

THE DISTRIBUTED COMPUTING COLUMN

BY

MARIO MAVRONICOLAS

Department of Computer Science, University of Cyprus
75 Kallipoleos St., CY-1678 Nicosia, Cyprus
mavronic@cs.ucy.ac.cy

AN INTRODUCTION TO POPULATION PROTOCOLS

James Aspnes*
Yale University

Eric Ruppert†
York University

Abstract

The population protocol model describes a collection of tiny mobile agents that interact with one another to carry out a computation. The agents are identically programmed finite state machines. Interactions between pairs of agents cause the two agents to update their states. These interactions are scheduled by an adversary, subject to a fairness constraint. Input values are initially distributed to the agents, and the agents must eventually converge to the correct output value. This framework can be used to model mobile *ad hoc* networks of tiny devices or collections of molecules undergoing chemical reactions. We survey results that describe what can be computed in various versions of the population protocol model.

1 Introduction

The population protocol model [3] was designed to represent sensor networks consisting of very limited mobile agents with no control over their own movement.

*Supported in part by NSF grant CNS-0435201.

†Supported in part by the Natural Sciences and Engineering Research Council of Canada.

It also bears a strong resemblance to models of interacting molecules in theoretical chemistry [16, 17]. The defining features of the basic model are:

1. Anonymous, finite-state agents. The system consists of a large population of indistinguishable finite-state agents.
2. Computation by direct interaction. In the original model, agents do not send messages or share memory; instead, an **interaction** between two agents updates both of their states according to a joint transition table. The actual mechanism of such interactions is abstracted away.
3. Unpredictable interaction patterns. The choice of which agents interact is made by an adversary. Agents have little control over which other agents they interact with, although the adversary may be limited to pairing only agents that are adjacent in an **interaction graph**, typically representing distance constraints. A strong global **fairness condition** is imposed on the adversary to ensure the protocol makes progress.
4. Distributed inputs and outputs. The input to a population protocol is distributed across the initial state of the entire population. Similarly, the output is distributed to all agents.
5. Convergence rather than termination. Population protocols generally cannot detect when they have finished; instead, the agents' outputs are required to converge after some finite time to a common, correct value.

A formal definition is given in Section 2.

The population protocol model was inspired in part by work by Diamadi and Fischer [13] on trust propagation in a social network. The **urn automata** of [2] can be seen as a first draft of the model that retained in vestigial form several features of classical automata: instead of interacting with each other, agents could interact only with a finite-state controller, complete with input tape. The motivation given for the current model in [3] was the study of sensor networks in which passive agents were carried along by other entities; the canonical example was sensors attached to a flock of birds. The name of the model was chosen by analogy to **population processes** [22] in probability theory.

A population protocol often looks like an amorphous soup of lost, nearly mindless, anonymous agents blown here and there at the whim of the adversary. But though individual agents lack much intelligence or control over their own destinies, the population as a whole is nonetheless capable of performing significant computations—under some conditions, it has the same power as a traditional computer with the same total storage capacity. Some examples of simple population protocols are given in Section 2.

Much of the work so far on population protocols has concentrated on characterizing what predicates on the input configuration can be computed in different variants of the model and under various assumptions, such as a bounded-degree interaction graph or random scheduling.

If the interaction graph is unrestricted, the worst case for computation turns out to be a complete interaction graph, since any other interaction graph can simulate a complete interaction graph by shuffling agents between the nodes [3]. In a complete interaction graph, all agents with the same state are indistinguishable, and only the counts of agents in each state affect the outcome of the protocol. The set of computable predicates in most variants of the basic model for such a graph is now known to be either exactly equal to or closely related to the set of **semilinear** predicates, those definable in **first-order Presburger arithmetic** [18, 25]. These results, which originally appeared in [1, 3, 5, 7, 8, 9, 12], are summarized in Sections 3, 4, 5, 7 and 9. Sometimes the structure of incomplete interaction graphs can be exploited to simulate a Turing machine, which implies that a restricted interaction graph can make the system stronger than the complete graph.

More recent work has concentrated on performance. Because the weak scheduling assumptions in the basic model allow the adversary to draw out a computation indefinitely, the worst-case adversary scheduler is replaced by a random scheduling assumption, where the pair of agents that interacts at each step is drawn uniformly from the population as a whole. This gives a natural notion of **time** equal to the total number of steps to convergence and **parallel time** equal to the average number of steps initiated by any one agent (essentially the total number of steps divided by the number of agents).

As with adversarial scheduling, for random scheduling the best-understood case is that of a complete interaction graph. In this case, it is possible to construct a register machine, where subpopulations of the agents hold tokens representing the various register values in unary. It is not hard to implement register operations like addition, subtraction, and comparison by local operations between pairs of agents; with the election of a leader, we can further construct a finite-state control. The main obstacle to implementing a complete register machine is to ensure that every agent completes any needed tasks for each instruction cycle before the next cycle starts. In [3], this was handled by having the leader wait a polynomial number of steps on average before starting the next cycle, a process which gives an easy proof of polynomially-bounded error but which also gives an impractically large slowdown. Subsequent work has reduced the slowdown to polylogarithmic by using epidemics both to propagate information quickly through the population and to provide timing [4,6]. These results are described in more detail in Section 6.

In addition to work on the basic population protocol model, there have been several extensions of the model to more accurately reflect the requirements of practical systems. The basic model requires coordinated two-way communication

between interacting agents; this assumption is relaxed in Section 4. Work on incorporating agent failures into the model are discussed in Sections 7 and 9. Versions of the model that give agents slightly increased memory capacity are discussed in Section 8.

