Population protocols with unreliable communication

Michael Raskin

Technical University of Munich

09–10.09.2021

Population protocols — original

Small sensors: constant memory, identical,
pairwise interaction when accidentally near

Accepting and rejecting states
Aiming for eventual consensus — unanimous acceptance or rejection
(of initial configuration)

[AADFP 2004]

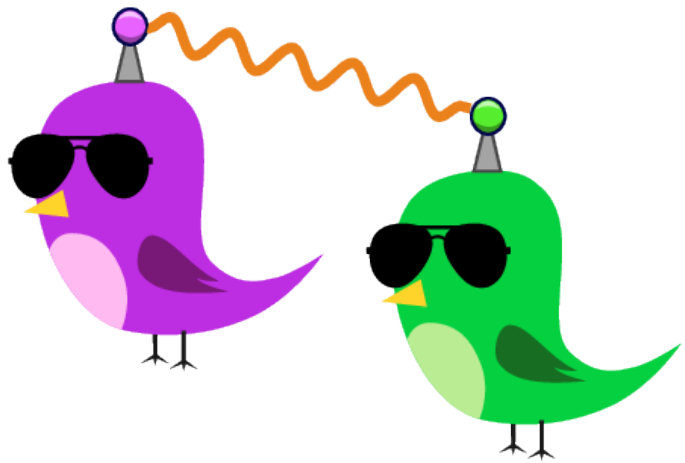

Population protocols — original

Small sensors: constant memory, identical,
pairwise interaction when accidentally near

Accepting and rejecting states
Aiming for eventual consensus — unanimous acceptance or rejection
(of initial configuration)

[AADFP 2004]

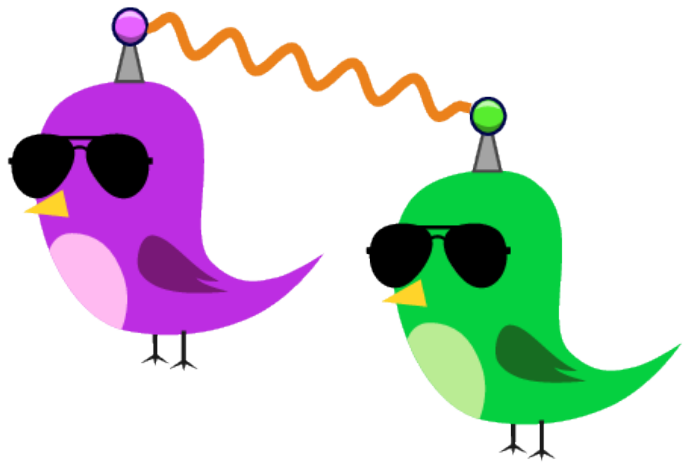

Population protocols — original

Small sensors: constant memory, identical,
pairwise interaction when accidentally near

Accepting and rejecting states
Aiming for eventual consensus — unanimous acceptance or rejection
(of initial configuration)

[AADFP 2004]

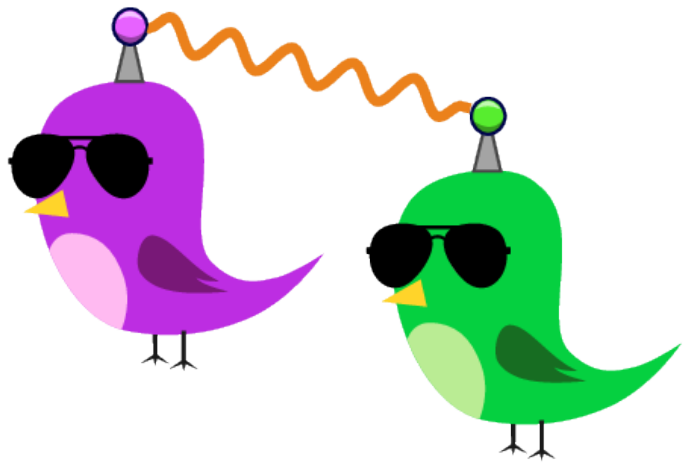

Population protocols — it got worse

Maybe messages can take time to arrive?
Maybe broadcast?
Or only multicast?

Population protocols — applicability

Our sensors are so small ...

Hard to find < 128 bytes RAM (agent ID: 4 bytes? 6 bytes?)

