

THE DISTRIBUTED COMPUTING COLUMN

BY

STEFAN SCHMID

Aalborg University, Denmark
schmiste@cs.aau.dk

SURVEY OF DISTRIBUTED DECISION*

Laurent Feuilloley and Pierre Fraigniaud

Institut de Recherche en Informatique Fondamentale
CNRS and University Paris Diderot

Abstract

We survey the recent distributed computing literature on checking whether a given distributed system configuration satisfies a given boolean predicate, i.e., whether the configuration is legal or illegal w.r.t. that predicate. We consider classical distributed computing environments, including mostly synchronous fault-free network computing (LOCAL and CONGEST models), but also asynchronous crash-prone shared-memory computing (WAIT-FREE model), and mobile computing (FSYNC model).

1 Introduction

The objective of this note is to survey the recent achievements in the framework of *distributed decision*: the computing entities of a distributed system aim at checking whether the system is in a legal state with respect to some boolean predicate. For instance, in a network, the computing entities may be aiming at checking whether the network satisfies some given graph properties.

Recall that, in a *construction task*, processes have to collectively compute a valid global state of a distributed system, as a collection of individual states, like, e.g., providing each node of a network with a color so that to form a proper coloring of that network. Instead, in a *decision task*, processes have to collectively check whether a given global state of a distributed system is valid or not, like, e.g., checking whether a given coloring of the nodes of a network is proper [25]. In general, a typical application of distributed decision is checking the validity of outputs produced by the processes w.r.t. a construction task that they were supposed to solved. This applies to various settings, including randomized algorithms as well as algorithms subject to any kind of faults susceptible to corrupt the memory of the processes.

*Both authors received additional support from Inria project-team GANG.

The global verdict on the legality of the system state is obtained as an aggregate of individual opinions produced by all processes. Typically, each process opinion is a single bit (i.e., *accept* or *reject*) expressing whether the system state looks legal or illegal from the perspective of the process, and the global verdict is the *logical conjunction* of these bits. Note that this mechanism reflects both decision procedures in which the individual opinions of the processes are collected by some centralized entity, and decision procedures where any process detecting some inconsistency in the system raises an alarm and/or launches a recovery procedure, in absence of any central entity. We will also briefly consider less common procedures where each process can send some limited information about its environment in the system, and a central authority gathers the information provided by the processes to forge its verdict about the legality of the whole system state.

The difficulty of distributed decision arises when the processes cannot obtain a global perspective of the system, which is typically the case if one insists on some form of locality in networks, or if the processes are asynchronous and subject to failures. In such frameworks, not all boolean predicates on distributed systems can be checked in a distributed manner, and one of the main issues of distributed decision is to characterize the predicates that can be distributedly checked, and at which cost. For predicates that cannot be checked, or for which checking is too costly, the system can be enhanced by providing processes with *certificates*, with the objective to help these processes for expressing their individual opinions. Such certificates could be produced by an external entity, but they might also well be produced by the processes themselves during a pre-computation phase. One typical framework in which the latter scenario finds application is self-stabilization. Indeed, a self-stabilizing algorithm may produce, together with its distributed output, a distributed certificate that this output is correct. Of course, the certificates are also corruptible, and thus not trustable. Hence, the checking procedure must involve a distributed verification algorithm in charge of verifying the collection of pairs (output, certificate) produced by all the processes. Some even more elaborated mechanisms for checking the legality of distributed system states are considered in the literature, and we survey such mechanisms as well.

We consider the most classical distributed computing models, including synchronous distributed network computing [49]. In this setting, processes are nodes of a graph representing a network. They all execute the same algorithm, they are fault-free, and they are provided with distinct identities in some ID-space (which can be bounded or not). All processes start simultaneously, and computation proceeds in synchronous rounds. At each round, every process exchanges messages with its neighboring processes in the network, and performs individual computation. The volume of communication each node can transmit and receive on each of its links at each round might be bounded or not. The CONGEST model typically assumes that at most $O(\log n)$ bits can be transferred along each link at each round

in n -node networks. (In this case, the ID-space is supposed to be polynomially bounded as a function of the network size). Instead the LOCAL model does not limit the amount of information that can be transmitted along each link at each round. So, a t -round algorithm \mathcal{A} in the LOCAL model can be transformed into another algorithm \mathcal{B} in which every node first collects all data available in the ball of radius t around it, and, second, simulate \mathcal{A} locally without communication.

We also consider other models like asynchronous distributed shared-memory computing [5]. In this setting, every process has access to a global memory shared by all processes. Every process accesses this memory via atomic read and write instructions. The memory is composed of registers, and each process is allocated a set of private registers. Every process can read all the registers, but can only write in its own registers. Processes are given distinct identities in $[n] = \{1, \dots, n\}$ for n -process systems. They runs asynchronously, and are subject to crashes. A process that crashes stops taking steps. An arbitrary large number of processes can crash. Hence, an algorithm must never include instructions leading a process to wait for actions by another process, as the latter process can crash. This model is thus often referred to as the WAIT-FREE model.

Finally, we briefly consider other models, including mobile computing [22], mostly in the fully-synchronous FSYNC model in graphs (where all mobile agents perform in lock-step, moving from nodes to adjacent nodes in a network), and distributed quantum computing (where processes have access to intricate variables).

2 Model and Definitions

Given a boolean predicate, a distributed decision algorithm is a distributed algorithm in which every process p must eventually output a value

$$\text{opinion}(p) \in \{\text{accept}, \text{reject}\}$$

such that the global system state satisfies the given predicate if and only if all processes accept. In other word, the global interpretation of the individual opinions produced by the processes is the logical conjunction of all these opinions:

$$\text{global verdict} = \bigwedge_p \text{opinion}(p).$$

Among the earliest references explicitly related to distributed decision, it is worth mentioning [1, 6, 42]. In this section, we describe the general framework of distributed decision, without explicit references to some specific underlying computational model.

The structure of the section is inspired from the structure of complexity classes in sequential complexity theory. Given the “base” class \mathbf{P} of languages that are

sequentially decidable by a Turing machine in time polynomial in the size of the input, the classes NP (for non-deterministic polynomial time) and BPP (for bounded probability polynomial time) are defined, as well as the classes Σ_k^P and Π_k^P , $k \geq 0$, of the polynomial hierarchy. In this section we assume given an abstract class BC (for *bounded distributed computing*), based on which larger classes can be defined. Such a base class BC could be a *complexity* class like, e.g., the class of graph properties that can be checked in constant time in the LOCAL model, or a *computability* class like, e.g., the class of system properties that can be checked in a shared-memory distributed system subject to crash failures. Given the “base” class BC, we shall define the classes NBC, BPBC, Σ_k^{BC} and Π_k^{BC} , that are to BC what NP, BPP, Σ_k^P and Π_k^P are to P, respectively.

