New Models for Population Protocols

Othon Michail

Joint work with:
Ioannis Chatzigiannakis
Stavros Nikolaou
Andreas Pavlogiannis
Paul Spirakis

Research Academic Computer Technology Institute (RACTI)

ACAC '10
August 2010

Outline I

1. Population Protocols
   - Intro - Population Protocols
   - Computational Power
   - Enhancing the PP model
2. Mediated Population Protocols
   - A Formal Model
   - Computational Power
3. Passively Mobile Machines
   - A Formal Model
   - Computational Power

Population Protocol Model
[Angluin, Aspnes, Diamadi, Fischer, and Peralta, PODC '04]

- Tiny sensor nodes (**agents**) move passively
- **Tiny:**
  - Protocol descriptions independent of the # agents (**uniformity**)
  - **Anonymity**
- **Passively mobile:**
  - Mobility stems from some natural phenomenon
    - wind, water flow, animals moving . . .
    - modeled by some **fair adversary scheduler** selecting ordered pairs of agents to interact
- **Fairness:**
  - $(C \to C') \wedge (C$ appears infinitely often$) \Rightarrow C'$ appears infinitely often
  - **Weak assumption**[Chatzigiannakis, Dolev, Fekete, Michail, and Spirakis, OPODIS '09]: All consistent probabilistic schedulers satisfy it with probability 1

A Formal Model

- finite **input and output alphabets** X and Y
- finite set of **states** Q
- **input function** $I : X \to Q$
- **output function** $O : Q \to Y$
- **transition function** $\delta : Q \times Q \to Q \times Q$

  $\delta(p, q) = (p', q')$ or simply $(p, q) \to (p', q')$ is called a **transition**
- Population protocols **do not halt**, instead we require their outputs to **stabilize**

Flock of Birds: A Canonical Example

- Assume a complete **communication graph** G
- Each agent senses the temperature of a distinct bird after a global start signal
- If detected elevated temperature input 1, else 0 (i.e. $X = \{0, 1\}$)
  "*Find if at least 5 sensors have detected elevated temperature*"
- We want every agent to eventually output
  - 1, if at least 5 birds were found sick
  - 0, otherwise

Computational Power

Theorem ([Angluin, Aspnes, Diamadi, Fischer, and Peralta, PODC '04] & [Angluin, Aspnes, and Eisenstat, PODC '06])

A predicate is computable in the basic population protocol model if and only if it is semilinear (definable in Presburger arithmetic)

Stably Computable (semilinear)

- "The number of *a*s is greater than 5" (i.e. $N_a > 5$)
- $(N_a = N_b) \vee (\neg(N_b > N_c))$

Non-stably computable (non-semilinear)

- "The number of *c*s is the product of the number of *a*s and the number of *b*s" (i.e. $N_c = N_a \cdot N_b$)

Enhancing the PP model

- ***SEM* is a small class**
- PPs can tolerate only $O(1)$ crash failures and 0 Byzantine agents [Delporte-Gallet, Fauconnier, Guerraoui, Ruppert, '06]
- **Major goal:** **Extend the PP model** with extra realistic and implementable assumptions in order to improve:
  - **computational power**
  - **fault tolerance**
  - **time to convergence**

Mediated Population Protocols
[Chatzigiannakis, Michail, Spirakis, ICALP '09]

- A MPP is a PP that additionally has
  - a finite set of **edge states** S
- and an extended transition function δ of the form
  - $\delta : Q \times Q \times S \to Q \times Q \times S$

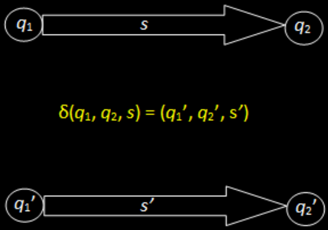

Figure: Each link is a constant storage

Question of Interest

What is the class of computable predicates here?

Computational Power

- Complete communication graphs (n denotes the **population size**)
- All edges are initially in a common state s_0
- *MPS*: the corresponding class
- For all $p \in MPS$, p is **symmetric**

Theorem ([Chatzigiannakis, Michail, Spirakis, ICALP '09] & [Chatzigiannakis, Michail, Nikolaou, Pavlogiannis, Spirakis, MFCS '10])

$p \in MPS$ iff p is symmetric and $p \in NSPACE(n^2)$

The Lower Bound

Theorem ([Chatzigiannakis, Michail, Nikolaou, Pavlogiannis, Spirakis, MFCS '10])