2 The basic model

In the basic population protocol model, a collection of agents are each given an input value, and agents have pairwise interactions in an order determined by a scheduler, subject to some fairness guarantee. Each agent is a kind of finite state machine and the “program” for the system describes how the states of two agents can be updated by an interaction. The agents’ output values change over time and must eventually converge to the correct output value for the inputs that were initially distributed to the agents.

A protocol is formally specified by

- Q , a finite set of possible states for an agent,
- Σ , a finite input alphabet,
- ι , an input map from Σ to Q , where $\iota(\sigma)$ represents the initial state of an agent whose input is σ ,
- ω , an output map from Q to the output range Y , where $\omega(q)$ represents the output value of an agent in state q , and
- $\delta \subseteq Q^4$, a transition relation that describes how pairs of agents can interact.

We now describe how a computation proceeds according to such a protocol. The computation takes place among n **agents**, where $n \geq 2$. Each agent is given an input value from Σ . Each agent’s initial state is determined by applying ι to its input value. This determines an initial configuration for an execution. A **configuration** of the system can be described by a vector of all the agents’ states. Because agents with the same state are indistinguishable, each configuration can be summarized as an unordered multiset of states.

An execution of a protocol proceeds from the initial configuration by interactions between pairs of agents. Suppose two agents in states q_1 and q_2 meet and have an interaction. They can change into states q'_1 and q'_2 as a result of the interaction if (q_1, q_2, q'_1, q'_2) is in the transition relation δ . We sometimes describe δ by listing all possible interactions using the notation $(q_1, q_2) \rightarrow (q'_1, q'_2)$. (We assume a null transition $(q_1, q_2) \rightarrow (q_1, q_2)$ if no others are specified with (q_1, q_2) on the left hand side.) If there is only one possible transition $(q_1, q_2) \rightarrow (q'_1, q'_2)$ for each pair (q_1, q_2) , then the protocol is **deterministic**. If C and C' are configurations, we write $C \rightarrow C'$ if C' can be obtained from C by a single interaction of two agents.

This means that C contains two states q_1 and q_2 and C' is obtained from C by replacing q_1 and q_2 by q'_1 and q'_2 , where (q_1, q_2, q'_1, q'_2) is in δ . An **execution** of the protocol is an infinite sequence of configurations C_0, C_1, C_2, \dots , where C_0 is an initial configuration and $C_i \rightarrow C_{i+1}$ for all $i \geq 0$. Thus, an execution is a sequence of snapshots of the system after each interaction occurs. In a real distributed execution, interactions could take place simultaneously, but when writing down an execution we can order those simultaneous interactions arbitrarily.

The order in which pairs of agents interact is unpredictable: we think of the schedule of interactions as being chosen by an adversary, so that protocols must work correctly under any schedule the adversary may choose. In order for meaningful computations to take place, we must put some restrictions on the adversarial scheduler; otherwise it could divide the agents into isolated groups and schedule interactions only between agents that belong to the same group.

The **fairness** condition that we impose on the scheduler is quite simple to state, but is somewhat subtle. Essentially, we do not allow the scheduler to avoid a possible step forever. More formally, if C is a configuration that appears infinitely often in an execution, and $C \rightarrow C'$, then C' must also appear infinitely often in the execution. Another way to think of this is that anything that is always possible eventually happens: it is equivalent to require that any configuration that is always reachable is eventually reached.

At any point during an execution of a population protocol, each agent's state determines its output at that time. If the agent is in state q , its output value is $\omega(q)$. Thus, an agent's output may change over the course of an execution. The fairness constraint allows the scheduler to behave arbitrarily for an arbitrarily long period of time, but does require that it behave nicely eventually. It is therefore natural to phrase correctness as a property to be satisfied eventually too. For example, the scheduler could schedule only interactions between agents 1 and 2, leaving the other $n - 2$ agents isolated, for millions of years, and it would be unreasonable to expect any sensible output during the period when only two agents have undergone state changes. Thus, for correctness, we require that all agents produce the correct output (for the input values that were initially distributed to the agents) at some time in the execution and continue to do so forever after that time.

To summarize, we say that a protocol computes a function f that maps multisets of elements of Σ to Y if, for every such multiset I and every fair execution that starts from the initial configuration corresponding to I , the output value of every agent eventually stabilizes to $f(I)$.

Example 1. Suppose each agent is given an input bit, and we want all agents to output the 'or' of those bits. There is a very simple protocol to accomplish this: each agent with input 0 simply outputs 1 as soon as it discovers that another agent had input 1. Formally, we have $\Sigma = Y = Q = \{0, 1\}$ and the input and output maps

are the identity functions. The only interaction in δ is $(0, 1) \rightarrow (1, 1)$. If all agents have input 0, no agent will ever be in state 1. If some agent has input 1 the number of agents with state 1 cannot decrease and fairness ensures that it will eventually increase to n . In both cases, all agents stabilize to the correct output value.

Example 2. Suppose the agents represent dancers. Each dancer is (exclusively) a leader or a follower. We wish to determine whether there are more leaders than followers. We use $Y = \{0, 1\}$, with 1 indicating that there are more leaders than followers. A centralized solution would count the leaders and the followers and compare the totals. A more distributed solution is to ask everyone to start dancing with a partner (who must dance the opposite role) and then see if any dancers are left without a partner. We formalize this cancellation procedure as a population protocol with $\Sigma = \{L, F\}$ and $Q = \{L, F, 0, 1\}$. The input map ι is the identity, and the output map ω maps L and 1 to 1 and maps F and 0 to 0. The transitions of δ are $(L, F) \rightarrow (0, 0)$, $(L, 0) \rightarrow (L, 1)$, $(F, 1) \rightarrow (F, 0)$ and $(0, 1) \rightarrow (0, 0)$. The first rule ensures that, eventually, either no L 's or no F 's will remain. At that point, if there are L 's remaining, the second rule ensures that all agents will eventually produce output 1. Similarly, the third rule takes care of the case where F 's remain. In the case of a tie, the third rule ensures that the output stabilizes to 0.