2.1 Distributed Languages

A system *configuration* \mathbf{C} is a (partial) description of a distributed system state. For instance, in distributed network computing, a configuration \mathbf{C} is of the form (G, ℓ) where G is a graph, and $\ell : V(G) \rightarrow \{0, 1\}^*$. Similarly, in shared memory computing, a configuration \mathbf{C} is of the form $\ell : [n] \rightarrow \{0, 1\}^*$ where n is the number of processes. The function ℓ is called *labeling* function, and $\ell(v)$ the *label* of v , which can be any arbitrary bit string. In the context of distributed decision, the label of a process is the input of that process.

For instance, the label of a node in a processor network can be a color, and the label of a process in a shared memory system can be a status like “elected” or “defeated”. Note that, in both examples, a configuration is oblivious to the content of the shared memory and/or to the message in transit. The labeling function ℓ may not describe the full state of each process, but only the content of some specific variables.

Definition 1. *Given a distributed computing model, a distributed language is a Turing-computable set of configurations compatible with this model.*

For instance, in the framework of network computing,

$$\text{PROPER-COLORING} = \{(G, \ell) : \forall \{u, v\} \in E(G), \ell(u) \neq \ell(v)\}$$

is the distributed language composed of all networks with a proper coloring of their nodes (the label $\ell(v)$ of node v is its color). Similarly, in the framework of crash-prone shared-memory computing,

$$\text{AGREEMENT} = \{\ell : \exists y \in \{0, 1\}^*, \forall i \in [n], x(i) = y \text{ or } x(i) = \perp\}$$

is the distributed language composed of all systems where agreement between the non-crashed processes is achieved (the label of process p_i is $\ell(i)$, and the symbol \perp refers to the scenario in which process p_i crashed).

For a fixed distributed language \mathcal{L} , a configuration in \mathcal{L} is said to be *legal*, and a configuration not in \mathcal{L} is said to be *illegal*. Any distributed language \mathcal{L} defines a *construction* task, in which every process must compute a label such that the collection of labels outputted by the processes form a legal configuration for \mathcal{L} . In the following, we are mostly interested in *decision* tasks, where the labels of the nodes are given, and the processes must collectively check whether these labels form a legal configuration.

Notation. Given a system configuration \mathbf{C} with respect to some distributed computing model, we denote by $V(\mathbf{C})$ the set of all computing entities (a.k.a. processes) in \mathbf{C} . This notation reflects the fact that, in the following, the set of processes will most often be identified as the vertex-set $V(G)$ of a graph G

2.2 Distributed Decision

Given a distributed computing model, let us define some *bounded computing* class \mathbf{BC} as a class of distributed languages that can be decided with a distributed algorithm \mathcal{A} using a bounded amount of resources. Such an algorithm \mathcal{A} is said to be *bounded*. What is meant by “resource” depends on the computing model. In most of the models investigated in this paper, the resource of interest is the number of rounds (as in the LOCAL and CONGEST models), or the number of read/write operations (as in the WAIT-FREE model). A distributed language \mathcal{L} is in \mathbf{BC} if and only if there exists a bounded algorithm \mathcal{A} such that, for any input configuration \mathbf{C} , the algorithm \mathcal{A} outputs $\mathcal{A}(\mathbf{C}, v)$ at each process v , and this output satisfies:

$$\mathbf{C} \in \mathcal{L} \iff \text{for every } v \in V(\mathbf{C}), \mathcal{A}(\mathbf{C}, v) = \text{accept.} \quad (1)$$

That is, for every $\mathbf{C} \in \mathcal{L}$, running \mathcal{A} on \mathbf{C} results in all processes accepting \mathbf{C} . Instead, for every $\mathbf{C} \notin \mathcal{L}$, running \mathcal{A} on \mathbf{C} results in at least one process rejecting \mathbf{C} .

Example. In the context of network computing, PROPER-COLORING can be decided in one round, by having each node merely comparing its color with the ones of its neighbors, and accepting if and only if its color is different from all these colors. Similarly, in the context of shared-memory computing, AGREEMENT can be decided by having each node performing just one read/write operation, accepting if and only if all labels different from \perp observed in memory are identical. In other words, assuming that \mathbf{BC} is a network computing class bounding algorithms to perform in a constant number of rounds, we have

$$\text{PROPER-COLORING} \in \mathbf{BC}$$

for any model allowing each process to send its color to all its neighbors in a constant number of rounds, like, e.g., the LOCAL model. Similarly, assuming that BC is a shared-memory computing class bounding algorithms to perform in a constant number of read/write operations, we have

$$\text{AGREEMENT} \in \text{BC}.$$

Notation. In the following, Eq. (1) will often be abbreviated with

$$\mathbf{C} \in \mathcal{L} \iff \mathcal{A}(\mathbf{C}) = \text{accept}$$

in the sense that \mathcal{A} accepts if and only if each of the processes accepts.

Note that the rule of distributed decision, i.e., the logical conjunction of the individual boolean outputs of the processes is not symmetric. For instance, deciding whether a graph is properly colored can be done locally, while deciding whether a graph is *not* properly colored may require long-distance communications. On the other hand, asking for other rules, like unanimous decision (where all processes must reject an illegal configuration) or even just majority decision, would require long-distance communications for most classical decision problems.

2.3 Probabilistic Distributed Decision

The bounded computing class BC is a base class upon which other classes can be defined. Given $p, q \in [0, 1]$, we define the class $\text{BPBC}(p, q)$, for *bounded probability bounded computing*, as the class of all distributed languages \mathcal{L} for which there exists a randomized bounded algorithm \mathcal{A} such that, for every configuration \mathbf{C} ,

$$\begin{cases} \mathbf{C} \in \mathcal{L} & \Rightarrow \Pr[\mathcal{A}(\mathbf{C}) = \text{accept}] \geq p; \\ \mathbf{C} \notin \mathcal{L} & \Rightarrow \Pr[\mathcal{A}(\mathbf{C}) = \text{reject}] \geq q. \end{cases} \quad (2)$$

Such an algorithm \mathcal{A} is called a (p, q) -decider for \mathcal{L} . Note that, as opposed to the class BPP of complexity theory, the parameters p and q are not arbitrary, in the sense that boosting the probability of success of a (p, q) -decider in order to get a (p', q') -decider with $p' > p$ and $q' > q$ is not always possible. Indeed, if \mathcal{A} is repeated many times on an illegal instance, say k times, it may well be the case that each node will reject at most once during the k repetitions, because, at each iteration of \mathcal{A} , rejection could come from a different node. As a consequence, classical boosting techniques based on repetition and taking majority do not necessarily apply.

Example. Let us consider the following distributed language, where each process can be labeled either white or black, i.e., $\ell : V(\mathbf{C}) \rightarrow \{\circ, \bullet\}$:

$$\text{AMOS} = \{\ell : |\{v \in V(\mathbf{C}) : \ell(v) = \bullet\}| \leq 1\}.$$

Here, AMOS stands for “at most one selected”, where a node v is selected if $\ell(v) = \bullet$. There is a trivial (p, q) -decider for AMOS as long as $p^2 + q \leq 1$, which works as follows. Every node v with $\ell(v) = \circ$ accepts (with probability 1). A node v with $\ell(v) = \bullet$ accepts with probability p , and rejects with probability $1 - p$. If $\mathbf{C} \in \text{AMOS}$, then $\Pr[\text{all nodes accept } \mathbf{C}] \geq p$. If $\mathbf{C} \notin \text{AMOS}$, then $\Pr[\text{at least one node rejects } \mathbf{C}] \geq 1 - p^2 \geq q$.