Any symmetric predicate in $NSPACE(n^2)$ belongs to MPS

1. Spanning Process

1. **Spanning process:** Agents become organized into a correctly labeled spanning line graph

1. Spanning Process

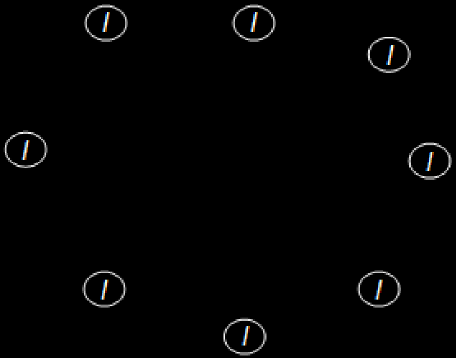

Figure: All agents are simple leaders

1. Spanning Process

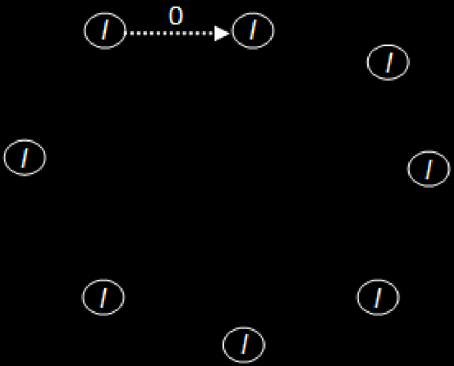

Figure: An interaction takes place

1. Spanning Process

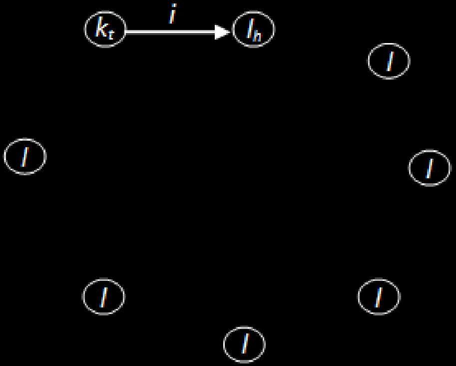

Figure: A line graph gets formed

1. Spanning Process

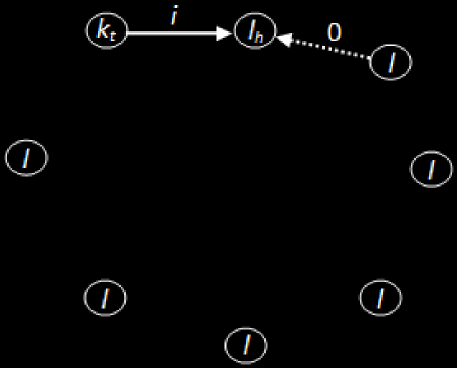

Figure: The line graph expands

1. Spanning Process

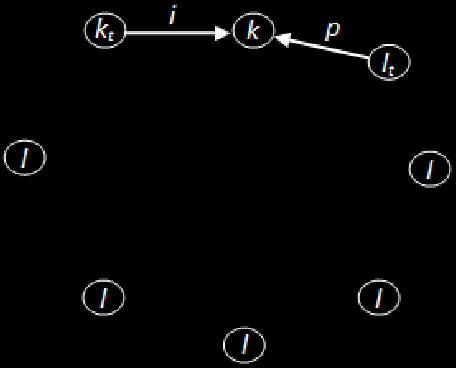

Figure: The line graph expands

1. Spanning Process

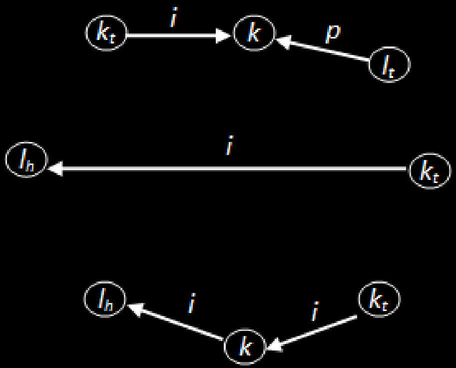

Figure: Separate line graphs formed

1. Spanning Process

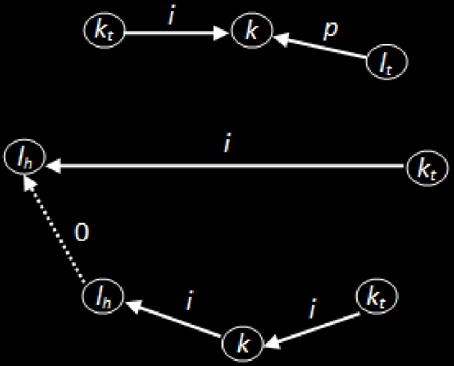

Figure: Line graphs get merged

1. Spanning Process

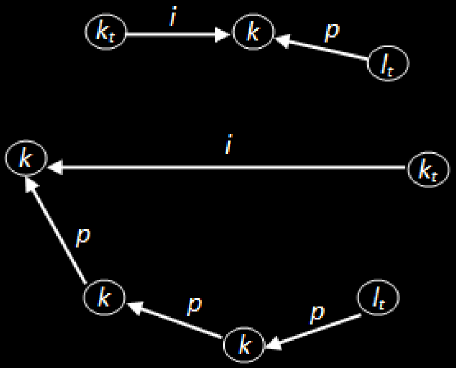

Figure: Line graphs get merged

1. Spanning Process

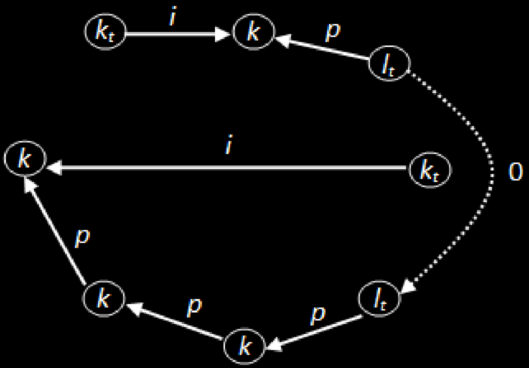

Figure: Line graphs get merged

1. Spanning Process

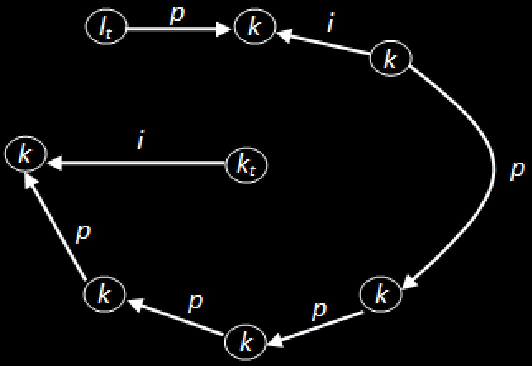

Figure: A correctly labeled spanning line graph

2. Reinitialization Process

- **2. Reinitialization Process:** The agents don't know when the spanning process ends
- Whenever a line graph is expanded they reinitialize the simulation

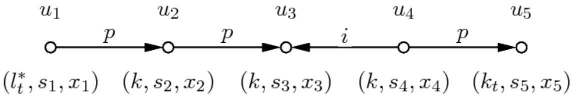

Figure: Just after merging. The leader endpoint has the special star mark. The reinitialization process begins.

