

Universal Systems of Oblivious Mobile Robots*

Paola Flocchini[†], Nicola Santoro[‡], Giovanni Viglietta[†], Masafumi Yamashita[§]

[†] University of Ottawa, Canada. paola.flocchini@uottawa.ca, gigliet@uottawa.ca

[‡] Carleton University, Ottawa, Canada. santoro@scs.carleton.ca

[§] Kyushu University, Fukuoka, Japan. mak@csce.kyushu-u.ac.jp

Abstract

An oblivious mobile robot is a stateless computational entity located in a spatial universe, capable of moving in that universe. When activated, the robot observes the universe and the location of the other robots, chooses a destination, and moves there. The computation of the destination is made by executing an algorithm, the same for all robots, whose sole input is the current observation. No memory of all these actions is retained after the move. When the spatial universe is a graph, distributed computations by oblivious mobile robots have been intensively studied focusing on the conditions for feasibility of basic problems (e.g., gathering, exploration) in specific classes of graphs under different schedulers. In this paper, we embark on a different, more general, type of investigation.

With their movements from vertices to neighboring vertices, the robots make the system transition from one configuration to another. Thus the execution of an algorithm from a given configuration defines in a natural way the computation of a discrete function by the system. Our research interest is to understand which functions are computed by which systems. In this paper we focus on identifying sets of systems that are *universal*, in the sense that they can collectively compute all finite functions. We are able to identify several such classes of fully synchronous systems. In particular, among other results, we prove the universality of the set of all graphs with at least one robot, of any set of graphs with at least two robots whose quotient graphs contain arbitrarily long paths, and of any set of graphs with at least three robots and arbitrarily large finite girths.

We then focus on the minimum size that a network must have for the robots to be able to compute all functions on a given finite set. We are able to approximate the minimum size of such a network up to a factor that tends to 2 as n goes to infinity.

The main technique we use in our investigation is the *simulation* between algorithms, which in turn defines *domination* between systems. If a system dominates another system, then it can compute at least as many functions. The other ingredient is constituted by *path* and *ring* networks, of which we give a thorough analysis. Indeed, in terms of implicit function computations, they are revealed to be fundamental topologies with important properties. Understanding these properties enables us to extend our results to larger classes of graphs, via simulation.

1 Introduction

Consider a network, represented as a finite graph G , where the vertices are unlabeled, and edge labels are possibly not unique. In G operate k *oblivious mobile robots* (or simply “robots”), that is,

*This work has been supported in part by the Natural Sciences and Engineering Research Council of Canada through the Discovery Grant program; by Prof. Flocchini’s University Research Chair; and by the Scientific Grant in Aid by the Ministry of Education, Culture, Sports, Science, and Technology of Japan.

indistinguishable computational entities with no memory, located at the vertices of the network, and capable of moving from vertex to neighboring vertex of G . Robots are activated by an adversarial scheduler S . Whenever activated, a robot observes the location of the other robots in the graph (the current *configuration*); it computes a destination (a neighboring vertex or the current location); and it moves there. The computation of the destination is made by executing an algorithm, the same for all robots, whose sole input is the current configuration. The current activity terminates after the move, and no memory of the computation is retained; in other words, the entities are *stateless*. The overall system is represented by the triplet (G, k, S) . Notice that, even if the algorithm A the robots execute is deterministic, its executions may still be non-deterministic. Indeed, since the network's port numbers may not be unique, it may be impossible for an algorithm to unambiguously indicate where each robot has to move. This model, introduced by Klasing, Markou, and Pelc [22] as an extension of the model of oblivious robots in continuous spaces (e.g., [14]), has been extensively employed and investigated, focusing on basic problems in specific classes of graphs under different schedulers: *gathering* and *scattering* (e.g., [4, 5, 6, 7, 10, 16, 17, 19, 21, 22, 25, 26]), and *exploration* and *traversal* (e.g., [1, 2, 3, 8, 9, 11, 12, 13, 23, 24]). Note that, with the exception of [3], the literature assumes unlabelled edges. In this paper, we consider both labelled and unlabelled edges, and focus on the fully synchronous scheduler \mathcal{F} , which simply activates every robot at every turn. We then embark on a different, more general, type of investigation.

Consider the system (G, k, \mathcal{F}) . Whenever the robots move in the graph according to algorithm A , the system transitions from the current configuration to a (possibly) different one. The obliviousness of the robots implies that always the same (or equivalent) transition occurs from the same given configuration. Consider now the *configuration graph* where there is a directed edge from one configuration to another if some algorithm dictates such a transition. Then the execution of A in (G, k, \mathcal{F}) from a given configuration is just a walk in this graph from that configuration. The execution can be viewed in a natural way as the computation of a discrete function f by the system, where f maps a configuration C into the configuration $f(C)$ reached by executing (one step of) A from C , defining a subgraph of the configuration graph, called *function graph*. The concept of function computation and function graph are formally defined in Section 2.

We seek to understand which functions are computed by which systems. Knowing the structure of such functions gives us information on the robots' behavior as they execute an algorithm, and what tasks the robots can and cannot perform in a network. For instance, if an algorithm computes a function whose graph has no cycles, it means that the robots will eventually be stationary regardless of their initial position; if the function has a unique fixed point, it means that the algorithm solves a *pattern formation* problem. On the other hand, if the function's graph has only cycles of length $p > 1$ (possibly with some "branches" attached), the robots are collectively implementing a *self-stabilizing clock* of period p . If such graphs can be embedded in the configuration graph, then we know that such algorithms exist, and that the corresponding problems are solvable in the system.

In this paper we focus on identifying sets of systems that are *universal*, in the sense that they compute all finite functions. In Section 4, we identify several classes of universal fully synchronous systems. In particular, among other results, we prove that

Theorem. *The following families of systems are universal:*

- (a) $\{(G, 1, \mathcal{F}) \mid G \text{ is an unlabeled network}\}$,
- (b) $\{(G_n, 2, \mathcal{F}) \mid \text{the quotient graph of } G_n \text{ contains a sub-path of length at least } n\}$,
- (c) $\{(G_n, 3, \mathcal{F}) \mid \text{the girth of } G_n \text{ is at least } n \text{ and finite}\}$.

In Section 5, we focus on computing discrete functions using the smallest possible networks, perhaps at the cost of employing a large numbers of robots. In particular, for a given finite set X , we study the minimum size that a network must have for the robots to be able to compute all

functions from X to X . We are able to approximate the minimum size of such a network up to a factor that tends to 2 as n goes to infinity.

The main tool we use in our investigation is the *simulation* between algorithms, which in turn defines *domination* between systems. If system Ψ dominates system Ψ' , then Ψ computes at least all the functions computed by Ψ' . The other tool is constituted by the *path* and *ring* graphs (Section 3). These are the main ingredients of all our stronger results, because rings and paths are fundamental topologies with important properties that can be extended to other graphs via simulation.

2 Definitions

In this section we introduce the models of mobile robots that we are going to study. Informally, we consider networks with port numbers, which are represented as graphs where each vertex has a label on each outgoing edge. Port numbers are not required to be unique, which allows us to model anonymous networks with unlabeled edges, as well.