2.4 Distributed Verification

Given a bounded computing class BC, we describe the class NBC, which is to BC what NP is to P in complexity theory. We define the class NBC, for *non-deterministic bounded computing*, as the class of all distributed languages \mathcal{L} such that there exists a bounded algorithm \mathcal{A} satisfying that, for every configuration \mathbf{C} ,

$$\mathbf{C} \in \mathcal{L} \iff \exists c : \mathcal{A}(\mathbf{C}, c) = \text{accepts} \quad (3)$$

where

$$c : V(\mathbf{C}) \rightarrow \{0, 1\}^*.$$

The function c is called the *certifying* function. It assigns a certificate to every process, and the certificates do not need to be identical. Note that the certificate $c(v)$ of process v must not be mistaken with the label $\ell(v)$ of that process.

The bounded algorithm \mathcal{A} is also known as a *verification* algorithm for \mathcal{L} , as it verifies a given proof c , which is supposed to certify that $\mathbf{C} \in \mathcal{L}$. At each process $v \in V(\mathbf{C})$, the verification algorithm takes as input the pair $(\ell(v), c(v))$. Note that the appropriate certificate c leading to accept a configuration $\mathbf{C} \in \mathcal{L}$ may depend on the given configuration \mathbf{C} . However, for $\mathbf{C} \notin \mathcal{L}$, the verification algorithm \mathcal{A} must systematically guaranty that at least one process rejects, whatever the given certificate function is.

Alternatively, one can interpret Eq. (3) as a game between a *prover* which, for every configuration \mathbf{C} , assigns a certificate $c(v)$ to each process $v \in V(\mathbf{C})$, and a *verifier* which checks that the certificates assigned by the prover collectively form a *proof* that $\mathbf{C} \in \mathcal{L}$. For a legal configuration (i.e., a configuration in \mathcal{L}) the prover must be able to produce a distributed proof leading the distributed verifier to accept, while, for an illegal configuration, the verifier must reject in at least one node whatever the proof provided by the prover is.

Example. Let us consider the distributed language

$$\text{ACYCLIC} = \{(G, \ell) : G \text{ has no cycles}\}$$

in the context of network computing. Note that ACYCLIC cannot be decided locally, even in the LOCAL model. However, ACYCLIC can be verified in just one round. If G is acyclic, i.e., G is a forest, then let us select an arbitrary node in each tree of G , and call it a root. Next, let us assign to each node $u \in V(G)$ the certificate $c(u)$ equal to its distance to the root of its tree. The verification algorithm \mathcal{A} then proceeds at every node u as follows. Node u exchanges its certificate with the ones of its neighbors, and checks that it has a unique neighbor v satisfying $c(v) = c(u) - 1$, and all the other neighbors $w \neq v$ satisfying $c(w) = c(u) + 1$. (If u has $c(u) = 0$, then it checks that all its neighbors w have $c(w) = 1$). If all tests are passed, then u accepts, else it rejects. If G is a acyclic, then, by construction, the verification accepts at all nodes. Instead, if G has a cycle, then, for every setting of the certifying function, some inconsistency will be detected by at least one node of the cycle, which leads this node to reject. Hence

$$\text{ACYCLIC} \in \text{NBC}$$

where BC bounds the number of rounds, for every distributed computing model allowing every node to exchange $O(\log n)$ bits along each of its incident edges at every round, like, e.g., the CONGEST model.

Notation. For any function $f : \mathbb{N} \rightarrow \mathbb{N}$, we define $\text{NBC}(f)$ as the class NBC where the certificates are bounded to be on at most $f(n)$ bits in n -node networks. For $f \in \Theta(\log n)$, $\text{NBC}(f)$ is rather denoted by $\log\text{-NBC}$.

2.5 Distributed Decision Hierarchy

In the same way the polynomial hierarchy PH is built upon P using alternating universal and existential quantifiers, one can define a hierarchy built upon base class BC . Given a class BC for some distributed computing model, we define the *distributed decision hierarchy* DH^{BC} as follows. We set $\Sigma_0^{\text{BC}} = \Pi_0^{\text{BC}} = \text{BC}$, and, for $k \geq 1$, we set Σ_k^{BC} as the class of all distributed languages \mathcal{L} such that there exists a bounded algorithm \mathcal{A} satisfying that, for every configuration \mathbf{C} ,

$$\mathbf{C} \in \mathcal{L} \iff \exists c_1 \forall c_2 \exists c_3 \dots Q c_k : \mathcal{A}(\mathbf{C}, c_1, \dots, c_k) = \text{accept}$$

where, for every $i \in \{1, \dots, k\}$, $c_i : V(\mathbf{C}) \rightarrow \{0, 1\}^*$, and Q is the universal quantifier if k is even, and the existential one otherwise. The class Π_k^{BC} is defined similarly, by having a universal quantifier as first quantifier, as opposed to an existential one as in Σ_k^{BC} . The c_i 's are called *certifying* functions. In particular, we have

$$\text{NBC} = \Sigma_1^{\text{BC}}.$$

Finally, we define

$$\text{DH}^{\text{BC}} = (\cup_{k \geq 0} \Sigma_k^{\text{BC}}) \cup (\cup_{k \geq 0} \Pi_k^{\text{BC}}).$$

As for NBC, a class Σ_k^{BC} or Π_k^{BC} can be viewed as a game between a prover (playing the existential quantifiers), a disprover (playing the universal quantifiers), and a verifier (running a verification algorithm \mathcal{A}).

Example. Let us consider the distributed language

$$\text{VERTEX-COVER} = \{(G, \ell) : \{v \in V(G) : \ell(v) = 1\} \text{ is a minimum vertex cover}\}$$

in the context of network computing. We show that $\text{VERTEX-COVER} \in \Pi_2^{\text{BC}}$, that is, there exists a bounded distributed algorithm \mathcal{A} such that

$$(G, \ell) \in \text{VERTEX-COVER} \iff \forall c_1 \exists c_2 : \mathcal{A}(G, \ell, c_1, c_2) = \text{accept}$$

where BC is any network computing class bounding algorithms to perform in a constant number of rounds. For any configuration (G, ℓ) , the disprover tries to provide a vertex cover $c_1 : V(G) \rightarrow \{0, 1\}$ of size smaller than the solution ℓ , i.e., $|\{v \in V(G) : c_1(v) = 1\}| < |\{v \in V(G) : \ell(v) = 1\}|$. On a legal configuration (G, ℓ) , the prover then reacts by providing each node v with a certificates $c_2(v)$ such that the c_2 -certificates collectively encode a spanning tree (and its proof) aiming at demonstrating that there is an error in c_1 (like c_1 is actually not smaller than ℓ , or c_1 is not covering some edge, etc.). It follows that

$$\text{VERTEX-COVER} \in \Pi_2^{\text{BC}}$$

for any model allowing each process to exchange $O(\log n)$ -bits messages with its neighbors in a constant number of rounds, like, e.g., the CONGEST model.