It may not be obvious why the protocol of in Example 2 must converge. Consider, for example, the transitions $\{\underline{L}, L, \underline{F}\} \rightarrow \{\underline{0}, \underline{L}, 0\} \rightarrow \{\underline{1}, L, \underline{0}\} \rightarrow \{0, \underline{L}, \underline{0}\} \rightarrow \{\underline{0}, L, \underline{1}\} \rightarrow \{0, L, 0\}$, where in each configuration, the agents that are about to interact are underlined. By repeating this loop, we obtain a non-converging execution in which every pair of agents interacts infinitely often. However, this execution is not fair: the configuration $\{0, L, 1\}$ appears infinitely often and $\{0, L, 1\} \rightarrow \{1, L, 1\}$, but $\{1, L, 1\}$ never appears. This is because the first two agents only interact at “inconvenient” times, *i.e.*, when third agent is in state 0. The definition of fairness rules this out. Thus, in some ways, the definition of fairness is stronger than saying that each pair of agents must interact infinitely often. (In fact, the two conditions are incomparable, since there can be fair executions in which two agents never meet. For example, an execution where every configuration is $\{L, L, L\}$ and all interactions take place between the first two agents is fair.)

Exercise 3. Show the protocol of Example 2 converges in every fair execution.

The definition of fairness was chosen to be quite weak (although it is still strong enough to allow useful computations). Many models of mobile systems assume that the mobility patterns of the agents follow some particular probability distribution. The goal of the population protocol model is to be more general. If there is an (unknown) underlying probability distribution on the interactions, which might even vary with time, and that distribution satisfies certain independence properties and ensures that every interaction's probability is bounded away

from 0, then an execution will be fair with probability 1. Thus, any protocol will converge to the correct output with probability 1. So the model captures computations that are correct with probability 1 for a wide range of probability distributions, even though the model definition does not explicitly incorporate probabilities.

Other predicates can be computed using an approach similar to Example 2.

Exercise 4. Design a population protocol to determine whether more than 40% of the dancers are leaders.

Some predicates, however, require a different approach.

Example 5. Suppose each agent is given an input from $\Sigma = \{0, 1, 2, 3\}$ and we wish to find the sum of the inputs, modulo 4. The protocol can gather the sum (modulo 4) into a single agent. Once an agent has given its value to another agent, its value becomes null, and it obtains its output values from the eventually unique agent with a non-null value. Formally, we have $Q = \{0, 1, 2, 3, \perp_0, \perp_1, \perp_2, \perp_3\}$, where \perp_v represents a null value with output v . Let $\iota(v) = v$ and $\omega(v) = \omega(\perp_v) = v$ for $v = 0, 1, 2, 3$. The transition rules of δ are $(v_1, v_2) \rightarrow (v_1 + v_2, \perp_{v_1+v_2})$ and $(v_1, \perp_{v_2}) \rightarrow (v_1, \perp_{v_1})$, where v_1 and v_2 are 0, 1, 2 or 3. (The addition is modulo 4.)

In some cases, agents may know when they have converged to the correct output, but in general they cannot. While computing the ‘or’ of input bits (Example 1), any agent in state 1 knows that its state will never change again: it has converged to its final output value. However, no agent in the protocol of Example 5 can ever be certain it has converged, since some additional agents with input 1 could always join the computation and change the required output value.

Two noteworthy properties of the population protocol model are its uniformity and anonymity. A protocol is **uniform** because its specification has no dependence on the number of agents that take part. In other words, no knowledge about the number of agents is required by the protocol. The system is **anonymous** because the agents are not equipped with unique identifiers and all agents are treated in the same way by the transition relation. Indeed, because the state set is finite and does not depend on the number of agents in the system, there is not even room in the state of an agent to store a unique identifier.

3 Computability

In studying what functions can be computed in the population protocol model, there is no loss of generality in restricting attention to predicates, *i.e.*, functions with range $Y = \{0, 1\}$. For any function f with range Y , let $P_{f,y}$ be a predicate defined by $P_{f,y}(x) = 1$ if and only if $f(x) = y$. Then, f is computable if and only

if $P_{f,y}$ is computable for each $y \in Y$. The “only if” part of this statement is trivial. For the converse, a protocol can compute all the predicates $P_{f,y}$ in parallel, using a separate component of each agent’s state for each y . Note that Y is finite because each distinct output value corresponds to at least one state in the original protocol.

For the basic population protocol model, there is an exact characterization of the computable predicates: they are precisely the **semilinear predicates**, which we now define. A multiset over the input alphabet Σ can also be thought of as a vector with $d = |\Sigma|$ components, where each component is a natural number representing the multiplicity of one input character. For example, the input multiset $\{a, a, a, b, b\}$ over the input alphabet $\Sigma = \{a, b, c\}$ can be represented by the vector $(3, 2, 0) \in \mathbb{N}^3$. A **semilinear set** is a subset of \mathbb{N}^d that is a finite union of **linear sets** of the form $\{\vec{b} + k_1\vec{a}_1 + k_2\vec{a}_2 + \dots + k_m\vec{a}_m\}$, where \vec{b} is a d -dimensional base vector, \vec{a}_1 through \vec{a}_m are basis vectors, and k_1 through k_m are non-negative coefficients. A **semilinear predicate** on inputs is one that is true precisely on a semilinear set.

An alternative characterization of semilinear predicates is that they can be described by first-order logical formulas in Presburger arithmetic, which is arithmetic on the natural numbers with addition but not multiplication [25]. Presburger arithmetic allows for **quantifier elimination**, replacing universal and existential quantifiers with formulas involving addition, $<$, the mod- k congruence relation \equiv_k for each constant k , and the usual logical connectives \wedge , \vee , and \neg .