On a network we may place any number of robots, which are indistinguishable mobile entities with no memory. At all times, each robot must be located at a vertex of the network, and any number of robots may occupy the same vertex. All robots follow the same algorithm, which takes as input the network and the robots' positions, and tells each robot to which adjacent vertex it has to move next (or it may tell it to stay still). Time is discretized, and we assume that robots can move to adjacent vertices instantaneously.

Even if algorithms are deterministic, their executions may still be non-deterministic. This is partly because the network's port numbers may not be unique, and therefore it may be impossible for an algorithm to unambiguously indicate where each robot has to move. Another potential source of non-determinism is the scheduler, which is an adversary that decides which robots are going to be activated next. In this paper we will focus on the fully synchronous scheduler, which simply activates every robot at every turn. We will also briefly discuss the semi-synchronous scheduler in Section 6.

Labeled Graphs. A *labeled graph* is a triplet $G = (V, E, \ell)$, where (V, E) is an undirected graph called the *base graph*, and ℓ is a function that maps each ordered pair (u, v) , such that $\{u, v\} \in E$, to a non-negative integer called *label*. A labeled graph is also referred to as a *network*. A network is *unlabeled* if all its labels are equal.

An *automorphism* of a labeled graph $G = (V, E, \ell)$ is a permutation α of V preserving adjacencies and labels, i.e., for all $u, v \in V$, if $\{u, v\} \in E$, then $\{\alpha(u), \alpha(v)\} \in E$ and $\ell(u, v) = \ell(\alpha(u), \alpha(v))$. If there exists an automorphism that maps a vertex u to a vertex v , then u and v are *equivalent vertices* in G . The *quotient graph* G^* is the labeled graph G obtained by identifying equivalent vertices, and preserving adjacencies and labels.

Configuration Spaces. Let $\mathbb{N}_n = \{0, 1, \dots, n-1\}$, for every $n \geq 1$. An *arrangement* of k robots on a network $G = (V, E, \ell)$ is a mapping from \mathbb{N}_k to V . An arrangement specifies the locations of k *distinguishable* robots on a network whose vertices are all *distinguishable*. However, we ultimately intend to model *identical* robots, which *cannot distinguish* between equivalent vertices of the network, unless such vertices are occupied by different amounts of robots. The following definition serves this purpose: two arrangements $a_1, a_2: \mathbb{N}_k \rightarrow V$, are *equivalent* if there exist an automorphism $\alpha: V \rightarrow V$ and a permutation π of \mathbb{N}_k such that $\alpha \circ a_1 = a_2 \circ \pi$.

The *configuration space* $\mathcal{C}(G, k)$, where G is a network and k is a positive integer, is the quotient of the set of arrangements of k robots on G under the above equivalence relation between

arrangements. The elements of the configuration space are called *configurations*.

Say that an arrangement a is equivalent to itself under an automorphism α and a permutation π , as defined above. Then, whenever $\alpha(v) = v'$ and $\pi(r) = r'$, we say that v and v' are *equivalent vertices* in a , and r and r' are *equivalent robots* in a .

A class of *indistinguishable* vertices U (respectively, a class of *indistinguishable* robots R) of a configuration $C \in \mathcal{C}(G, k)$ is a mapping from each arrangement $a \in C$ to an equivalence class of vertices U_a (respectively, an equivalence class of robots R_a) of a such that, for all $a_1, a_2 \in C$ and all automorphisms α and permutations π under which a_1 and a_2 are equivalent, $\alpha(U_{a_1}) = U_{a_2}$ (respectively, $\pi(R_{a_1}) = R_{a_2}$).

Configuration Graphs. While the configuration space contains all the configurations that are distinguishable, either by the base graph's topology, or by the labels, or by the robots' positions, the *configuration graph* specifies which configurations can reach which other configurations "in one step". Of course, this depends on a notion of algorithm, and on a notion of scheduler.

An *algorithm* for k robots on a network G is a function that maps a pair (C, U) into a set U' , where $C \in \mathcal{C}(G, k)$ (describing the network's configuration at the moment the algorithm is executed), and U and U' are classes of indistinguishable vertices of C (indicating the executing robot's location and its destination, respectively) such that, for every arrangement $a \in C$ and every vertex $u \in U(a)$, there exists a vertex $u' \in U'(a)$ such that either $u = u'$ or u' is adjacent to u . According to this definition, a robot can only specify its destination as a class of indistinguishable vertices, representing either a null movement or a movement to some adjacent vertex.

An *execution* for k robots in a network G is a sequence of configurations of $\mathcal{C}(G, k)$. A *scheduler* for k robots in a network G is a binary relation between algorithms and executions. The *possible* executions of an algorithm under some scheduler are the executions that correspond to the algorithm under the relation specified by such a scheduler. A *system of oblivious mobile robots* is a triplet $\Psi = (G, k, S)$, where G is a labeled graph, $k \geq 1$, and S is a scheduler for k robots in G .

The *configuration graph* $\mathcal{G}(\Psi) = (\mathcal{C}(G, k), \mathcal{E}(\Psi))$, where $\Psi = (G, k, S)$ is a system of oblivious mobile robots, is a directed graph on the configuration space $\mathcal{C}(G, k)$, where $(C, C') \in \mathcal{E}(G, k)$ if there is an algorithm A and a possible execution $E = (C_i)_{i \geq 0}$ of A under S , such that there exists an index i satisfying $C = C_i$ and $C' = C_{i+1}$.

The *deterministic configuration graph* $\mathcal{G}'(\Psi) = (\mathcal{C}(G, k), \mathcal{E}'(\Psi))$, where $\Psi = (G, k, S)$ is a system of oblivious mobile robots, is a directed graph on the configuration space $\mathcal{C}(G, k)$, where $(C, C') \in \mathcal{E}'(G, k)$ if there is an algorithm A such that, for all possible executions $E = (C_i)_{i \geq 0}$ of A under S , and for every index i satisfying $C = C_i$, we have $C' = C_{i+1}$.

Intuitively, $\mathcal{G}'(\Psi)$ is a subgraph of $\mathcal{G}(\Psi)$ whose edges represent moves that can be deterministically done by the robots, i.e., on which all the scheduler's choices yield the same result. If $\mathcal{G}(\Psi) = \mathcal{G}'(\Psi)$, then Ψ is said to be a *deterministic* system.

Fully Synchronous Scheduler. Given an algorithm A for k robots on a network, we say that a configuration C' *yields* from a configuration C under algorithm A if, for every arrangement $a \in C$ there is an arrangement $a' \in C'$ such that, for every $r \in \mathbb{N}_k$, either $a(r) = a'(r)$ or $a(r)$ is adjacent to $a'(r)$ and, if U is the class of indistinguishable vertices of C such that $a(r) \in U(a)$, then $a'(r) \in U'(a)$, where $U' = A(C, U)$. The *fully synchronous scheduler* \mathcal{F} is defined as follows: $(A, E = (C_i)_{i \geq 0}) \in \mathcal{F}$ if, for every $i \geq 0$, C_{i+1} yields from C_i . In the rest of the paper, we will write $\mathcal{F}(G, k)$ instead of (G, k, \mathcal{F}) .