Notation. Similarly to the class NBC, for any function $f : \mathbb{N} \rightarrow \mathbb{N}$, we define $\Sigma_k^{\text{BC}}(f)$ (resp., $\Pi_k^{\text{BC}}(f)$) as the class Σ_k^{BC} (resp., Π_k^{BC}) where all certificates are bounded to be on at most $f(n)$ bits in n -node networks. For $f \in \Theta(\log n)$, these classes are denoted by $\log\text{-}\Sigma_k^{\text{BC}}$ and $\log\text{-}\Pi_k^{\text{BC}}$, respectively. The classes $\text{DH}^{\text{BC}}(f)$ and $\log\text{-}\text{DH}^{\text{BC}}$ are defined similarly.

3 Distributed Decision in Networks

In this section, we focus on languages defined as collections of configurations of the form (G, ℓ) where G is a simple connected n -node graph, and $\ell : V(G) \rightarrow \{0, 1\}^*$ is a labeling function assigning to every node v a label $\ell(v)$. Recall that an algorithm \mathcal{A} is *deciding* a distributed language \mathcal{L} if and only if, for every configuration (G, ℓ) ,

$$(G, \ell) \in \mathcal{L} \iff \mathcal{A}(G, \ell) \text{ accepts at all nodes.}$$

3.1 LOCAL model

3.1.1 Local Distributed Decision (LD and BPLD)

In their seminal paper [48], Naor and Stockmeyer define the class LCL, for *locally checkable labelings*. Let $\Delta \geq 0$, $k \geq 0$, and $t \geq 0$, and let \mathcal{B} be a set of balls of radius at most t with nodes of degree at most Δ , labeled by labels in $[k]$. Note that \mathcal{B} is finite. Such a set \mathcal{B} defines the language \mathcal{L} consisting of all configurations (G, ℓ) where G is a graph with maximum degree Δ , and $\ell : V(G) \rightarrow [k]$, such that all balls of radius t in (G, ℓ) belong to \mathcal{B} . The set \mathcal{B} is called the set of *good balls* for \mathcal{L} . LCL is the class of languages that can be defined by a set of good balls, for some parameters Δ , k , and t . For instance the set of k -colored graphs with maximum degree Δ is a language in LCL. The good balls of this LCL language are simply the balls of radius 1 where the center node is labeled with a color different from all the colors of its neighbors.

A series of results were achieved in [48] about LCL languages. In particular, it is Turing-undecidable whether any given $\mathcal{L} \in \text{LCL}$ has a construction algorithm running in $O(1)$ rounds in the LOCAL model. Also, [48] showed that the node IDs play a limited role in the context of LCL languages. Specifically, [48] proves that, for every $r \geq 0$, if a language $\mathcal{L} \in \text{LCL}$ has a r -round construction algorithm, then it has also a r -round *order invariant* construction algorithm, where an algorithm is order invariant if the *relative order* of the node IDs may play a role, but not the actual *values* of these IDs. The assumption $\mathcal{L} \in \text{LCL}$ can actually be discarded, as long as \mathcal{L} remains defined on constant degree graphs with constant labels. That is, [3] proved that, in constant degree graphs, if a language with constant size labels has a r -round construction algorithm, then it has also a r -round order invariant construction algorithm. Last but not least, [48] established that randomization is of little help in the context of LCL languages. Specifically, [48] proves that if a language $\mathcal{L} \in \text{LCL}$ has a randomized Monte-Carlo construction algorithm running in $O(1)$ rounds, then \mathcal{L} also has a deterministic construction algorithm running in $O(1)$ rounds.

The class LD, for *local decision* was defined in [33] as the class of all distributed languages that can be decided in $O(1)$ rounds in the LOCAL model. The class LD is the basic class playing the role of BC in the context of local decision. Hence $\text{LCL} \subseteq \text{LD}$ since the set of good balls of a language in LCL is, by definition, finite. On the other hand, $\text{LCL} \subset \text{LD}$, where the inclusion is strict since LD does not restrict the graphs to be of bounded degree, nor the labels to be of bounded size. Given $p, q \in [0, 1]$, the class $\text{BPLD}(p, q)$, for *bounded probability local decision*, was defined in [33] as the class of languages for which there is a (p, q) -decider running in $O(1)$ rounds in the LOCAL model. For $p^2 + q \leq 1$, $\text{BPLD}(p, q)$ is shown to include languages that cannot be even decided deterministically in $o(n)$

rounds. On the other hand, [33] also establishes a derandomization result, stating that, for $p^2 + q > 1$, if $\mathcal{L} \in \text{BPLD}(p, q)$, then $\mathcal{L} \in \text{LD}$. This result however holds only for languages closed under node deletion, and it is proved in [27] that, for any every $c \geq 2$, there exists a language \mathcal{L} with a (p, q) -decider satisfying $p^c + q > 1$ and running in a single round, which cannot be decided deterministically in $o(\sqrt{n})$ rounds. On the other hand, [27] proves that, for $p^2 + q > 1$, we have $\text{BPLD}(p, q) = \text{LD}$ for all languages restricted on paths.

On the negative side, it was proved in [27] that boosting the probability of success for decision tasks is not always achievable in the distributed setting, by considering the classes

$$\text{BPLD}_k = \bigcup_{p^{1+1/k} + q > 1} \text{BPLD}(p, q) \text{ and } \text{BPLD}_\infty = \bigcup_{p+q > 1} \text{BPLD}(p, q)$$

for any $k \geq 1$, and proving that, for every $k \geq 1$, $\text{BPLD}_k \subset \text{BPLD}_\infty$, and $\text{BPLD}_k \subset \text{BPLD}_{k+1}$, where all inclusions are strict.

On the positive side, it was proved in [20] that the result in [48] regarding the derandomization of construction algorithms can be generalized from LCL to BPLD. Namely, [20] proves that, for languages on bounded degree graphs and bounded size labels, for every $p > \frac{1}{2}$ and $q > \frac{1}{2}$, if $\mathcal{L} \in \text{BPLD}(p, q)$ has a randomized Monte-Carlo construction algorithm running in $O(1)$ rounds, then \mathcal{L} has also a deterministic construction algorithm running in $O(1)$ rounds.

3.1.2 Identity-Oblivious Algorithms (LDO)

In the LOCAL model, a distributed algorithm is *identity-oblivious*, or simply *ID-oblivious*, if the outputs of the nodes are not impacted by the identities assigned to the nodes. That is, for any two ID-assignments given to the nodes, the output of every node must be identical in both cases. Note that an *identity-oblivious* algorithm may use the IDs of the nodes (e.g., to distinguish them), but the output must be oblivious to these IDs.