- Whenever a line graph expands or two line graphs get merged the simulation is reinitialized in all agents and all outgoing edges

3. Simulation Process

- **3. Simulation Process:** The remaining $O(n^2)$ edges are used as tape cells to simulate a TM

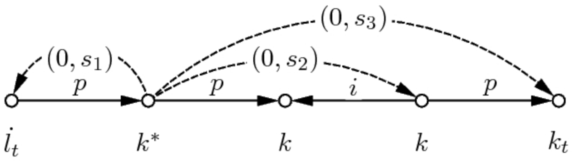

Figure: The agent in k^* controls now the simulation.

- **Nondeterminism:** Stems from the nondeterminism of the interaction pattern

The Upper Bound

Theorem ([Chatzigiannakis, Michail, Spirakis, ICALP '09])

$p \in MPS$ *implies that* $p \in NSPACE(n^2)$

- A configuration consists of $O(n^2)$ states of constant size
- To compute p in $O(n^2)$ space we perform a nondeterministic search on the transition graph of the protocol that stably computes it (by always storing at most one configuration)

Passively Mobile Machines
[Chatzigiannakis, Michail, Nikolaou, Pavlogiannis, Spirakis, FRONTS-TR '10]

- In the PM model each agent
  - is a **multitape Turing machine**
  - has tapes unbounded to the right

Passively Mobile Machines

- X is the **input alphabet**, where $\sqcup \notin X$ ($\sqcup$ is the **blank symbol**)
- Γ is the **tape alphabet**, where $\sqcup \in \Gamma$ and $X \subset \Gamma$
- Q is the set of **states**
- $\delta : Q \times \Gamma^4 \to Q \times \Gamma^4 \times \{L, R\}^4 \times \{0, 1\}$ is the **internal transition function**
- $\gamma : Q \times Q \to Q \times Q$ is the **external transition function**
- $q_0 \in Q$ is the **initial state**

Question of Interest

What is the class of computable predicates for each space bound?