In other words, the fully synchronous scheduler lets every robot move at every turn to the destination it computes. However, if a robot's destination consists of several indistinguishable

vertices, the scheduler may arbitrarily decide to move the robot to any of those vertices, provided that it can be reached in at most one hop. All these choices are made by the scheduler at each turn and for each robot, independently.

Simulating Algorithms. To define the concept of simulation, we preliminarily define a relation on executions. Given an execution $E = (C_i)_{i \geq 0}$ for k robots on a network G , an execution $E' = (C'_i)_{i \geq 0}$ for k' robots on a network G' , and a surjective partial function $\varphi: \mathcal{C}(G, k) \rightarrow \mathcal{C}(G', k')$, we say that E is *compliant* with E' under φ if either φ is undefined on C_0 , or there exists a weakly increasing surjective function $\sigma: \mathbb{N} \rightarrow \mathbb{N}$ such that, for every $i \in \mathbb{N}$, φ is defined on C_i , and $\varphi(C_i) = C'_{\sigma(i)}$.

An algorithm A under system Ψ *simulates* an algorithm A' under system Ψ' if there is a surjective partial function $\varphi: \mathcal{C}(G, k) \rightarrow \mathcal{C}(G', k')$ such that each execution of A under Ψ is compliant under φ with at least one execution of A' under Ψ' .

In this definition, φ “interprets” some configurations of the simulating system Ψ as configurations of the simulated system Ψ' , in such a way that every configuration of Ψ' is represented by at least one configuration of Ψ . Moreover, the definition of compliance ensures that the simulating algorithm A makes configurations transition in a way that agrees with A' under φ .

Computing functions. We define the implicit computation of a function as the simulation of a system consisting in a single robot on a network whose shape is given by the function itself.

The network *induced* by a function $f: X \rightarrow X$ is defined as $\Gamma_f = (X, f, \ell)$, where $\ell: (u, v) \mapsto v$. Hence the base graph of Γ_f has edges of the form $(x, f(x))$, and the labeling ℓ makes all vertices of Γ_f distinguishable from each other. The algorithm A_f *associated* to the function f is the algorithm for one robot on Γ_f that always makes the robot move from any vertex $x \in X$ to the vertex $f(x)$.

We say that an algorithm A *computes* a function $f: X \rightarrow X$ under system Ψ if it simulates the algorithm A_f under $\mathcal{F}(\Gamma_f, 1)$.

What this definition intuitively means is that each element of X is represented by a set of robot configurations; an algorithm computes f if any execution from a configuration representing $x \in X$ eventually yields a configuration representing $f(x)$ without passing through configurations that represent other elements of X (or that represent no element of X).

If an algorithm under system Ψ computes a function f (respectively, simulates an algorithm A'), then we say that Ψ computes f (respectively, simulates A'). Moreover, a system Ψ *dominates* Ψ' if every algorithm under Ψ' is simulated by some algorithm under Ψ .

We use the notation $X \preceq Y$ to indicate all the concepts defined above: X may be a function computed by an algorithm Y (under some system), or it can be an algorithm simulated by Y , or a system dominated by a system Y , etc.

3 Basic Results

Proposition 1. *The relation \preceq is transitive.*

Proof. Let us prove that, if $A'' \preceq A'$ and $A' \preceq A$, then $A'' \preceq A$, where A is an algorithm under system $\Psi = (G, k, S)$, A' is an algorithm under system $\Psi' = (G', k', S')$, and A'' is an algorithm under system $\Psi'' = (G'', k'', S'')$. All other cases (i.e., those involving functions or systems instead of only algorithms) trivially follow from this basic case.

Let the surjective partial function $\varphi: \mathcal{C}(G, k) \rightarrow \mathcal{C}(G', k')$ express the fact that A simulates A' , and let the surjective partial function $\varphi': \mathcal{C}(G', k') \rightarrow \mathcal{C}(G'', k'')$ express the fact that A' simulates

A'' . Let us show that $\mu = \varphi' \circ \varphi$, which is a surjective partial function from $\mathcal{C}(G, k)$ to $\mathcal{C}(G'', k'')$, expresses the fact that A simulates A'' .

Let $E = (C_i)_{i \geq 0}$ be an execution of A under Ψ . If μ is undefined on C_0 , then E complies with any execution of A'' , by definition. So, let us assume that μ is defined on C_0 , and therefore φ is defined on C_0 , and φ' is defined on $\varphi(C_0)$. This implies that, for all $i \geq 0$, φ is defined on C_i , by definition of simulation. Moreover, E is compliant under φ with some execution $E' = (C'_i)_{i \geq 0}$ of A' under Ψ' . So there is a weakly increasing surjective function $\sigma: \mathbb{N} \rightarrow \mathbb{N}$ such that $\varphi(C_i) = C'_{\sigma(i)}$.

In particular, $\varphi(C_0) = C'_{\sigma(0)} = C'_0$. But we know that φ' is defined on $\varphi(C_0)$, and hence on C'_0 . By definition of computation, φ' is therefore defined on each C'_i , with $i \geq 0$. Moreover, E' is compliant under φ' with some execution $E'' = (C''_i)_{i \geq 0}$ of A'' under Ψ'' . So there is a weakly increasing surjective function $\sigma': \mathbb{N} \rightarrow \mathbb{N}$ such that $\varphi'(C'_i) = C''_{\sigma'(i)}$.

Now, let $i \in \mathbb{N}$ be an index. By the above, μ is defined on C_i , and

$$\mu(C_i) = \varphi'(\varphi(C_i)) = \varphi'(C'_{\sigma(i)}) = C''_{\sigma'(\sigma(i))}.$$

Since $\sigma' \circ \sigma: \mathbb{N} \rightarrow \mathbb{N}$ is itself a weakly increasing surjective function, this implies that E is compliant with E'' under μ , and hence A simulates A'' . \square

Corollary 2. *If a system Ψ dominates a system Ψ' , then all functions computed by Ψ' are also computed by Ψ .*

Proof. Suppose that $\Psi' \preceq \Psi$. Then, for any function f such that $f \preceq \Psi'$, the transitivity of \preceq implies that $f \preceq \Psi$. \square

3.1 General Graphs

Proposition 3. *For every network G , the system $\mathcal{F}(G, 1)$ is deterministic, and its configuration graph is isomorphic to the graph obtained from the quotient graph G^* by replacing each unoriented edge $\{u, v\}$ with the two oriented edges (u, v) and (v, u) , and adding a self-loop (v, v) to each vertex v .*

Proof. If there is just one robot in the network, all non-deterministic choices of a neighboring vertex among equivalent ones yield equivalent configurations. This proves that $\mathcal{F}(G, 1)$ is deterministic.

Note that the arrangements of one robot on G consist of all the mappings of the form $0 \mapsto v$, with v a vertex of G . For each such v , let C_v be the configuration corresponding to the arrangement $0 \mapsto v$, and let $[v]$ be the equivalence class of G^* such that $v \in [v]$. By definition of equivalence, for any two arrangements $0 \mapsto u$ and $0 \mapsto v$, $C_u = C_v$ if and only if v and v' are equivalent in G . This implies that the relation φ that maps C_v to $[v]$ for every vertex v is a well-defined bijection between the configuration space $\mathcal{C}(G, 1)$ and the vertex set of G^* .