Angluin *et al.* [3] gave protocols to compute any threshold predicate or remainder predicate; the protocols are similar to those in Examples 2 and 5. This gives $<$ and \equiv_k . They further observed that addition is trivially obtained by renaming states: to compute $A + B$ from A and B we pretend that any A or B token is really an $A+B$ token. Finally, Boolean combinations (\wedge , \vee) of these predicates can be computed by running the protocols for each of the basic predicates in parallel, using separate components of the agents’ states, and negation (\neg) simply involves relabeling the output values. It follows that population protocols can compute all predicates definable with Presburger formulas, *i.e.*, all semilinear predicates.

Amazingly, the converse also holds: only semilinear predicates can be computed by population protocols. For the basic model this result was shown by Angluin, Aspnes, and Eisenstat [5] by applying results from partial order theory. The proof is quite involved, but the essential idea is that, like finite-state automata, population protocols can be “pumped” by adding extra input tokens that turn out not to affect the final output. By carefully considering exactly when this is possible, it can be shown that the positive inputs to a population protocol (considered as sets of vectors) can be separated into a collection of cones over some finite set of minimal positive inputs, and that each of these cones can be further expressed using only a finite set of basis vectors. This is sufficient to show that the predicate corresponds to a semilinear set as described above. We thus have:

Theorem 6 ([3,5,7]). *A predicate is computable in the basic population protocol model if and only if it is semilinear.*

Similar results with weaker classes of predicates hold for restricted models with various forms of one-way communication [8]; we describe these results in more detail in Section 4. Indeed, these results were a precursor to the semilinearity theorem of [5]. The journal paper [7] combines and extends these results.

A useful property of Theorem 6 is that it continues to hold unmodified in many simple variants of the basic model. The reason is that any change that weakens the agents can only decrease the set of computable predicates, while any model that is still strong enough to compute congruence modulo k and comparison can still compute all the semilinear predicates. So the semilinear predicates continue to be those that are computable when the inputs are not given immediately but stabilize after some finite time [1] or when one agent in an interaction can see the other's state but not vice versa [7], as in each case it is still possible to compute congruence and threshold in the limit. A similar result holds when a small number of agents can fail [12]; here a slight modification must be made to allow for partial predicates that can tolerate the loss of part of the input. We describe all of these results in later sections.

4 One-way communication

In the basic population protocol model, it is assumed that two interacting agents can simultaneously learn each other's state before updating their own states as a result of the interaction. This requires two-way communication between the two agents. Angluin *et al.* [7] studied several weaker interaction models where, in an interaction, information flows in one direction only. A **receiver** agent learns the state of a **sender** agent, but the sender learns nothing about the state of the receiver. The power of a system with such one-way communication depends on the nature of the communication mechanism.

The model is called a **transmission** model if the sender is aware that an interaction has happened (and can update its own state, although the update cannot depend on the state of the receiver). In an **observation** model, on the other hand, the sender's state is passively observed by the receiver. Another independent attribute is whether an interaction happens instantaneously (**immediate transmission** and **immediate observation** models) or requires some interval of time (**delayed transmission** and **delayed observation** models). The **queued transmission** model is similar to the delayed transmission model, except that receivers can temporarily refuse incoming messages so that they are not overwhelmed with more incoming information than they can handle. The queued transmission model is the closest to traditional message-passing models of distributed computing.

The weakest of these one-way models is the delayed observation model: protocols can detect whether any particular input symbol is present or absent (and compute any predicate that depends only on this information) simply by observing other agents' input values. Nothing else can be computed: there is no way to distinguish between a sequence of observations of several agents with the same input and a sequence of observations of a single agent.

The immediate observation model is slightly stronger: it can count the number of agents with a particular input, up to some constant threshold. For example, a protocol can determine whether the number of copies of input symbol a is 0, 1, 2, 3 or more than 3. Consequently, any predicate that depends only on this kind of information can be computed. A kind of pumping lemma can be used to show that no other predicates are computable.

Angluin *et al.* also showed that the immediate and delayed transmission models are equivalent in power. They gave a characterization of the computable predicates that shows the power of these models is intermediate between the immediate observation model and the standard two-way model.

Finally, the queued transmission model is equivalent in power to the standard two-way model: any protocol designed for the two-way model can be simulated using queued transmission and vice versa. This holds even though the set of configurations in queued transmission is in principle unbounded; the ability to generate large numbers of buffered messages does not help the protocol, largely because there is no guarantee of where or when they will be delivered.

5 Restricted interaction graphs

In some cases the mobility of agents will have physical limitations, and this will limit the possible interactions that can occur. We can represent this information in an **interaction graph**, where nodes represent agents and edges represent possible interactions. The basic model corresponds to the case where the graph is complete. In this model, a configuration is always represented as a vector of states. (The agents are no longer indistinguishable, so we cannot use a multiset.) If C and C' are configurations, we write $C \rightarrow C'$ if C' can be obtained from C through a single interaction of *adjacent* agents, and the definitions of executions and fairness are as before, using this modified notion of a step.

Having a non-complete (but connected) interaction graph does not make the model any weaker, since adjacent agents can swap states to simulate free movement [3]. For some interaction graphs, the model becomes strictly more powerful. For example, consider a straight-line graph. It is not difficult to simulate a linear-space Turing machine by using each agent to represent one square of the Turing machine tape. This allows computation of any function or predicate in

LINSPACE, many of which are not semilinear and thus not computable in the complete interaction graph of the basic model.

In addition to computing predicates on the inputs to agents, it also makes sense in this model to ask whether properties of the interaction graph itself can be computed by the agents in the system. Such problems, which were studied by Angluin *et al.* [1], could have useful applications in determining the network topology induced by an *ad hoc* deployment of mobile agents.

