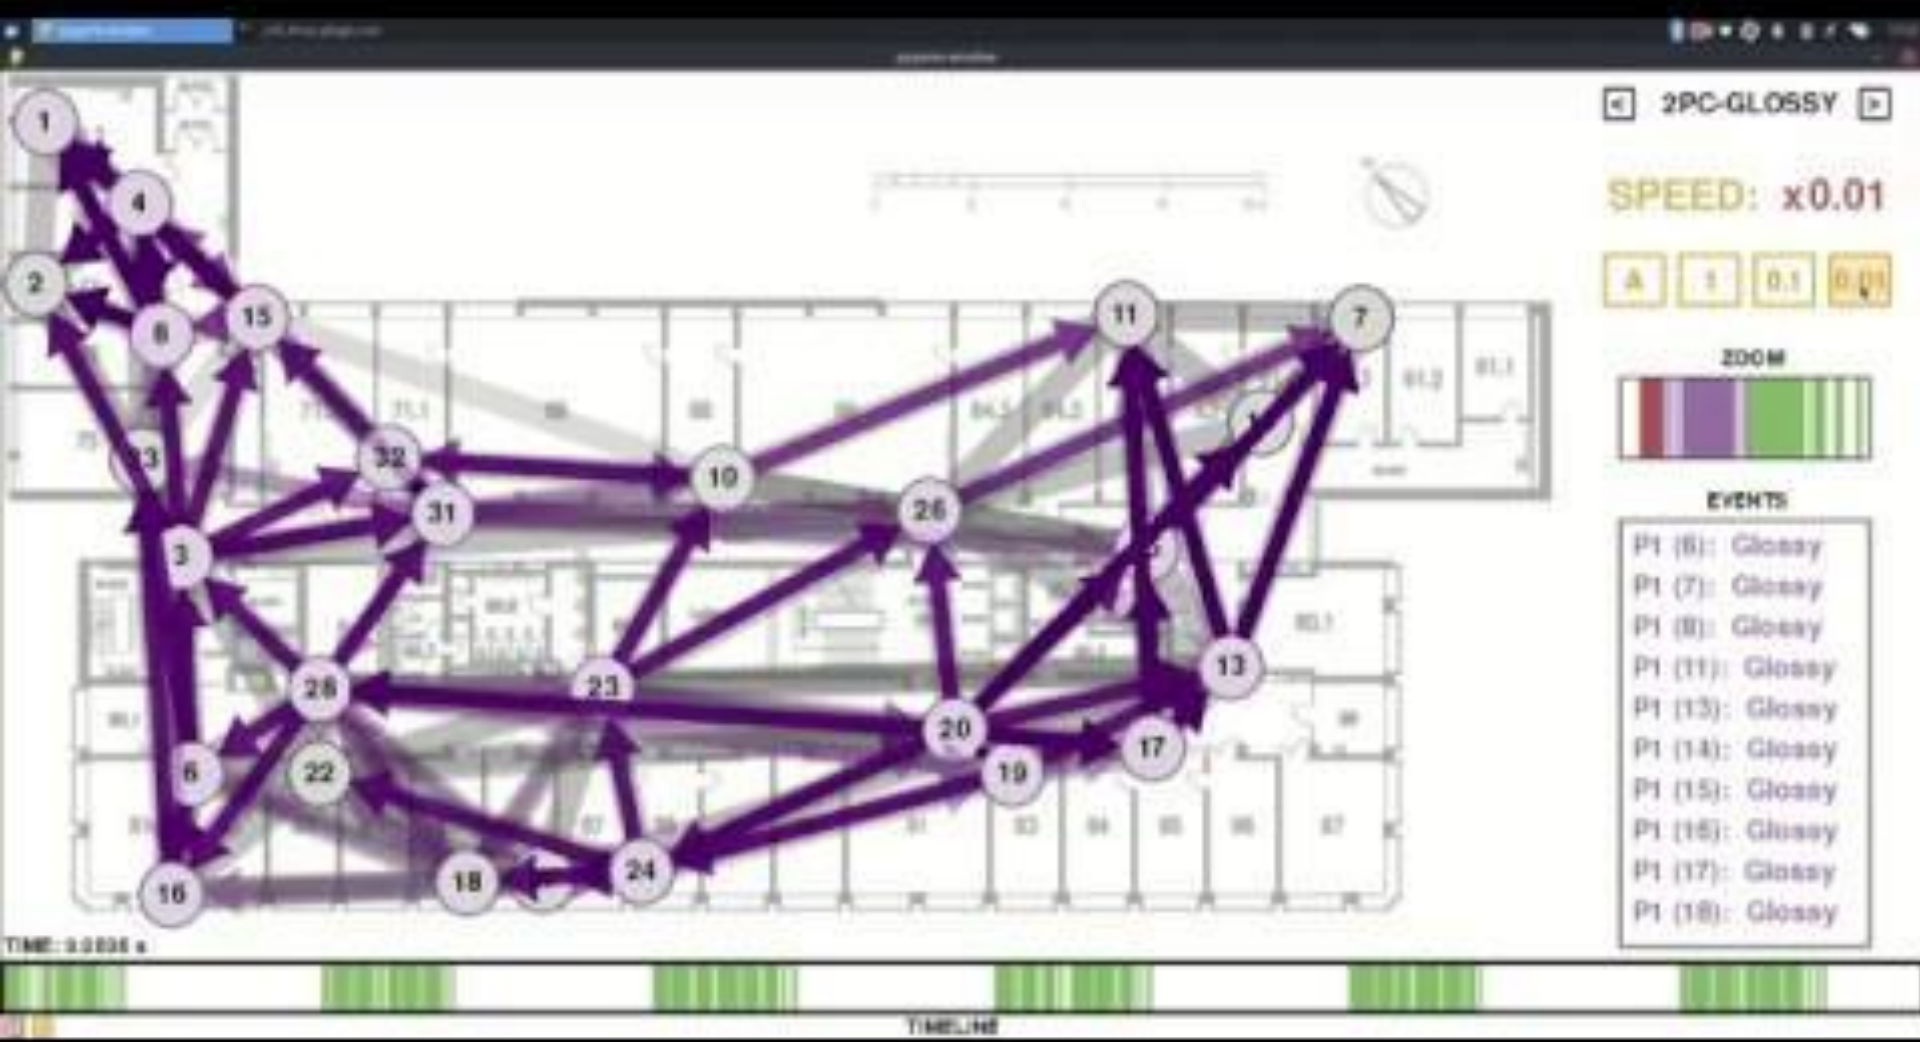
2PC-GLOSSY
SPEED: x0.01
A
1
0.1
0.01
ZOOM
EVENTS
P1 (6): Glossy
P1 (7): Glossy
P1 (8): Glossy
P1 (11): Glossy
P1 (13): Glossy
P1 (14): Glossy
P1 (15): Glossy
P1 (16): Glossy
P1 (17): Glossy
P1 (18): Glossy
TIMELINE