Let us prove that φ induces a graph isomorphism, if the edges of G^* are modified as in the observation's statement. Let (C_u, C_v) be an edge of the configuration graph of $\mathcal{F}(G, 1)$. By definition of fully synchronous scheduler, this holds if and only if there is an algorithm A_{uv} that takes C_u and the class of indistinguishable vertices of C_u containing u , and maps it to the class of indistinguishable vertices of C_u containing v . Such an algorithm exists if and only if there is a vertex u' indistinguishable from u in C_u and a vertex v' indistinguishable from v in C_u such that either $u' = v'$ or u' and v' are neighbors in G . However, since there is a unique robot in the network, $u = u'$, and v and v' have the same distance from u . In other words, either $[u] = [v]$, or u and v are non-equivalent neighbors in G . In the first case, the edge (C_u, C_v) is correctly mapped into the self-loop on $[u]$ in the modified G^* . The second case is equivalent to $[u]$ and $[v]$ being adjacent in G^* , which is true if and only if there is the directed edge $([u], [v])$ in the modified G^* . \square

A fundamental question is whether adding robots to a network allows to compute more functions. We can at least prove that adding robots does not reduce the set of computable functions, provided that the network is not pathologically small.

Theorem 4. *For all networks G with at least three vertices and all $k \geq 1$, $\mathcal{F}(G, k+1) \not\preceq \mathcal{F}(G, k)$.*

Proof. It suffices to show that $|\mathcal{C}(G, k+1)| > |\mathcal{C}(G, k)|$. For each configuration in $\mathcal{C}(G, k)$, choose a vertex that contains the largest number of robots, and add one robot to it. This way we obtain $|\mathcal{C}(G, k)|$ distinct configurations of $\mathcal{C}(G, k+1)$. We can generate yet another configuration by placing $\lfloor (k+1)/2 \rfloor$ robots on a vertex, $\lfloor (k+1)/2 \rfloor$ robots on another vertex, and the remainder on a third vertex. \square

We can also show that a single robot does not compute more functions than $k \geq 1$ robots, in any network G .

Theorem 5. *For all networks G and all $k \geq 1$, $\mathcal{F}(G, 1) \preceq \mathcal{F}(G, k)$.*

Proof. In the simulation we use only the configurations of $\mathcal{F}(G, k)$ in which all robots lie in equivalent vertices of G . Then each robot pretends to be the only robot in the network, and makes the move that the unique robot of $\mathcal{F}(G, 1)$ would make. This is a well-defined simulation even if $\mathcal{F}(G, k)$ is not deterministic, due to Proposition 3. \square

We can extend this idea to show that $\mathcal{F}(G, k) \preceq \mathcal{F}(G, 2k)$, provided that $\mathcal{F}(G, 2k)$ is deterministic.

Theorem 6. *$\mathcal{F}(G, k) \preceq \mathcal{F}(G, k')$, provided that $\mathcal{F}(G, k')$ is deterministic and $k' \geq 2k$.*

Proof. The configurations of $\mathcal{F}(G, k')$ that we use in our simulation are only those in which there is a (unique) vertex v occupied by at least $k' - k + 1$ robots. Each of these configurations is mapped to the configuration of $\mathcal{F}(G, k)$ that is obtained by removing $k' - k$ robots from v . This mapping is surjective. The simulation can be carried out because $\mathcal{F}(G, k')$ is deterministic, and therefore robots occupying the same vertex can never be separated, implying that there is always going to be a vertex with at least $k' - k + 1$ robots. \square

Finally, we conjecture that adding robots increases a system's computational capabilities.

Conjecture 7. *For all networks G and all $k \geq 1$, $\mathcal{F}(G, k) \preceq \mathcal{F}(G, k+1)$.*

3.2 Path Graphs

The first special type of network we consider is the one whose base graph consists of a single *path*. This fundamental configuration will turn out to be of great importance in Sections 4 and 5, when studying universal classes of systems. In terms of labeling, we focus on two extreme cases: a labeling that gives a consistent orientation to the whole network (i.e., each vertex in the path has port labels indicating which neighbor is on the “left” and which one is on the “right”), and the anonymous unlabeled network. In the first case we have an *oriented path*, and in the second case we have an *unoriented path*. By $\vec{\mathcal{P}}_n^k$ and \mathcal{P}_n^k we denote, respectively, the oriented and the unoriented path with n vertices and k robots, under the fully synchronous scheduler.

Oriented Paths. Let us study the configuration graph of $\vec{\mathcal{P}}_n^k$. Since the path has an orientation, no two vertices are equivalent. Therefore, by Proposition 3, $\mathcal{G}(\vec{\mathcal{P}}_n^1)$ consists of a path of length n with bidirectional edges and a self-loop on each vertex. In general, the configuration space of $\vec{\mathcal{P}}_n^k$ is in bijection with the set of weakly increasing k -tuples of integers in \mathbb{N}_n . Hence, for a fixed k , the size of the configuration space is

$$\binom{n+k-1}{k} \sim \frac{n^k}{k!}.$$

If these k -tuples are thought of as points of \mathbb{R}^k , they constitute the set of lattice points in the k -dimensional simplex whose $k+1$ vertices have the form $(0, 0, \dots, 0, 1, 1, \dots, 1)$. This simplex has $k+1$ facets, two of which correspond to configurations in which the first or the last vertex of the network is occupied by a robot, while the other $k-1$ facets correspond to configurations in which exactly $k-1$ vertices are occupied (i.e., exactly two robots share a vertex).

The edges of the configuration graph (that are not self-loops) connect bidirectionally all pairs of points whose Chebyshev distance is at most 1, with the exception of the points that lie on the aforementioned $k-1$ facets. Indeed, since no algorithm can separate two robots that occupy the same vertex, it follows that those facets (as well as all their intersections) can never be left once they are reached. Figure 1(a) shows the configuration graph of $\vec{\mathcal{P}}_5^2$.

Property 8. For all $n \geq 1$, $\mathcal{G}(\vec{\mathcal{P}}_{2n}^2)$ contains an $n \times n$ grid with bidirectional edges and self-loops.

Since an oriented path gives the robots a sense of direction, an algorithm can unambiguously indicate to which neighbor each robot is supposed to move. Therefore $\mathcal{G}(\vec{\mathcal{P}}_n^k) = \mathcal{G}'(\vec{\mathcal{P}}_n^k)$.

Property 9. For all $n, k \geq 1$, the system $\vec{\mathcal{P}}_n^k$ is deterministic.