As a simple example, one might want to determine whether the interaction graph has maximum degree k or more, for some fixed k . This can be done by electing a single moving leader token. Initially, all agents hold a leader token. When two leader tokens interact, the tokens coalesce, and when a leader agent interacts with a non-leader agent the leader token may change places. To test the maximum degree, the leader may instead choose to mark up to k distinct neighbors of its current node. By counting how many nodes it successfully marks, the leader can get a lower bound on the degree of the node.

A complication is that the leader has no way to detect when it has interacted with all neighbors of the current node. The best it can do is nondeterministically wait for some arbitrary but finite time before gathering in its marks and trying again. In doing so it relies on the fairness condition to eventually drive it to a state where it has correctly computed the maximum degree. Because the original population model used deterministic transitions, some additional machinery is needed to simulate a nondeterministic transition function by exploiting the nondeterminism of the interaction schedule. Note that the unmarking step does not require nondeterminism: since the leader keeps track of how many marks it has placed, it can simply wait until it has encountered each marked neighbor again.

During the initial leader election phase, two leaders deploying marks could interfere with each other. To handle this, the survivor of any interaction between two leaders collects all outstanding marks from both and resets its degree estimate.

A similar mechanism can be used to assign unique colors to all neighbors of each node in a bounded-degree graph: a wandering **colorizer** token deploys pairs of marks to its neighbors and recolors any it finds with the same color. Once this process converges, the resulting **distance-2 coloring** (so called because all nodes at distance 2 have distinct colors) effectively provides local identifiers for the neighbors of each node. These can be used to carry out arbitrary distributed computations using standard techniques (subject to the $O(1)$ space limit at each node). An example given in [1] is the construction of a rooted spanning tree, which can be used to simulate a Turing machine tape (as in the case of a line graph) by threading the Turing machine tape along a traversal of the tree. It follows that arbitrary LINSPACE-computable properties of bounded-degree graphs can be computed by population protocols.

6 Random interactions

An alternative assumption that also greatly increases the power of the model is to replace the adversarial (but fair) scheduler of the basic model with a more constrained interaction pattern. The simplest such variant assumes **random interactions**: each pair of agents is equally likely to interact at each step.

Protocols for random scheduling were given in the initial population protocol paper of Angluin *et al.* [3], based in part on similar protocols for the related model of urn automata [2]. The central observation was that the main limitation observed in trying to build more powerful protocols in the basic model was the inability to detect the absence of agents with a particular state. However, if a single leader agent were willing to wait long enough, it could be assured (with reasonably high probability) that it would meet every other agent in the population, and thus be able to verify the presence or absence of particular token values stored on the other agents by direct inspection. The method used was to have the leader issue a single special marked token to some agent; when the leader encountered this special agent k times in a row it could be reasonably confident that the number of intervening interactions was close to $\Theta(n^{k+1})$. This is sufficient to build unary counters supporting the usual increment, decrement, and zero test operations (the last probabilistic). With counters, a register machine with an $O(\log n)$ bit random-access memory can be simulated using a classic technique of Minsky [23].

The cost of this simulation is a polynomial blowup for the zero test and a further polynomial blowup in the simulation of the register machine. A faster simulation was given by Angluin, Aspnes, and Eisenstat [4], based on epidemics to propagate information quickly through the population. This simulation assumes a single designated leader agent in the initial configuration, which acts as the finite-state controller for the register machine. Register values are again stored in unary as tokens scattered across the remaining agents.

To execute an operation, the leader initiates an epidemic containing an operation code. This opcode is copied through the rest of the population in $\Theta(n \log n)$ interactions on average and with high probability; the latter result is shown to follow by a reduction to a concentration bound for coupon collector due to Kamath *et al.* [21]. Arithmetic operations such as addition, comparison, subtraction, and multiplication and division by constants can be carried out by the non-leader agents in $O(n \log^c n)$ interactions (or $O(\log^c n)$ parallel time units) each. Some of these algorithms are quite simple (adding A to B requires only adding a new B token to each agent that already holds an A token, possibly with an additional step of unloading extra B tokens onto empty agents to maintain $O(1)$ space per agent), while others are more involved (comparing two values in [4] involves up to $O(\log n)$ alternating rounds of doubling and cancellation, because simply having A and B tokens cancel each other as in Example 2 might require as many as

$\Theta(n^2)$ expected interactions for the last few survivors to meet). The most expensive operation is division, at $O(n \log^4 n)$ interactions (or $O(\log^4 n)$ parallel time units).

Being able to carry out individual arithmetic operations is of little use if one cannot carry out more than one. This requires that the leader be able to detect when an operation has finished, which ultimately reduces down to being able to detect when $\Theta(n \log n)$ interactions have occurred. Here the trick of issuing a single special mark is not enough, as the wait needed to ensure a low probability of premature termination is too long.

Instead, a **phase clock** based on successive waves of epidemics is used. The leader starts by initiating a phase 0 epidemic which propagates through the population in parallel to any other activity. When the leader meets an agent that is already infected with phase 0, it initiates a phase 1 epidemic that overwrites the phase 0 epidemic, and similarly with phase 2, 3, and so on, up to some fixed maximum phase $m - 1$ that is in turn overwritten by phase 0 again. Angluin *et al.* show that, while the leader might get lucky and encounter one of a small number of newly-infected agents in a single phase, the more typical case is that a phase takes $\Theta(n \log n)$ interactions before the next is triggered, and over m phases the probability that all are too short is polynomially small. It follows that for a suitable choice of m , the phase clock gives a high-probability $\Theta(n \log n)$ -interaction clock, which is enough to time the other parts of the register machine simulation.