The class LDO, for *local decision oblivious* was defined in [28, 29], as the class of all distributed languages that can be decided in $O(1)$ rounds by an ID-oblivious algorithm in the LOCAL model. The class LDO is the basic class playing the role of BC in the context of ID-oblivious local decision. It is shown in [29] that $\text{LDO} = \text{LD}$ when restricted to languages that are closed under node deletion. However, it is proved in [28] that $\text{LDO} \subset \text{LD}$, where the inclusion is strict. In the language $\mathcal{L} \in \text{LD} \setminus \text{LDO}$ used in [28] to prove the strict inclusion $\text{LDO} \subset \text{LD}$, each node label includes a Turing machine M . Establishing $\mathcal{L} \in \text{LD}$ makes use of an algorithm simulating M at each node, for a number of rounds equal to the identity of the node. Establishing $\mathcal{L} \notin \text{LDO}$ makes use of the fact that an ID-oblivious algorithm can be sequentially simulated, and therefore, if an ID-oblivious algorithm would allow

to decide \mathcal{L} , then by simulation of this algorithm, there would exist a sequential algorithm for separating the set of Turing machines that halts and output 0 from the set of Turing machines that halts and output 1, which is impossible.

In [29,30], the power of IDs in local decision is characterized using *oracles*. An oracle is a trustable party with full knowledge of the input, who can provide nodes with information about this input. It is shown in [29] that $\text{LDO} \subseteq \text{LD} \subseteq \text{LDO}^{\#\text{node}}$ where $\#\text{node}$ is the oracle providing each node with an arbitrary large upper bound on the number of nodes. A *scalar* oracle f returns a list $f(n) = (f_1, \dots, f_n)$ of n values that are assigned arbitrarily to the n nodes in a one-to-one manner. A scalar oracle f is large if, for any set of k nodes, the largest value provided by f to the nodes in this set grows with k . [30] proved that, for any computable scalar oracle f , we have $\text{LDO}^f = \text{LD}^f$ if and only if f is large, where LD^f (resp., LDO^f) is the class of languages that can be locally decided in $O(1)$ rounds in the LOCAL model by an algorithm (resp., by an ID-oblivious algorithm) which uses the information provided by f available at the nodes.

3.1.3 Anonymous Networks

Derandomization results were achieved in [19] in the framework of anonymous network (that is, nodes have no IDs). Namely, for every language \mathcal{L} that can be decided locally in any anonymous network, if there exists a randomized anonymous construction algorithm for \mathcal{L} , then there exists a deterministic anonymous construction algorithm for \mathcal{L} , provided that the latter is equipped with a 2-hop coloring of the input network.

3.2 CONGEST model

3.2.1 Non-Local Algorithms

In [44] and [17] the authors consider decision problems such as checking whether a given set of edges forms a spanning tree, checking whether a given set of edges forms a minimum-weight spanning tree (MST), checking various forms of connectivity, etc. All these decision tasks require essentially $\Theta(\sqrt{n} + D)$ rounds (the lower bound is typically obtained using reduction to communication complexity). In particular, [17] proved that checking whether a given set of edges is a spanning tree requires $\Omega(\sqrt{n} + D)$ rounds, which is much more than what is required to construct a spanning tree ($O(D)$ rounds, using a simple breadth-first search). However, [17] proved that, for some other problems (e.g., MST), lower bounds on the round-complexity of the decision task consisting in checking whether a solution is valid yield lower bounds on the round-complexity of the corresponding construction task, and this holds also for the construction of approximate solutions.

The *congested clique* model is the CONGEST model restricted to complete graphs. Deciding whether a graph given as input contains some specific patterns as subgraphs has been considered in [16] and [18] for the congested clique. In particular, [16] provides an algorithm for deciding the presence of a k -node cycle C_k running in $O(2^{O(k)}n^{0.158})$ -rounds.

3.2.2 Local Algorithms

Very few distributed languages on graphs can be checked locally in the CONGEST model. For instance, even just deciding whether G contains a triangle cannot be done in $O(1)$ rounds in the CONGEST model. *Distributed property testing* is a framework recently introduced in [15]. Let $0 < \epsilon < 1$ be a fixed parameter. Recall that, according to the usual definition borrowed from *property testing* (in the so-called *sparse* model), a graph property P is ϵ -far from being satisfied by an m -edge graph G if applying a sequence of at most ϵm edge-deletions or edge-additions to G cannot result in a graph satisfying P . We say that a distributed algorithm \mathcal{A} is a distributed *testing* algorithm for P if and only if, for any graph G modeling the actual network,

$$\begin{cases} G \text{ satisfies } P \implies \Pr[\mathcal{A} \text{ accepts } G \text{ in all nodes}] \geq \frac{2}{3}; \\ G \text{ is } \epsilon\text{-far from satisfying } P \implies \Pr[\mathcal{A} \text{ rejects } G \text{ in at least one node}] \geq \frac{2}{3}. \end{cases}$$

Among other results, [15] proved that, in bounded degree graphs, bipartiteness can be distributedly tested in $O(\text{polylog } n)$ rounds in the CONGEST model. Moreover, it is also proved that triangle-freeness can be distributedly tested in $O(1)$ rounds. (The dependence in ϵ is hidden in the big- O notation). This latter result has been recently extended in [40] to testing H -freeness, for every 4-node graph H , in $O(1)$ rounds. On the other hand, it is not known whether distributed testing K_5 -freeness or C_5 -freeness can be achieved in $O(1)$ rounds, and [40] proves that “natural” approaches based on DFS or BFS traversals do not work.

3.3 General Interpretation of Individual Outputs

In [3, 4], a generalization of distributed decision is considered, where every node outputs not just a single bit (accept or reject), but can output an arbitrary bit-string. The global verdict is then taken based on the multi-set of all the binary strings outputted by the nodes. The concern is restricted to decision algorithms performing in $O(1)$ rounds in the LOCAL model, and the objective is to minimize the size of the outputs. The corresponding basic class BC for outputs on $O(1)$ bits is denoted by ULD, for *universal LD*. (It is universal in the sense that the global interpretation of the individual outputs is not restricted to the logical conjunction). It is proved in [3] that, for any positive even integer Δ , every distributed decision algorithm for

cycle-freeness in connected graphs with degree at most Δ must produce outputs of size at least $\lceil \log \Delta \rceil - 1$ bits. Hence, cycle-freeness does not belong to ULD in general, but it does belong to ULD for constant degree graphs.

In [11] the authors consider a model in which each node initially knows the IDs of its neighbors, while the nodes do not communicate through the edges of the network but via a public whiteboard. The concern of [11] is mostly restricted to the case in which every node can write only once on the whiteboard, and the objective is to minimize the size of the message written by each node on the whiteboard. The global verdict is then taken based on the collection of messages written on the whiteboard. It is shown that, with just $O(\log n)$ -bit messages, it is possible to rebuild the whole graph from the information on the whiteboard as long as the graph is planar or, more generally, excluding a fixed minor. Variants of the model are also considered, in which problems such as deciding triangle-freeness or connectivity are considered. See also [43] for deciding the presence of induced subgraphs.