Unoriented Paths. Let us study the configuration graph of \mathcal{P}_n^k . Since the network in this system is unlabeled, if two vertices are symmetric with respect to the center of the path, they are equivalent. So, the configuration space is in bijection with the set of weakly increasing k -tuples of integers in \mathbb{N}_n , where each k -tuple (a_1, \dots, a_k) is identified with its ‘‘symmetric’’ one, $(n - a_k - 1, \dots, n - a_1 - 1)$. Elementary computations reveal that, if k is fixed, these k -tuples are

$$\frac{1}{2} \cdot \binom{n+k-1}{k} + O(n^{\lfloor k/2 \rfloor}) \sim \frac{n^k}{2 \cdot k!}.$$

Geometrically, the configuration space of \mathcal{P}_n^k can be represented as the set of lattice points in a truncated k -dimensional simplex, which is obtained by cutting the simplex of $\vec{\mathcal{P}}_n^k$ roughly in half, along a suitable hyperplane. Figures 1(b) and 1(c) show the configuration graphs of \mathcal{P}_5^2 and \mathcal{P}_6^2 .

If n is even, we have $\mathcal{G}(\mathcal{P}_n^k) = \mathcal{G}'(\mathcal{P}_n^k)$, because each robot has a unique closest endpoint of the path, which it can use to specify unambiguously in which direction it intends to move. However, if $n > 1$ is odd and $k \geq 2$, the two graphs differ. Indeed, if the configuration is symmetric and the central vertex is occupied by more than one robot, then it is impossible to guarantee that all the central robots will move in the same direction: the adversary will decide how many of these robots go left, and how many go right. For instance, $\mathcal{G}'(\mathcal{P}_5^2)$ differs from $\mathcal{G}(\mathcal{P}_5^2)$ in that the vertex in $(2, 2)$ has no outgoing edges in $\mathcal{G}'(\mathcal{P}_5^2)$, because these correspond to non-deterministic moves.

Property 10. For all $n, k \geq 1$, the system \mathcal{P}_{2n}^k is deterministic.




Figure 1: Configuration graphs of some oriented and unoriented paths and rings. Dashed arrows represent non-deterministic moves. For clarity, self-loops have been omitted from all vertices.

3.3 Ring Graphs

Now we consider *ring* networks, which are networks whose base graph is a single cycle. This is another fundamental class of networks, which will have a great importance in Sections 4 and 5. Like a path, a ring can be *oriented* if its labeling gives a consistent sense of direction to the robots in the network (i.e., each vertex has port labels indicating which neighbor lies in the “clockwise” direction, and which one lies in the “counterclockwise” direction), and *unoriented* if the network is unlabeled. Therefore we have the two systems $\vec{\mathcal{R}}_n^k$ and \mathcal{R}_n^k , denoting, respectively, the oriented and the unoriented ring with n vertices and k robots, under the fully synchronous scheduler.

Oriented Rings. Let us study the structure of $\mathcal{G}(\vec{\mathcal{R}}_n^k)$. Note that, in a ring network, all vertices are equivalent. Therefore, by Proposition 3, $\mathcal{G}(\vec{\mathcal{R}}_n^1)$ consists of a single vertex with a self-loop.

In the case of $k = 2$ robots, a configuration is uniquely identified by the distance d of the two robots on the ring, which may be any integer between 0 and $\lfloor n/2 \rfloor$. If $d = 0$, the robots are bound to remain on the same vertex. If $d = n/2$ (hence n is even), the robots are located on indistinguishable vertices, and they are bound to remain on indistinguishable vertices. In all other cases, it is possible to distinguish the two robots and move them independently, thanks to the orientation of the ring. Therefore, if $0 < d < n/2$, an algorithm may move the two robots independently in any direction, thus adding any integer between -2 and $+2$ to d (subject to the $0 \leq d \leq \lfloor n/2 \rfloor$ constraint). Figures 1(d) and 1(e) show the configuration graphs of $\vec{\mathcal{R}}_{14}^2$ and $\vec{\mathcal{R}}_{15}^2$.

In general, the configuration space of $\vec{\mathcal{R}}_n^k$ is in bijection with the set of binary *necklaces* of length n and density k , i.e., the binary strings having k zeros and $n - k$ ones, taken modulo rotations. To count them, the Pólya enumeration theorem can be applied, as in [15]. If k is fixed, the configuration space has size

$$\frac{1}{n+k} \cdot \sum_{d|\gcd(k,n)} \phi(d) \cdot \binom{(n+k)/d}{k/d} \sim \frac{n^{k-1}}{k!},$$

where ϕ is Euler's totient function. Since the ring is oriented, we have $\mathcal{G}(\vec{\mathcal{R}}_n^k) = \mathcal{G}'(\vec{\mathcal{R}}_n^k)$.

Property 11. *For all $n, k \geq 1$, the system $\vec{\mathcal{R}}_n^k$ is deterministic.*

Unoriented Rings. The structure of $\mathcal{G}(\mathcal{R}_n^k)$ is similar to that of $\mathcal{G}(\vec{\mathcal{R}}_n^k)$, except that now two configurations are indistinguishable also if they are “reflections” of each other. The case $k = 1$ is again trivial and yields a single configuration, but the case $k = 2$ is more interesting. As before, a configuration of \mathcal{R}_n^2 is identified by an integer d with $0 \leq d \leq \lfloor n/2 \rfloor$, but the two robots are now indistinguishable, and therefore they must always make symmetric moves. Hence d may only change by -2 or $+2$; the only exception is when n is odd and $d = \lfloor n/2 \rfloor$, which can change to $d = \lfloor n/2 \rfloor - 1$ (as well as to $d = \lfloor n/2 \rfloor - 2$), and vice versa. So, if n is odd, $\mathcal{G}(\mathcal{R}_n^2)$ is isomorphic to $\mathcal{G}(\vec{\mathcal{P}}_{\lfloor n/2 \rfloor}^1)$. If n is even, $\mathcal{G}(\mathcal{R}_n^2)$ consists of two connected components: the one corresponding to even d 's is isomorphic to $\mathcal{G}(\vec{\mathcal{P}}_{\lfloor (n+2)/4 \rfloor}^1)$, and the one corresponding to odd d 's is isomorphic to $\mathcal{G}(\vec{\mathcal{P}}_{\lfloor (n+2)/4 \rfloor}^1)$. Figures 1(f), 1(g), and 1(h) show the configuration graphs of \mathcal{R}_{14}^2 , \mathcal{R}_{15}^2 , and \mathcal{R}_{16}^2 .

Property 12. *For all $n, k \geq 1$, $\mathcal{G}(\mathcal{R}_n^k)$ consists of either a single path or two disjoint paths.*

In general, instead of representing the configurations of \mathcal{R}_n^k with necklaces as before, we use *bracelets* of length n and density k , i.e., binary strings having k zeros and $n - k$ ones, taken modulo rotations and reflections. The size of the configuration space can be computed again with the Pólya enumeration theorem, this time using the dihedral group instead of the cyclic group. For fixed k , its size is

$$\frac{1}{2(n+k)} \cdot \sum_{d|\gcd(k,n)} \phi(d) \cdot \binom{(n+k)/d}{k/d} + O(n^{\lfloor k/2 \rfloor}) \sim \frac{n^{k-1}}{2 \cdot k!}.$$

With unoriented rings, the deterministic configuration graph is slightly different. If $d = 0$ or $d = n/2$, the adversary may choose to keep d unvaried, by making both robots always move in the same direction. Therefore, the configurations corresponding to $d = 0$ and $d = n/2$ have no outgoing edges in $\mathcal{G}'(\mathcal{R}_n^k)$. Other than that, the two graphs are the same.