A curious result in [4] is that even though the register machine simulation has a small probability of error, the same techniques can compute semilinear predicates in polylogarithmic expected parallel time with no error in the limit. The trick is to run a fast error-prone computation to get the answer quickly most of the time, and then switch to the result of a slower, error-free computation using the mechanisms of [3] after some polynomially long interval. The high time to converge for the second algorithm is apparent only when the first fails to produce the correct answer; but as this occurs only with polynomially small probability, it disappears in the expectation.

This simulation leaves room for further improvement. An immediate task is to reduce the overhead of the arithmetic operations. In [6], the same authors show how to drop the cost of the worst-case arithmetic operation to $O(n \log^2 n)$ interactions by combining a more clever register encoding with a fast **approximate majority** primitive based on dueling epidemics. This protocol has only three states: the decision values x and y , and b (for “blank”). When an x token meets a y token or vice versa, the second token turns blank. When an x or y token meets a blank agent, it converts the blank token to its own value. Much of the technical content of [6] involves showing that this process indeed converges to the majority value in $O(n \log n)$ interactions with high probability, which is done using a probabilistic potential function argument separated into several interleaved cases. The authors

suggest that simplifying this argument would be a very useful target for future research. It is also possible that further improvements could reduce the overhead for arithmetic operations down to the $O(n \log n)$ interactions needed simply for all tokens to participate.

A second question is whether the distinguished leader in the initial configuration could be replaced. The coalescing leader election algorithm of [3] takes $\Theta(n^2)$ interactions to converge, which may dwarf the time for simple computations. A heuristic leader-election method is proposed in [6] that appears to converge much faster, but more analysis is needed. The authors also describe a more robust version of the phase clock of [4] that, by incorporating elements of the three-state majority protocol, appears to self-stabilize in $O(n \log n)$ interactions once the number of leaders converges to a polynomial fraction, but to date no proof of correctness for this protocol is known.

7 Self-stabilization and related problems

A series of papers [9, 10, 15] have examined the question of when population protocols can be made self-stabilizing [14], or at least can be made to tolerate input values that fluctuate over some initial part of the computation. Either condition is a stronger property than the mere convergence of the basic model, as both require that the population eventually converge to a good configuration despite an unpredictable initial configuration. Many of the algorithms designed to start in a known initial configuration (even if it is an inconvenient one, with, say, all agents in the same state) will not work if started in a particularly bad one. An example is leader election by coalescence: this algorithm can reduce a population of many would-be leaders down to a single unique leader, but it cannot create a new leader if the initial population contains none.

Angluin *et al.* [9] gave the first self-stabilizing protocols for the population protocol model, showing how to carry out various tasks from previous papers without assuming a known initial configuration. These include a distance-2 coloring protocol for bounded-degree graphs based on local handshaking instead of a wandering colorizer token (which is vulnerable to being lost). Their solution has each node track whether it has interacted with a neighbor of each particular color an odd or even number of times; if a node has two neighbors of the same color, eventually its count will go out of sync with that of one or the other, causing both the node and its neighbor to choose new colors. This protocol is applied in a framework that allows self-stabilizing protocols to be composed, to give additional protocols such as rooted spanning tree construction for networks with a single special node. This last protocol is noteworthy in part because it requires $O(\log D)$ bits of storage per node, where D is the diameter of the network; it is

thus one of the earliest examples of pressure to escape the restrictive $O(1)$ -space assumption of the original population protocol model. Other results in this paper include a partial characterization of which network topologies do or do not support self-stabilizing leader election.

This work was continued by Angluin, Fischer, and Jiang [10], who considered the issue of solving the classic **consensus problem** [24] in an environment characterized by unpredictable communication, with the goal of converging to a common consensus value at all nodes eventually (as in a population protocol) rather than terminating with one. The paper gives protocols for solving consensus in this stabilizing sense with both crash and Byzantine failures. The model used deviates from the basic population protocol model in several strong respects: agents have identities (and the $O(\log n)$ -bit memories needed to store them), and though the destinations to which messages are delivered are unpredictable, communication itself is synchronous.

Fischer and Jiang [15] return to the anonymous, asynchronous, and finite-state world of standard population protocols to consider the specific problem of leader election. As observed above, a difficulty with the simple coalescence algorithm for leader election is that it fails if there is no leader to begin with. Fischer and Jiang propose adding to the model a new **eventual leader detector**, called $\Omega?$, which acts as an oracle that eventually correctly informs the agents if there is no leader. Self-stabilizing leader election algorithms based on $\Omega?$ are given for complete interaction graphs and rings. Curiously, the two cases distinguish between the standard global fairness condition assumed in most population protocol work and a local fairness condition that requires only that each action occurs infinitely often (but not necessarily in every configuration in which it is enabled). The latter condition is sufficient to allow self-stabilizing leader election in a complete graph but is provably insufficient in a ring. Many of these results are further elaborated in Hong Jiang's Ph.D. dissertation [20].

8 Larger states

The assumption that each agent can only store $O(1)$ bits of information is rather restrictive. One direction of research is to slowly relax this constraint to obtain other models that are closer to real mobile systems while still keeping the model simple enough to allow for a complete analysis.

Unique identifiers As noted in Section 2, the requirements that population protocols be independent of n and use $O(1)$ space per agent imply that agents cannot have unique identifiers. This contrasts with the vast majority of models of distributed computing, in which processes do have unique identifiers that are often

a crucial component of algorithms. Guerraoui and Ruppert investigated a model, called **community protocols**, that preserve the tiny nature of agents in population protocols, but allow agents to be initially assigned unique identifiers drawn from a large set [19]. Each agent is equipped with $O(1)$ memory locations that can each store an identifier. It is assumed that transition rules cannot be dependent on the values of the identifiers: the identifiers are atomic objects that can only be tested for equality with one another. (For example, bitwise operations on identifiers are not permitted.) This preserves the property that protocols are independent of n . They gave the following precise characterization of what can be computed in this model.