Additional

XPC with Chaos

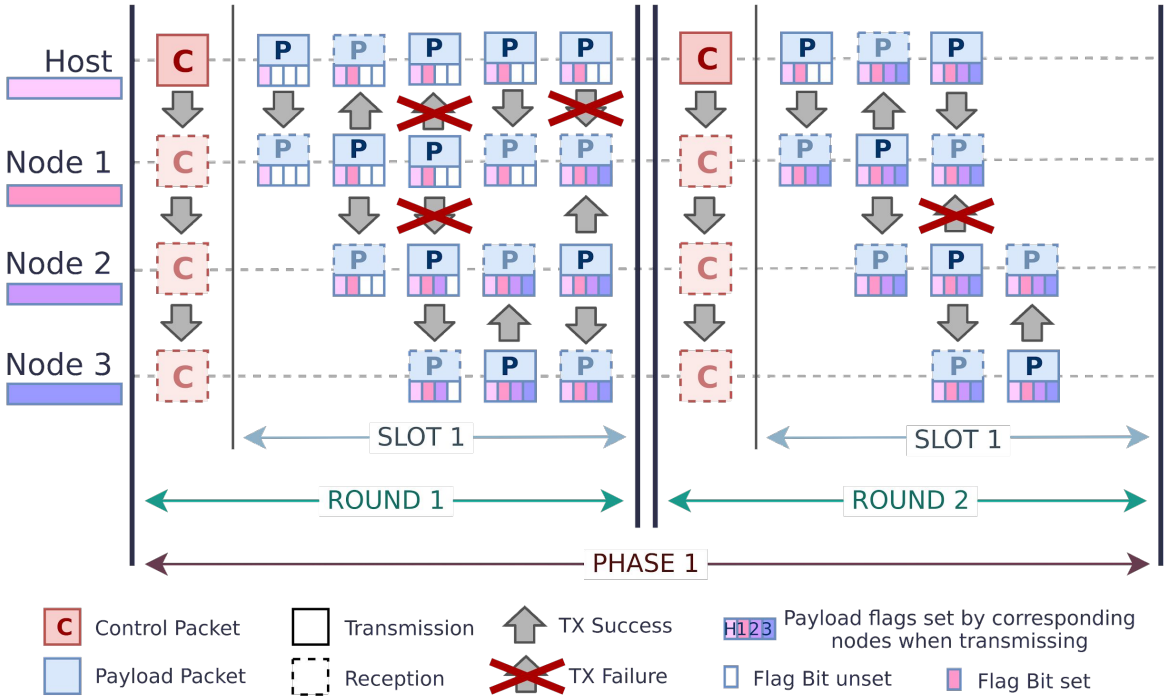


2PC-CHAOS
SPEED: x0.005
A
1
0.1
0.01
ZOOM
EVENTS
P1 (10): Control
P1 (10): Chaos
TIMELINE