3.4 Domination Relations Between Paths and Rings

Next we describe how different systems of paths and rings dominate each other.

Theorem 13. *For all $n, k \geq 1$ and $k' \geq 2k$, $\vec{\mathcal{P}}_n^k \preceq \vec{\mathcal{P}}_n^{k'}$, $\mathcal{P}_{2n}^k \preceq \mathcal{P}_{2n}^{k'}$, and $\vec{\mathcal{R}}_n^k \preceq \vec{\mathcal{R}}_n^{k'}$.*

Proof. By Properties 9, 10, and 11, the systems $\vec{\mathcal{P}}_n^{k'}$, $\mathcal{P}_{2n}^{k'}$, and $\vec{\mathcal{R}}_n^{k'}$ are deterministic. Thus all relations follow from Theorem 6. \square

Theorem 14. *For all $n, k \geq 1$, $\vec{\mathcal{P}}_n^k \preceq \mathcal{P}_{2n}^k$.*

Proof. We place the k robots of \mathcal{P}_{2n}^k on the first n vertices of the path, leaving the other half empty. This gives an implicit orientation to the path, and now the simulation is trivial. \square

Theorem 15. *For all $n \geq 1$ and $k \geq 2$, $\vec{\mathcal{P}}_n^k \preceq \vec{\mathcal{R}}_{2n}^{k+1}$.*

Proof. We use only the configurations of $\vec{\mathcal{R}}_{2n}^{k+1}$ having a vertex occupied by a single robot, followed in clockwise order by $n - 1$ empty vertices. The simulation is performed by the other k robots on the remaining n vertices, which are always unambiguously identified because $k \geq 2$. \square

Theorem 16. For all $n \geq 1$ and $k \geq 2$, $\mathcal{P}_n^k \preceq \mathcal{R}_{3n-1}^{k+1}$.

Proof. We use only the configurations of \mathcal{R}_{3n-1}^{k+1} having a vertex occupied by one robot, surrounded on both sides by sequences of $n - 1$ empty vertices. The simulation is performed by the other k robots on the remaining n vertices, which are always unambiguously identified because $k \geq 2$. \square

4 Universality

A system Ψ is *universal* for \mathbb{N}_n if it computes every function on \mathbb{N}_n . In this case, we write $\mathbb{N}_n \preceq \Psi$. Note that this extension of the relation \preceq preserves its transitivity. A set of systems Υ is *universal* if, for every $n \geq 1$ and every function $f: \mathbb{N}_n \rightarrow \mathbb{N}_n$, there is a system $\Psi \in \Upsilon$ that computes f .

One robot is sufficient for universality, even on unlabeled networks.

Theorem 17. $\{\mathcal{F}(G, 1) \mid G \text{ is an unlabeled network}\}$ is universal.

Proof. Given a function $f: \mathbb{N}_n \rightarrow \mathbb{N}_n$, take the complete graph K_n and attach i dangling vertices to its i -th vertex, for all $0 \leq i < n$. To compute f , instruct the robot to move from the vertex with i dangling vertices to the one with $f(i)$ dangling vertices (which is always distinguishable). \square

However, complete graphs are very demanding networks. If no vertex in the network has more than two neighbors, universality requires more robots: two for paths and oriented rings, and three for unoriented rings.

Theorem 18. $\{\vec{\mathcal{P}}_n^1 \mid n \geq 1\}$, $\{\mathcal{P}_n^1 \mid n \geq 1\}$, $\{\vec{\mathcal{R}}_n^1 \mid n \geq 1\}$, and $\{\mathcal{R}_n^2 \mid n \geq 1\}$ are not universal.

Proof. This is trivially true for $\{\vec{\mathcal{R}}_n^1 \mid n \geq 1\}$, since $\vec{\mathcal{R}}_n^1$ has only one configuration.

Let the *cycle function* $\lambda_m: \mathbb{N}_m \rightarrow \mathbb{N}_m$ be defined as $\{(i, i + 1) \mid i \in \mathbb{N}_{m-1}\} \cup \{(m - 1, 0)\}$. Note that, if a system computes λ_m , its configuration graph must contain a cycle of length at least m . By Proposition 3 and Property 12, the configuration graphs of $\vec{\mathcal{P}}_n^1$, \mathcal{P}_n^1 , and \mathcal{R}_n^2 consist of either a single path or two disjoint paths. It follows that none of these systems can compute λ_m , for any $m \geq 3$. \square

Next we show that two robots are indeed sufficient to compute every function on arbitrarily long oriented and unoriented paths and oriented rings, and that three robots are sufficient for unoriented rings. First we introduce a definition: an algorithm is *sequential* if it never instructs two robots to move at the same time.

Lemma 19. For all $n \geq 1$, $\mathbb{N}_n \preceq \vec{\mathcal{P}}_{2n}^2$. The algorithms used in the computations are sequential.

Proof. Let $f: \mathbb{N}_n \rightarrow \mathbb{N}_n$, and consider the base graph G of the network induced by f . Each connected component of G consists of a single cycle (possibly degenerating in a vertex with a self-loop) to which some disjoint trees are attached. Each tree is rooted at the vertex that it shares with the cycle, and all edges in the tree are directed towards the root. In particular, G is planar.

We would like to reduce the maximum degree of G to 3. We do so by a series of local transformations of G , as shown in Figure 2. If G has a vertex v of degree greater than 3, we transform it as follows. If (u, v) and (v, w) are edges of G , we delete them, we create a new vertex v_u , and we add the edges (v, v_u) , (u, v_u) , and (v_u, w) . This transformation creates a new vertex of degree 3, and decreases by 1 the degree of v . If v has initial indegree $d > 3$, this operation is performed on v exactly $d - 3$ times, adding $d - 3$ vertices to the graph.




Figure 2: Reducing the degree of v .

Therefore, by repeatedly applying this transformation, we obtain a new planar graph G' on at most $2n$ vertices, with maximum degree at most 3, such that G is a graph minor of G' . As shown in [20], G' admits a planar orthogonal grid drawing on an $n \times n$ grid. So, G' is a minor of some graph G'' (obtained by adding vertices on some edges of G') that in turn is a subgraph of an $n \times n$ grid with self-loops and bidirectional edges.

By Properties 8 and 9, this $n \times n$ grid is a subgraph of $\mathcal{G}'(\vec{\mathcal{P}}_{2n}^2)$. It is easy to check that $\vec{\mathcal{P}}_{2n}^2$ computes f under the surjective partial function induced by the composition of these graph minor relations.

Indeed, whenever the computation reaches a configuration corresponding to a vertex of G that has been splitted during the first set of transformations, it just follows the edges generated during the splitting. This is possible because all these edges represent deterministic moves, so they can be chosen by the algorithm and have to be complied by the fully synchronous scheduler.