Theorem 7 ([19]). *A predicate is computable in the community protocol model if and only if it is in $NSPACE(n \log n)$ and permuting the input characters does not affect the output value.*

The necessity of the second condition (symmetry) follows immediately from the fact that the identifiers cannot be used to order the input symbols. The proof that any computable predicate is in $NSPACE(n \log n)$ uses a non-deterministic search of the graph whose nodes are configurations of the community protocol and whose edges represent transitions between configurations.

Conversely, the proof that any symmetric predicate in $NSPACE(n \log n)$ can be computed by a community protocol uses Schönhage's pointer machines [26] as a bridge. A pointer machine is a sequential machine model that runs a program using only a directed graph structure as its memory. A community protocol can emulate a pointer machine by having each agent represent a node in the graph data structure. Some care must be taken to organize the agents to work together to simulate the sequential machine. It was known that a pointer machine that uses $O(n)$ nodes can simulate a Turing machine that uses $O(n \log n)$ space [27].

It follows that the restriction that agents can use their additional memory space only for storing $O(1)$ identifiers can essentially be overcome: the agents can do just as much as they could if they each had $O(\log n)$ bits of storage that could be used arbitrarily.

Heterogeneous systems One interesting direction for future research is allowing some heterogeneity in the model, so that some agents have more computational power than others. As an extreme example, consider a network of weak sensors that interact with one another, but also with a base station that has unlimited capacity.

Beauquier *et al.* [11] studied a scenario like this, focusing on the problem of having the base station compute n , the number of mobile agents. They replaced the fairness condition of the population protocol model by a requirement that all pairs

of agents interact infinitely often. They considered a self-stabilizing version of the model, where the mobile agents are initialized arbitrarily. (Otherwise the problem can be trivially solved by having the base station mark each mobile agent as it is counted.) The problem cannot be solved if each agent's memory is constant size: they proved a tight lower bound of n on the number of possible states the mobile agents must be able to store.

9 Failures

The work described so far assumes that the system experiences no failures. This assumption is somewhat unrealistic in the context of mobile systems of tiny agents, and was made to obtain a clean model as a starting point. Some work has studied fault-tolerant population protocols, although this topic is still largely unexplored.

Crash failures Crash failures are a relatively benign type of failure: faulty agents simply cease having any interactions at some time during the execution. Delporte-Gallet *et al.* [12] examined how crash failures affect the computational power of population protocols. They showed how to transform any protocol that computes a function in the failure-free model into a protocol that can tolerate $O(1)$ crash failures. However, this requires some inevitable weakening of the problem specification.

To understand how the problem specification must change when crash failures are introduced, consider the majority problem described in Example 2. We saw that this problem is solvable if there are no failures. Now consider a version of the majority problem where up to 5 agents may crash. Consider an execution with m followers and $m + 5$ leaders. According to the original problem specification, the output of any such execution must be 1. Suppose, however, that the agents associated with 5 of the $m + 5$ leaders crash before having any interactions. There is no way that the non-faulty agents can distinguish such an execution from a failure-free execution involving m followers and m leaders. In the latter execution, the output must be 0. So, the majority problem, in its original form, cannot be solved when crash failures occur. Nevertheless, we can solve a closely related problem. Suppose that we impose preconditions on the problem, requiring that the margin of the majority is at least 5. More precisely, we impose the requirement on the inputs that either the number of leaders exceeds the number of followers by more than 5 or the number of followers exceeds the number of leaders by at least 5. Under this precondition, it can be shown that the majority problem becomes solvable even when up to 5 agents may crash.

The above example can be generalized in a natural way: If we wish to solve a problem in a way that tolerates up to f crash failures, where f is a constant,

we must impose a precondition that says the removal of f of the input values cannot change the output value. To prove that this is sufficient to make the predicate computable in a fault-tolerant way (assuming that the original predicate is computable in the failure-free model), Delporte-Gallet *et al.* [12] designed an automatic transformation that converts a protocol P for the failure-free model into a protocol P' that will tolerate up to f failures.

The transformation uses replication. In P' , agents are divided (in a fault-tolerant way) into $\Theta(f)$ groups, each of size $\Theta(n/f)$. Each group simulates an execution of P on the entire set of inputs. Each agent of P' can store, in its own memory, the simulated states of $O(f)$ agents of P , since f is a constant, so each group of $\Theta(n/f)$ agents has sufficient memory space to collectively simulate all agents of P . To get a group's simulation started, agents within the group gather the initial states (in P) of *all* agents. Up to f agents may crash before giving their initial states to anyone within that group, but the precondition ensures that this will not affect the output of the simulated run. Thus, any group whose members do not experience any crashes will eventually produce the correct output. The number of groups is chosen to ensure that a majority of the groups run correct simulations, and this allows each agent to determine the correct output.

A variant of the simulation handles a combination of a finite number of transient failures (where an agent spontaneously changes state) and crash failures [12]. It can also be used in the community protocol model described in Section 8 [19].

Byzantine failures An agent that has a Byzantine failure may behave arbitrarily: it can interact with all other agents, pretending to be in any state for each interaction. This behavior can cause havoc in a population protocol since none of the usual techniques used in distributed computing to identify and contain the effects of Byzantine agents can be used. Indeed, it is known that no non-trivial predicate can be computed by a population protocol in a way that tolerates even one Byzantine agent [19]. Two ways of circumventing this fact have been studied.

In the community protocol model of Section 8, if we assume that agent identifiers cannot be tampered with, then some failure detection is possible. Guerraoui and Ruppert give a protocol that solves the majority problem, tolerating a constant number of Byzantine failures, if the margin of the majority is sufficiently wide [19].