4 Distributed Verification in Networks

In this section, we still focus on languages defined as collections of configurations of the form (G, ℓ) where G is a simple connected n -node graph, and $\ell : V(G) \rightarrow \{0, 1\}^*$ is a labeling function. Recall that an algorithm \mathcal{A} is *verifying* a distributed language \mathcal{L} if and only if, for every configuration (G, ℓ) ,

$$(G, \ell) \in \mathcal{L} \iff \exists c : \mathcal{A}(G, \ell, c) \text{ accepts at all nodes} \quad (4)$$

where $c : V(G) \rightarrow \{0, 1\}^*$, and $c(v)$ is called the *certificate* of node $v \in V(G)$. Again, the certificate $c(v)$ of node v must not be mistaken with the label $\ell(v)$ of node v . Also, the notion of certificate must not be confused with the notion of *advice*. While the latter are trustable information provided by an oracle [26, 31, 32], the former are proofs that must be verified.

We survey the results about the class $\text{NBC} = \Sigma_1^{\text{BC}}$ where the basic class BC is LD, LDO, ULD, etc.

4.1 LOCAL model

It is crucial to distinguish two cases in Eq. (4), depending on whether the certificates can depend on the identities assigned to the nodes, or not, as reflected in Eq. (5) and (6) below.

4.1.1 Local Distributed Verification (Σ_1^{LD} , PLS, and LCP)

A distributed language \mathcal{L} satisfies $\mathcal{L} \in \Sigma_1^{\text{LD}}$ if and only if there exists a verification algorithm \mathcal{A} running in $O(1)$ rounds in the LOCAL model such that, for every configuration (G, ℓ) , we have

$$\begin{cases} (G, \ell) \in \mathcal{L} & \Rightarrow \forall \text{ID}, \exists c, \mathcal{A}(G, \ell, c) \text{ accepts at all nodes} \\ (G, \ell) \notin \mathcal{L} & \Rightarrow \forall \text{ID}, \forall c, \mathcal{A}(G, \ell, c) \text{ rejects in at least one node} \end{cases} \quad (5)$$

where $c : V(G) \rightarrow \{0, 1\}^*$, and where, for $(G, \ell) \in \mathcal{L}$, the assignment of the certificates to the nodes may depend on the identities given to these nodes. This notion has actually been introduced under the terminology *proof-labeling scheme* in [47], where the concern is restricted to verification algorithms running in just a single round, with the objective of minimizing the size of the certificates. In particular, it is proved that minimum-weight spanning tree can be verified with certificates on $O(\log^2)$ bits in n -node networks, and this bound is tight [45] (see also [44]). Interestingly, the $\Omega(\log^2 n)$ bits lower bound on the certificate size can be broken, and reduced to $O(\log n)$ bits, to the price of allowing verification to proceed in $O(\log n)$ rounds [46]. There are tight connections between proof-labeling schemes and compact silent self-stabilizing algorithms [13], and proof-labeling schemes can even be used as a basis to semi-automatically derive compact time-efficient self-stabilizing algorithms [12]. Let PLS be the class of distributed languages for which there exists a proof-labeling scheme. We have

$$\text{PLS} = \text{ALL}$$

where ALL is the class of all distributed languages on networks (i.e., with configurations of the form (G, ℓ)). This equality is however achieved using certificates on $O(n^2 + nk)$ bits in n -node networks, where k is the maximum size of the labels in the given configuration (G, ℓ) . The $O(n^2)$ bits are used to encode the adjacency matrix of the network, and the $O(nk)$ bits are used to encode the inputs to the nodes.

The notion of proof-labeling scheme has been extended in [41] to the notion of *locally checkable proofs*, which is the same as proof-labeling scheme but where the verification algorithm is not bounded to run in a single round, but may perform an arbitrarily large constant number of rounds. Let LCP be the associate class of distributed languages. By definition, we have

$$\text{LCP} = \Sigma_1^{\text{LD}},$$

and, more specifically,

$$\text{LCP}(f) = \Sigma_1^{\text{LD}}(f)$$

for every function f bounding the size of the certificates. Moreover, since $\text{PLS} = \text{ALL}$, it follows that

$$\text{PLS} = \text{LCP} = \Sigma_1^{\text{LD}} = \text{ALL}.$$

Yet, allowing more rounds for the verification may enable to save space in the certificate size. This is indeed the case for some languages [10], that is there are functions f for which

$$\text{PLS}(f) \subset \text{LCP}(f)$$

with strict inclusions. It is proved in [41] that there are natural languages (e.g., the set of graphs with a non-trivial automorphism, 3-non-colorability, etc.) which require certificates on $\tilde{\Omega}(n^2)$ bits in n -node networks. Recently, [9] introduced a mechanism enabling to reduce exponentially the amount of communication in proof-labeling schemes, using randomization. See also [51] for applications of locally checkable proofs to software-defined networks.

4.1.2 Identity-Oblivious Algorithms (Σ_1^{LDO} and NLD)

A distributed language \mathcal{L} satisfies $\mathcal{L} \in \Sigma_1^{\text{LDO}}$ if and only if there exists a verification algorithm \mathcal{A} running in $O(1)$ rounds in the LOCAL model such that, for every configuration (G, ℓ) , we have

$$\begin{cases} (G, \ell) \in \mathcal{L} & \Rightarrow \exists c, \forall \text{ID}, \mathcal{A}(G, \ell, c) \text{ accepts at all nodes} \\ (G, \ell) \notin \mathcal{L} & \Rightarrow \forall c, \forall \text{ID}, \mathcal{A}(G, \ell, c) \text{ rejects in at least one node} \end{cases} \quad (6)$$

where $c : V(G) \rightarrow \{0, 1\}^*$, and, for $(G, \ell) \in \mathcal{L}$, the assignment of the certificates to the nodes must not depend on the identities given to these nodes. In [33], the class NLD, for *non-deterministic local decision* is introduced. In NLD, even if the certificates must not depend on the identities of the nodes, the verification algorithm is not necessarily identity-oblivious. Yet, it was proved in [29] that restricting the verification algorithm to be identity-oblivious does not restrict the power of the verifier. Hence,

$$\text{NLD} = \Sigma_1^{\text{LDO}}$$

Σ_1^{LDO} is characterized in [29] as the class of languages that are *closed under lift*, where H is a k -lift of G if there exists an homomorphism from H to G preserving radius- k balls. Hence,

$$\Sigma_1^{\text{LDO}} \subset \text{ALL}$$

where the inclusion is strict. However, it was proved in [33] that, for every distributed language \mathcal{L} , and for every p, q such that $p^2 + q \leq 1$, there is a non-deterministic (p, q) -decider for \mathcal{L} . In other words, for every p, q such that $p^2 + q \leq 1$, we have

$$\text{BPNLD}(p, q) = \text{ALL}.$$

In [33], a complete problem for NLD was identified. However, it was recently noticed in [7] that the notion of local reduction used in [33] is way too strong, enabling to bring languages outside NLD into NLD. A weaker notion of local

reduction was thus defined in [7], preserving the class NLD. A language is proved to be NLD-complete under this weaker type of local reduction.