On the other hand, if the computation reaches a configuration corresponding to an edge of G that has been extended in order to embed it in the grid, the computation follows the extended edge. Again, this is possible because these edges represent deterministic moves. \square

Theorem 20. $\{\vec{\mathcal{P}}_n^2 \mid n \geq 1\}$, $\{\mathcal{P}_n^2 \mid n \geq 1\}$, $\{\vec{\mathcal{R}}_n^2 \mid n \geq 1\}$, and $\{\mathcal{R}_n^3 \mid n \geq 1\}$ are universal.

Proof. For $\{\vec{\mathcal{P}}_n^2 \mid n \geq 1\}$, $\{\mathcal{P}_n^2 \mid n \geq 1\}$, and $\{\mathcal{R}_n^3 \mid n \geq 1\}$, this immediately follows from Lemma 19, Theorems 14 and 16, and the transitivity of \preceq .

For oriented rings, we need a different argument. The shape of the configuration graph of $\vec{\mathcal{R}}_n^2$ is illustrated in Figures 1(d) and 1(e). Given a function f whose induced network is Γ_f , we can embed Γ_f in $\mathcal{G}(\vec{\mathcal{R}}_n^2)$, provided that n is large enough, for instance as shown in Figure 3. In general, each connected component of the base graph of Γ_f is a cycle L with some trees attached to it. for each leaf of such a tree, connected to L via a path P , we embed a copy of L in $\mathcal{G}(\vec{\mathcal{R}}_n^2)$ and attach a copy of P to it, mapping each vertex to the appropriate vertex of $L \cup P$. The resulting surjective partial function φ indeed defines a simulation. \square

We can generalize this result in two directions. First, to systems of two robots on networks whose quotient graphs contain arbitrarily long paths.

Theorem 21. $\{\mathcal{F}(G_n, 2) \mid G_n^* \text{ contains a sub-path of length at least } n\}$ is universal.

Proof. We prove that $\vec{\mathcal{P}}_n^2 \preceq \mathcal{F}(G_n, 2)$, and the universality follows from Lemma 19. By assumption, the robots can agree on an oriented path P of length n in G_n^* , since all its vertices are distinguishable, due to the definition of quotient graph. Then, a robot located at some vertex of G_n is interpreted as lying on the corresponding vertex of G_n^* and follows whatever algorithm it is simulating, remaining on vertices of G_n corresponding to P ; this way, the two robots can simulate $\vec{\mathcal{P}}_n^2$ on G_n . \square




Figure 3: Computing a function by two robots in an oriented ring.

Theorem 20 can also be generalized to systems of three robots on networks with arbitrarily long *girths* (the girth of a graph being the length of its shortest cycle, or infinity if there are no cycles).

Theorem 22. $\{\mathcal{F}(G_n, 3) \mid \text{the girth of } G_n \text{ is at least } n \text{ and finite}\}$ is universal.

Proof. We show that $\mathbb{N}_n \preceq \mathcal{F}(G_{8n}, 3)$. By Lemma 19, it is sufficient to simulate any *sequential* algorithm A for $\vec{\mathcal{P}}_{2n}^2$. In particular, A does not make the two robots move if they lie on the same vertex, or both would have to move, and the algorithm would not be sequential. To do the simulation, we choose a shortest cycle L in G_{8n} , and we initially place all three robots on it: one robot a will remain fixed in one vertex, then we will keep $2n - 1$ empty vertices on L after a , while the other two robots b and c simulate A on the next $2n$ vertices, with c always farthest from a . This makes the robots distinguishable (unless b and c coincide), because L has length at least $8n$ (see Figure 4).

Since L is a shortest cycle, there is a unique shortest path P in G_{8n} containing all three robots. Indeed, if this were not the case, there would be two vertices of L connected by two disjoint paths of length at most $4n - 1$, implying the existence of a cycle of length at most $8n - 2$, and contradicting the fact that the girth is at least $8n$.

Due to the uniqueness of P , if b and c have to move within P , they can deterministically do so. This may cause b and c to switch positions, in which case we simply rename them and we pretend they have not moved.

If c has to move away from b (hence b and c are not on the same vertex, and b stays still because A is sequential), it proceeds along a shortest cycle that contains P . This could be a non-deterministic move, but it will indeed keep the robots on a same shortest cycle, although maybe not L . All the configurations obtained this way are mapped to the corresponding configuration of $\vec{\mathcal{P}}_{2n}^2$, and this yields a surjective partial function that describes the simulation. \square

Our conjecture is that the above results characterize the universal classes of systems with at least three robots on unlabeled networks.




Figure 4: Proof of Theorem 22.

Conjecture 23. *The set $\{\mathcal{F}(G_i, k) \mid i \geq 0\}$, where $k \geq 3$ and every G_i is an unlabeled network, is universal if and only if either the quotient graphs G_i^* have unboundedly long sub-paths, or the graphs G_i have unboundedly long shortest cycles.*

5 Optimizing Network Sizes

In this section, our goal is to compute all the functions on \mathbb{N}_n under the fully synchronous scheduler, using the smallest possible network, and perhaps a large number of robots. We are able to approximate the minimum size of such a network up to a factor that tends to 2 as n goes to infinity, using very short oriented paths. Nonetheless, thanks to the simulation tools developed in Section 3.4, we could as well use unoriented paths or rings, again achieving the optimum size up to factors that tend to small constants.

Lemma 24. *For all $n, k \geq 1$, $\vec{\mathcal{P}}_{n!}^k \preceq \vec{\mathcal{P}}_{kn}^{kn(n-1)/2}$.*

Proof. We divide the base graph of $\vec{\mathcal{P}}_{kn}^{kn(n-1)/2}$ into k sub-paths of length n . In every sub-path, we place a different amount of robots on each vertex, from 0 to $n - 1$ robots. The possible placements of such robots within a sub-path correspond to the $n!$ permutations of n distinct objects. It is well known that the set of permutations can be ordered in such a way that each permutation is obtained from the previous one by swapping only two adjacent objects [18]. If we let the i -th permutation under this ordering encode the i -th vertex of $\vec{\mathcal{P}}_{n!}^k$, we can simulate a move of i -th robot of $\vec{\mathcal{P}}_{n!}^k$ by simply swapping the robots occupying two adjacent vertices of the i -th sub-path of $\vec{\mathcal{P}}_{kn}^{kn(n-1)/2}$. \square

Theorem 25. *For all $n \geq 1$, $\mathbb{N}_n \preceq \vec{\mathcal{P}}_{2m}^{m(m-1)}$, with $(m - 1)! < 2n \leq m!$.*

Proof. Immediate from Lemma 19, Lemma 24 with $k = 2$, and the transitivity of \preceq . \square

This tells us that, on a network with $|V|$ vertices, all the functions on a set of size $2(|V|/2)!$ can be computed, provided that enough robots are available. We can also show that, on the same network, it is impossible to compute all functions on a set of size $|V|! + 1$.

Theorem 26. For all networks $G = (V, E, \ell)$ and all $n, k \geq 1$, if $\mathbb{N}_n \preceq \mathcal{F}(G, k)$, then $|V| \geq n$.