Byzantine agents also appear in the random-scheduling work of [6], where it is shown that the approximate majority protocol quickly converges to a configuration in which nearly all non-faulty agents possess the correct decision value despite the actions of a small ($o(\sqrt{n})$) minority of Byzantine agents. Here there is no extension of the basic population protocol model to include identifiers, but the convergence condition is weak, and the Byzantine agents can eventually—after

exponential time—drive the protocol to any configuration, including stable configurations in which no agent holds a decision value. Determining the full power of random scheduling in the presence of Byzantine agents remains open.

References

- [1] D. Angluin, J. Aspnes, M. Chan, M. J. Fischer, H. Jiang, and R. Peralta. Stably computable properties of network graphs. In *Proc. Distributed Computing in Sensor Systems: 1st IEEE International Conference*, pages 63–74, 2005.
- [2] D. Angluin, J. Aspnes, Z. Diamadi, M. J. Fischer, and R. Peralta. Urn automata. Technical Report YALEU/DCS/TR-1280, Yale University Department of Computer Science, Nov. 2003.
- [3] D. Angluin, J. Aspnes, Z. Diamadi, M. J. Fischer, and R. Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, 18(4):235–253, Mar. 2006.
- [4] D. Angluin, J. Aspnes, and D. Eisenstat. Fast computation by population protocols with a leader. In *Proc. Distributed Computing, 20th International Symposium*, pages 61–75, Sept. 2006.
- [5] D. Angluin, J. Aspnes, and D. Eisenstat. Stably computable predicates are semilinear. In *Proc. 25th Annual ACM Symposium on Principles of Distributed Computing*, pages 292–299, 2006.
- [6] D. Angluin, J. Aspnes, and D. Eisenstat. A simple protocol for fast robust approximate majority. In *Proc. Distributed Computing, 21st International Symposium*, pages 20–32, 2007.
- [7] D. Angluin, J. Aspnes, D. Eisenstat, and E. Ruppert. The computational power of population protocols. *Distributed Computing*. To appear; published online in 2007.
- [8] D. Angluin, J. Aspnes, D. Eisenstat, and E. Ruppert. On the power of anonymous one-way communication. In *Proc. Principles of Distributed Systems, 9th International Conference*, pages 396–411, 2005.
- [9] D. Angluin, J. Aspnes, M. J. Fischer, and H. Jiang. Self-stabilizing population protocols. In *Proc. Principles of Distributed Systems, 9th International Conference*, pages 103–117, 2005.
- [10] D. Angluin, M. J. Fischer, and H. Jiang. Stabilizing consensus in mobile networks. In *Proc. Distributed Computing in Sensor Systems, 2nd IEEE International Conference*, pages 37–50, 2006.
- [11] J. Beauquier, J. Clement, S. Messika, L. Rosaz, and B. Rozoy. Self-stabilizing counting in mobile sensor networks. Technical Report 1470, LRI, Université Paris-Sud 11, 2007.

- [12] C. Delporte-Gallet, H. Fauconnier, R. Guerraoui, and E. Ruppert. When birds die: Making population protocols fault-tolerant. In *Proc. 2nd IEEE International Conference on Distributed Computing in Sensor Systems*, pages 51–66, 2006.
- [13] Z. Diamadi and M. J. Fischer. A simple game for the study of trust in distributed systems. *Wuhan University Journal of Natural Sciences*, 6(1–2):72–82, Mar. 2001. Also appears as Yale Technical Report TR–1207, Jan. 2001.
- [14] E. W. Dijkstra. Self-stabilizing systems in spite of distributed control. *Communications of the ACM*, 17(11):643–644, 1974.
- [15] M. J. Fischer and H. Jiang. Self-stabilizing leader election in networks of finite-state anonymous agents. In *Proc. Principles of Distributed Systems, 10th International Conference*, pages 395–409, 2006.
- [16] D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, 81(25):2340–2361, 1977.
- [17] D. T. Gillespie. A rigorous derivation of the chemical master equation. *Physica A*, 188:404–425, 1992.
- [18] S. Ginsburg and E. H. Spanier. Semigroups, Presburger formulas, and languages. *Pacific Journal of Mathematics*, 16:285–296, 1966.
- [19] R. Guerraoui and E. Ruppert. Even small birds are unique: Population protocols with identifiers. Technical Report CSE-2007-04, Department of Computer Science and Engineering, York University, 2007.
- [20] H. Jiang. *Distributed Systems of Simple Interacting Agents*. PhD thesis, Yale University, 2007.
- [21] A. P. Kamath, R. Motwani, K. Palem, and P. Spirakis. Tail bounds for occupancy and the satisfiability threshold conjecture. *Random Structures and Algorithms*, 7:59–80, 1995.
- [22] T. G. Kurtz. *Approximation of Population Processes*. Number 36 in CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics, Philadelphia, 1981.
- [23] M. L. Minsky. *Computation: Finite and Infinite Machines*. Prentice-Hall, Inc., 1967.
- [24] M. Pease, R. Shostak, and L. Lamport. Reaching agreements in the presence of faults. *Journal of the ACM*, 27(2):228–234, Apr. 1980.
- [25] M. Presburger. Über die Vollständigkeit eines gewissen Systems der Arithmetik ganzer Zahlen, in welchem die Addition als einzige Operation hervortritt. In *Comptes-Rendus du I Congrès de Mathématiciens des Pays Slaves*, pages 92–101, Warszawa, 1929.
- [26] A. Schönhage. Storage modification machines. *SIAM J. Comput.*, 9(3):490–508, Aug. 1980.
- [27] P. van Emde Boas. Space measures for storage modification machines. *Inf. Process. Lett.*, 30(2):103–110, Jan. 1989.