4.1.3 Anonymous Networks

Distributed verification in the context of fully anonymous networks (no node-identities, and no port-numbers) has been considered in [23].

4.2 CONGEST model ($\log\text{-}\Sigma_1^{\text{LD}}$ and $\log\text{-LCP}$)

The class $\log\text{-LCP}$, that is, $\log\text{-}\Sigma_1^{\text{LD}}$, i.e., Σ_1^{LD} with certificates of size $O(\log n)$ bits, was investigated in [41]. This class fits well with the CONGEST model, which allows to exchange messages of at most $O(\log n)$ bits at each round. For instance, non-bipartiteness is in $\log\text{-LCP}$. Also, restricted to bounded-degree graphs, there are problems in $\log\text{-LCP}$ that are not contained in NP, but $\log\text{-LCP} \subseteq \text{NP/poly}$, i.e., NP with a polynomial-size non-uniform advice. Last but not least, [41] shows that existential MSO on connected graphs is included in $\log\text{-LCP}$.

4.3 General Interpretation of Individual Outputs

As already mentioned in Section 3.3, a generalization of distributed decision was considered in [3,4], where every node outputs not just a single bit (accept or reject), but can output an arbitrary bit-string. The global verdict is then taken based on the multi-set of all the binary strings outputted by the nodes. The concern is restricted to decision algorithm performing in $O(1)$ rounds in the LOCAL model, and the objective is to minimize the size of the output. The certificates must not depend on the node IDs, that is, verification proceed as specified in Eq. (6). For constant size outputs, it is shown in [4] that the class $\text{UNLD} = \Sigma_1^{\text{ULD}}$ satisfies

$$\text{UNLD} = \text{ALL}$$

with just 2-bit-per-node outputs, which has to be consider in contrast to the fact that NLD is restricted to languages that are closed under lift (cf. Section 4.1.2). This result requires using certificates on $O(n^2 + nk)$ bits in n -node networks, where k is the maximum size of the labels in the given configuration (G, ℓ) , but [4] shows that this is unavoidable. Also, while verifying cycle-freeness using the logical conjunction of the 1-bit-per-node outputs requires certificates on $\Omega(\log n)$ bits [41], it is proved in [4] that, by simply using the conjunction and the disjunction operators together, on only 2-bit-per-node outputs, one can verify cycle-freeness using certificates of size $O(1)$ bits.

5 Local Hierarchies in Networks

In this section, we survey the results about the hierarchies Σ_k^{BC} and Π_k^{BC} , $k \geq 0$, for different basic classes BC, including LD, LDO, etc.

5.1 LOCAL model (DH^{LD} and DH^{LDO})

We have seen in Section 4.1.1 that $\Sigma_1^{\text{LD}} = \text{ALL}$, which implies that the local distributed hierarchy DH^{LD} collapses at the first level. On the other hand, we have also seen in Section 4.1.2 that $\Sigma_1^{\text{LDO}} \subset \text{ALL}$, where the inclusion is strict as Σ_1^{LDO} is restricted to languages that are closed under lift. It was recently proved in [7] that

$$\text{LDO} \subset \Pi_1^{\text{LDO}} \subset \Sigma_1^{\text{LDO}} = \Sigma_2^{\text{LDO}} \subset \Pi_2^{\text{LDO}} = \text{ALL}$$

where all inclusions are strict. Hence, the local ID-oblivious distributed hierarchy collapses at the second level. Moreover, it is shown that Π_2^{LDO} has a complete problem for local label-preserving reductions. (A complete problem for ALL was also identified in [33], but using an inappropriate notion of local reduction).

In the context of a general interpretation of individual outputs (see Section 4.3), [4] proved that $\Sigma_1^{\text{ULD}} = \text{ALL}$.

5.2 CONGEST model (log-DH^{LD})

We have previously seen that $\Sigma_1^{\text{LD}} = \text{ALL}$. However, this requires certificates of polynomial size. In order to fit with the constraints of the CONGEST model, the local distributed hierarchy with certificate of logarithmic size was recently investigated in [21]. While it follows from [45] that $\text{MST} \notin \log\text{-}\Sigma_1^{\text{LD}}$, it is shown in [21] that

$$\text{MST} \in \log\text{-}\Pi_2^{\text{LD}}.$$

In fact, [21] proved that, for any $k \geq 1$,

$$\log\text{-}\Sigma_{2k}^{\text{LD}} = \log\text{-}\Sigma_{2k-1}^{\text{LD}} \quad \text{and} \quad \log\text{-}\Pi_{2k+1}^{\text{LD}} = \log\text{-}\Pi_{2k}^{\text{LD}},$$

and thus focused only on the hierarchy $(\Lambda_k)_{k \geq 0}$ defined by $\Lambda_0 = \text{LD}$, and, for $k \geq 1$,

$$\Lambda_k = \begin{cases} \log\text{-}\Sigma_k^{\text{LD}} & \text{if } k \text{ is odd} \\ \log\text{-}\Pi_k^{\text{LD}} & \text{if } k \text{ is even.} \end{cases}$$

It is proved that if there exists $k \geq 0$ such that $\Lambda_{k+1} = \Lambda_k$, then $\Lambda_{k'} = \Lambda_k$ for all $k' \geq k$. That is, the hierarchy collapses at the k -th level. Moreover, there exists a distributed language on 0/1-labelled oriented paths that is outside the Λ_k -hierarchy, and thus outside $\log\text{-}\text{DH}^{\text{LD}}$. However, deciding whether a given solution to several

optimisation problems such as maximum independent set, minimum dominating set, maximum matching, max-cut, min-cut, traveling salesman, etc., is optimal are all in $\text{co-}\Lambda_1$, and thus in $\text{log-}\Pi_2^{\text{LD}}$. The absence of a non-trivial automorphism is proved to be in Λ_3 , that is $\text{log-}\Sigma_3^{\text{LD}}$ — recall that this language requires certificated of $\tilde{\Omega}(n^2)$ bits to be placed in Σ_1^{LD} (see [41]). It is however not known whether $\Lambda_3 \neq \Lambda_2$, that is whether $\text{log-}\Pi_2^{\text{LD}} \subset \text{log-}\Sigma_3^{\text{LD}}$ with a strict inclusion.

5.3 Distributed Graph Automata (DH^{DGA})

An analogue of the polynomial hierarchy, where sequential polynomial-time computation is replaced by distributed local computation was recently investigated in [50]. The model in [50] is called *distributed graph automata*. This model assumes a finite-state automaton at each node (instead of a Turing machine), and assumes anonymous computation (instead of the presence of unique node identities). Also, the model assumes an arbitrary interpretation of the outputs produced by each automaton, based on an arbitrary mapping from the collection of all automata states to $\{\text{true}, \text{false}\}$. The main result in [50] is that the hierarchy DH^{DGA} coincides with MSO on graphs.