Proof. Let A be an algorithm that computes the cycle function $\lambda_n = \{(i, i + 1) \mid i \in \mathbb{N}_{n-1}\} \cup \{(n - 1, 0)\}$ under $\mathcal{F}(G, k)$, according to the surjective partial function by $\varphi: \mathcal{C}(G, k) \rightarrow \mathbb{N}_n$. For any execution $E = (C_i)_{i \geq 0}$ of A such that $C_0 \in \varphi^{-1}(0)$, the sequence $(\varphi(C_i))_{i \geq 0}$ must span the whole range \mathbb{N}_n infinitely often. Moreover, since $\mathcal{C}(G, k)$ is finite, there is a configuration C that occurs infinitely many times in E . Hence there are two such occurrences, say $C_j = C_{j'} = C$, between which E spans at least $d \geq n$ different configurations. During this fragment of the execution, no two separate sets of robots may end up on the same vertex at the same time, or they become impossible to separate deterministically, contradicting the fact that C must be reached again. In particular, the number of vertices that are occupied by some robots remains the same, q . It follows that d cannot be larger than $|V|!/(|V| - q)!$. Thus $n \leq d \leq |V|!/(|V| - q)! \leq |V|!$, as desired. \square

6 Further Work

In addition to the fully synchronous scheduler, also a semi-synchronous and an asynchronous one can be defined. Both schedulers may activate any subset of the robots at each turn, keeping the others quiescent. The asynchronous scheduler may even delay the robots, making them move based on obsolete observations of the network.

Our universality results can be extended to both these schedulers, by observing that all the algorithms we used in our simulations are sequential. We can also prove a weaker version of Theorem 25 for these schedulers: $\mathbb{N}_n \preceq \vec{\mathcal{P}}_{2m}^m$, with $m = O(\log n)$; we have a matching lower bound for both schedulers, as well. These results indicate that the semi-synchronous and asynchronous schedulers, albeit not drastically reducing the robots' computing powers, make them somewhat less efficient.

We leave two open problems: Conjectures 7 and 23.

References

- [1] L. Blin, A. Milani, M. Potop-Butucaru, and S. Tixeuil. Exclusive perpetual ring exploration without chirality. In *Proc. 24th Int. Symp. on Distributed Computing (DISC)*, 312–327, 2010.
- [2] F. Bonnet, A. Milani, M. Potop-Butucaru, and S. Tixeuil. Asynchronous exclusive perpetual grid exploration without sense of direction. In *Proceedings of 15th International Conference on Principles of Distributed System (OPODIS)*, 251–265, 2011.
- [3] J. Chalopin, P. Flocchini, B. Mans, and N. Santoro. Network exploration by silent and oblivious robots. In *Proceedings of 36th International Workshop on Graph Theoretic Concepts in Computer Science (WG)*, pages 208–219, 2010.
- [4] G. D'Angelo, G. Di Stefano, R. Klasing, and A. Navarra. Gathering of robots on anonymous grids without multiplicity detection. *Theoretical Computer Science*, 610: 158–168, 2016.
- [5] G. D'Angelo, G. Di Stefano, and A. Navarra. Gathering six oblivious robots on anonymous symmetric rings. *J. Discrete Algorithms*, 26: 16–27, 2014.
- [6] G. D'Angelo, G. Di Stefano, and A. Navarra. Gathering on rings under the Look-Compute-Move model. *Distributed Computing*, 27(4): 255–285, 2014.

- [7] G. D’Angelo, G. Di Stefano, A. Navarra, N. Nisse, and K. Suchan. Computing on rings by oblivious robots: A unified approach for different tasks. *Algorithmica* 72(4): 1055–1096, 2015.
- [8] S. Devismes, A. Lamani, F. Petit, and S. Tixeuil. Optimal torus exploration by oblivious mobile robots. INRIA Technical Report HAL-00926573, 2014.
- [9] S. Devismes, F. Petit, and S. Tixeuil. Optimal probabilistic ring exploration by semi-synchronous oblivious robots. *Theoretical Computer Science*, 498: 10–27, 2013.
- [10] Y. Elor and A. M. Bruckstein. Uniform multi-agent deployment on a ring. *Theoretical Computer Science*, 412: 783–795, 2011.
- [11] P. Flocchini, D. Ilcinkas, A. Pelc, and N. Santoro. Remembering without memory: Tree exploration by asynchronous oblivious robots. *Theoretical Computer Science*, 411(14–15): 1583–1598, 2010.
- [12] P. Flocchini, D. Ilcinkas, A. Pelc, and N. Santoro. How many oblivious robots can explore a line. *Information Processing Letters*, 111(20): 1027–1031, 2011.
- [13] P. Flocchini, D. Ilcinkas, A. Pelc, and N. Santoro. Ring exploration by asynchronous oblivious robots. *Algorithmica*, 65(3): 562–583, 2013.
- [14] P. Flocchini, G. Prencipe, and N. Santoro. *Distributed Computing by Oblivious Mobile Robots*. Morgan & Claypool, 2012.
- [15] E. Gilbert and J. Riordan. Symmetry types of periodic sequences. *Illinois Journal of Mathematics*, 5(4): 657–665, 1961.
- [16] S. Guilbault and A. Pelc. Gathering asynchronous oblivious agents with local vision in regular bipartite graphs. *Theoretical Computer Science*, 509: 86–96, 2013.
- [17] T. Izumi, T. Izumi, S. Kamei, and F. Ooshita. Mobile robots gathering algorithm with local weak multiplicity in rings. In *Proceedings of 17th International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, 101–113, 2010.
- [18] S. Johnson. Generation of permutations by adjacent transposition. *Mathematics of Computation*, 17: 282–285, 1963.
- [19] S. Kamei, A. Lamani, F. Ooshita, and S. Tixeuil. Gathering an even number of robots in an odd ring without global multiplicity detection. In *Proceedings of 37th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, 542–553, 2012.
- [20] G. Kant. Drawing planar graphs using the canonical ordering. *Algorithmica*, 16(1): 4–32, 1996.
- [21] R. Klasing, A. Kosowski, and A. Navarra. Taking advantage of symmetries: Gathering of many asynchronous oblivious robots on a ring. *Theoretical Computer Science*, 411: 3235–3246, 2010.
- [22] R. Klasing, E. Markou, and A. Pelc. Gathering asynchronous oblivious mobile robots in a ring. *Theoretical Computer Science*, 390: 27–39, 2008.
- [23] A. Kosowski and A. Navarra. Graph decomposition for improving memoryless periodic exploration. In *Proceedings of 34th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, 501–512, 2009.

- [24] A. Lamani, M. Gradinariu Potop-Butucaru, and S. Tixeuil. Optimal deterministic ring exploration with oblivious asynchronous robots. In *Proceedings of 17th Int. Colloquium on Structural Information and Communication Complexity (SIROCCO)*, 183–196, 2010.
- [25] L. Millet, M. Potop-Butucaru, N. Sznajder, and S. Tixeuil. On the synthesis of mobile robots algorithms: The case of ring gathering. In *Proceedings of 16th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)*, 237–251, 2014.
- [26] F. Ooshita and S. Tixeuil. On the self-stabilization of mobile oblivious robots in uniform rings. *Theoretical Computer Science*, 568: 84–96, 2015.