Computational Power

- Complete communication graphs
- $PMSPACE(f(n))$: the corresponding class when each agent uses space $O(f(n))$
- $PMSPACE(c) = SEM$
- For all f and $p \in PMSPACE(f(n))$, p is **symmetric**

Theorem ([Chatzigiannakis, Michail, Nikolaou, Pavlogiannis, Spirakis, FRONTS-TR '10])

$\forall f$ *s.t.* $f(n) = \Omega(\log n)$, $p \in PMSPACE(f(n))$ *iff* p *is symmetric and* $p \in NSPACE(nf(n))$

Another Natural Question

What happens below $\log n$**?**

$\log n$: A Computational Threshold

Theorem ([Chatzigiannakis, Michail, Nikolaou, Pavlogiannis, Spirakis, FRONTS-TR '10])

For any $f : \mathbb{N} \to \mathbb{N}$, any predicate in PMSPACE$(f(n))$ is also in SNSPACE$(2^{f(n)}(f(n) + \log n))$

- $2^{O(f(n))}$ different agent configurations (internal configurations)
- each of size $O(f(n))$
- $O(f(n)2^{f(n)})$ space, together with a number per agent configuration representing # of agents in that agent configuration

$\log n$: A Computational Threshold

Some examples:

- $f(n) = \log n$, $O(2^{\log n}(2\log n)) = O(n\log n)$ (was $n\log n$)
- $f(n) = \log\log n$, $O(2^{\log\log n}(\log\log n + \log n)) = O(\log^2 n)$ (was $n\log\log n$)
- $f(n) = n$, $O(2^n(n + \log n)) = O(n2^n)$ (was n^2)

$\log n$: A Computational Threshold

Theorem (Symmetric Space Hierarchy Theorem)

For any function $f : \mathbb{N} \rightarrow \mathbb{N}$, a symmetric language L exists that is decidable in $O(f(n))$ (non)deterministic space but not in $o(f(n))$ (non)deterministic space

Proof.

Follows immediately from the unary (tally) separation language presented in [Geffert, TCS '03] and the fact that any unary language is symmetric □

$\log n$: A Computational Threshold

Theorem

For any $f(n) = o(\log n)$ *it holds that* $PMSPACE(f(n)) \subsetneq SNSPACE(n f(n))$

Proof.

By considering the previous theorems, it suffices to show that $2^{f(n)}(f(n) + \log n) = o(nf(n))$ for $f(n) = o(\log n)$. We have that

$$2^{f(n)}(f(n) + \log n) = 2^{o(\log n)}O(\log n) = o(n)O(\log n),$$

which obviously grows slower than $nf(n) = n \cdot o(\log n)$. □

Conclusion: $f(n) = \Theta(\log n)$ **acts as a threshold**

Conclusions

- We have proposed **2 new theoretical models for passively mobile sensor networks**
- Both **MPP** and **PM with $\Omega(\log n)$ available space per agent** **can use the whole memory for the simulation of a NTM that decides symmetric languages**
  - **Population protocols do not achieve this** as they have $O(n)$ space but cannot simulate linear-space NTMs
  - We showed that **PM with $o(\log n)$ space also does not**
- **$\log n$ memory is more realistic than constant and it is also an extremely small requirement**
- Due to its **threshold behavior** it seems to be **the best memory selection**

Open Problems

- **Fault tolerance** of both models (preconditions?)
- **Expected time complexity** of predicates under some **probabilistic scheduling assumption**
- **Protocol verification** (see e.g. [Chatzigiannakis, Michail, Spirakis, SSS '10] for PPs)
- **Stable decidability of properties of the communication graph** (see e.g. [Chatzigiannakis, Michail, Spirakis, SSS '10 - 2] for a first attempt for MPPs)
- **Exact characterization of $PMSPACE(f(n))$ for all $f(n) = o(\log n)$**
  - At a first glance it seems that $\log \log n$ is another threshold and that between $\log \log n$ and $\log n$ the power depends on the # of agents that can be assigned uids

FRONTS

- This work has been partially supported by the ICT Programme of the European Union under contract number ICT-2008-215270 (FRONTS).
- FRONTS is a joint effort of eleven academic and research institutes in foundational algorithmic research in Europe.
- The effort is towards establishing the **foundations of adaptive networked societies of tiny artefacts**.

SEVENTH FRAMEWORK PROGRAMME

Thank You!