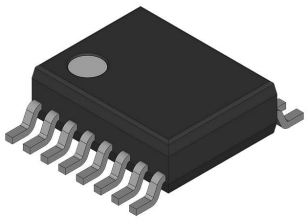

Edge of decidability
Even beefy servers like simplicity!

Population protocols — applicability

Our sensors are so small ...

Hard to find < 128 bytes RAM (agent ID: 4 bytes? 6 bytes?)

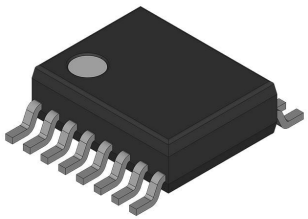

Edge of decidability

Even beefy servers like simplicity!

Population protocols — applicability

Our sensors are so small ...

Hard to find < 128 bytes RAM (agent ID: 4 bytes? 6 bytes?)

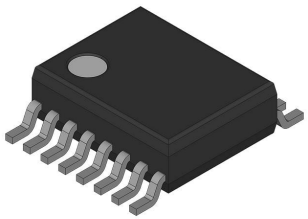

Edge of decidability
Even beefy servers like simplicity!

Are population protocols cheap?

Low memory — nice

Atomic interaction ...expensive in distributed systems
What if atomicity is lost?

Are population protocols cheap?

Low memory — nice

Atomic interaction ...expensive in distributed systems
What if atomicity is lost?

Unreliable communication:setting

Goal: generic upper bound on consensus predicates

Constant memory per agent
Slow messages from fixed finite language

Few restrictions on interaction rules
(newcomers can observe but not prevent interactions)
Wide range of scheduling

But! No atomicity (communication is unreliable)

Unreliable communication:setting

Goal: generic upper bound on consensus predicates

Constant memory per agent
Slow messages from fixed finite language

Few restrictions on interaction rules
(newcomers can observe but not prevent interactions)
Wide range of scheduling

But! No atomicity (communication is unreliable)

Unreliable communication:setting

Goal: generic upper bound on consensus predicates

Constant memory per agent
Slow messages from fixed finite language

Few restrictions on interaction rules
(newcomers can observe but not prevent interactions)
Wide range of scheduling

But! No atomicity (communication is unreliable)

Unreliable communication:setting

Goal: generic upper bound on consensus predicates

Constant memory per agent
Slow messages from fixed finite language

Few restrictions on interaction rules
(newcomers can observe but not prevent interactions)
Wide range of scheduling

But! No atomicity (communication is unreliable)

Unreliable communication: unreliability

You never know if interaction succeeded
Some agents update while others might not
But if broadcast fails, nobody receives it

Protocols with unreliable communication: expressivity

What predicates can be consensus-computed?

Boolean combinations of constant thresholds
There are ≥ 3 *red birds or* ≤ 5 *black ones*

Same as immediate observation/one-way communication protocols
which have PSPACE model checking, semilinear reachability...

Protocols with unreliable communication: expressivity

What predicates can be consensus-computed?

Boolean combinations of constant thresholds
There are ≥ 3 *red birds or* ≤ 5 *black ones*

Same as immediate observation/one-way communication protocols
which have PSPACE model checking, semilinear reachability...

Protocols with unreliable communication: expressivity

What predicates can be consensus-computed?

Boolean combinations of constant thresholds
There are ≥ 3 *red birds or* ≤ 5 *black ones*

Same as immediate observation/one-way communication protocols
which have PSPACE model checking, semilinear reachability...

Power inversion: message models

Assume each step is either receipt or transmission
of some messages by one agent

Reliable communication: strictly more expressive than IO
Unreliable communication: strictly less expressive than IO

Power inversion: message models

Assume each step is either receipt or transmission
of some messages by one agent

Reliable communication: strictly more expressive than IO
Unreliable communication: strictly less expressive than IO

Techniques

General case: strong generalisation of copycat arguments
Copycat doesn't copy the path, but we guarantee
the endpoint is copied

Message-based case: saturation
Flood what you can, fail to update while receiving rare messages

Techniques

General case: strong generalisation of copycat arguments
Copycat doesn't copy the path, but we guarantee
the endpoint is copied

Message-based case: saturation
Flood what you can, fail to update while receiving rare messages

Conclusion

Scope, interest, and missing directions
around population protocols

Unreliability of communication gives same expressivity
to wide range of models

IO protocols are nice

Thanks for your attention

Questions?

Conclusion

Scope, interest, and missing directions
around population protocols

Unreliability of communication gives same expressivity
to wide range of models

IO protocols are nice