6 Other Computational Models

6.1 Wait-Free Computing

The class WFD defined as the class of all distributed languages that are wait-free decidable was characterized in [36] as the class of languages satisfying the so-called *projection-closeness* property. For non projection-closed languages, [37] investigated more general interpretation of the individual opinions produced by the processes, beyond the logical conjunction of boolean opinions. In [35], it is proved that k -set agreement requires that the processes must be allowed to produce essentially k different opinions to be wait-free decided. The class Σ_1^{WFD} has been investigated in [38, 39], with applications to the space complexity of failure detectors. Interestingly, it is proved in [14] that wait-free decision finds applications to run-time verification.

6.2 Mobile Computing

The class MAD, for *mobile agent decision* has been considered in [34], as well as the class MAV = Σ_1^{MAD} , for *mobile agent verification*. It is proved that MAV has a complete language for a basic notion of reduction. The complement classes of

MAD and MAV have been recently investigated in [8] together with sister classes defined by other ways of interpreting the opinions of the mobile agents.

6.3 Quantum Computing

Distributed decision in a framework in which nodes can have access to extra resources, such as shared randomness, or intricate variables (in the context of quantum computing) is discussed in [2].

7 Conclusion

Distributed decision and distributed verification are known to have applications to very different contexts of distributed computing, including self-stabilization, randomized algorithms, fault-tolerance, runtime verification, etc. In this paper, our aim was to survey the results targeting distributed decision and verification per se. Beside the many interesting problems left open in each of the references listed in this paper, we want to mention two important issues.

Lower bounds in decision problems are often based on spatial or temporal arguments. Typically, the lack of information about far away processes, or the lack of information about desynchronized (or potentially crashed) processes, prevents processes to forge a consistent opinion about the global status of the distributed system. In the context of shared resources, such type of arguments appears however to be too weak (cf. [2]). Similarly, lower bounds in verification problems are often based on reduction to communication complexity theory. However, such reductions appear to be difficult to apply to higher classes in the local hierarchy, like separating the class at the third level from the class at the second level of the local hierarchy with $O(\log n)$ -bit certificates (cf. [21]).

This paper has adopted a systematic approach for presenting the results related to distributed decision and verification from the literature. This approach was inspired from sequential complexity and sequential computability theories. Such an approach provides a framework that enables to clearly separate decision from verification, as well as clearly separate the results obtained under different assumption (ID-oblivious, size of certificates, etc.). As already mentioned in [24], we believe that distributed decision provides a framework in which bridges between very different models might be identified, as decision tasks enables easy reductions between languages, while construction tasks are harder to manipulate because of the very different natures of their outputs.

References

- [1] Yehuda Afek, Shay Kutten, and Moti Yung. The local detection paradigm and its application to self-stabilization. *Theor. Comput. Sci.*, 186(1-2):199–229, 1997.
- [2] Heger Arfaoui and Pierre Fraigniaud. What can be computed without communications? *SIGACT News*, 45(3):82–104, 2014.
- [3] Heger Arfaoui, Pierre Fraigniaud, David Ilcinkas, and Fabien Mathieu. Distributedly testing cycle-freeness. In *40th International Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, pages 15–28, 2014.
- [4] Heger Arfaoui, Pierre Fraigniaud, and Andrzej Pelc. Local decision and verification with bounded-size outputs. In *15th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)*, pages 133–147, 2013.
- [5] Hagit Attiya and Jennifer Welch. *Distributed Computing: Fundamentals, Simulations, and Advanced Topics*. Wiley, 2004.
- [6] Baruch Awerbuch, Boaz Patt-Shamir, and George Varghese. Self-stabilization by local checking and correction (extended abstract). In *32nd Symposium on Foundations of Computer Science (FOCS)*, pages 268–277, 1991.
- [7] Alkida Balliu, Gianlorenzo D’Angelo, Pierre Fraigniaud, and Dennis Olivetti. Local distributed verification. *CoRR*, abs/1605.03892, 2016.
- [8] Evangelos Bampas and David Ilcinkas. On mobile agent verifiable problems. In *12th Latin American Symposium on Theoretical Informatics (LATIN)*, pages 123–137, 2016.
- [9] Mor Baruch, Pierre Fraigniaud, and Boaz Patt-Shamir. Randomized proof-labeling schemes. In *ACM Symposium on Principles of Distributed Computing (PODC)*, pages 315–324, 2015.
- [10] Mor Baruch, Rafail Ostrovsky, and Will Rosenbaum. Space-time tradeoffs for distributed verification. Brief Announcement at the 35th ACM Symposium on Principles of Distributed Computing, 2016.
- [11] Florent Becker, Adrian Kosowski, Martín Matamala, Nicolas Nisse, Ivan Rapaport, Karol Suchan, and Ioan Todinca. Allowing each node to communicate only once in a distributed system: shared whiteboard models. *Distributed Computing*, 28(3):189–200, 2015.
- [12] Lélia Blin and Pierre Fraigniaud. Space-optimal time-efficient silent self-stabilizing constructions of constrained spanning trees. In *35th IEEE International Conference on Distributed Computing Systems (ICDCS)*, pages 589–598, 2015.
- [13] Lélia Blin, Pierre Fraigniaud, and Boaz Patt-Shamir. On proof-labeling schemes versus silent self-stabilizing algorithms. In *16th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)*, pages 18–32, 2014.

- [14] Borzoo Bonakdarpour, Pierre Fraigniaud, Sergio Rajsbaum, David Rosenbluth, and Corentin Travers. Decentralized asynchronous crash-resilient runtime verification. Technical Report CAS-16-02-BB, Department of Computing and Software, McMaster University, 2016.
- [15] Keren Censor-Hillel, Eldar Fischer, Gregory Schwartzman, and Yadu Vasudev. Fast distributed algorithms for testing graph properties. *CoRR*, abs/1602.03718, 2016.
- [16] Keren Censor-Hillel, Petteri Kaski, Janne H. Korhonen, Christoph Lenzen, Ami Paz, and Jukka Suomela. Algebraic methods in the congested clique. In *ACM Symposium on Principles of Distributed Computing (PODC)*, pages 143–152, 2015.
- [17] Atish Das-Sarma, Stephan Holzer, Liah Kor, Amos Korman, Danupon Nanongkai, Gopal Pandurangan, David Peleg, and Roger Wattenhofer. Distributed verification and hardness of distributed approximation. *SIAM J. Comput.*, 41(5):1235–1265, 2012.
- [18] Danny Dolev, Christoph Lenzen, and Shir Peled. "tri, tri again": Finding triangles and small subgraphs in a distributed setting - (extended abstract). In *26th International Symposium on Distributed Computing (DISC)*, pages 195–209, 2012.
- [19] Yuval Emek, Christoph Pfister, Jochen Seidel, and Roger Wattenhofer. Anonymous networks: randomization = 2-hop coloring. In *ACM Symposium on Principles of Distributed Computing (PODC)*, pages 96–105, 2014.
- [20] Laurent Feuilloley and Pierre Fraigniaud. Randomized local network computing. In *27th ACM on Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 340–349, 2015.
- [21] Laurent Feuilloley, Pierre Fraigniaud, and Juho Hirvonen. A hierarchy of local decision. In *43rd International Colloquium on Automata, Languages and Programming (ICALP)*, 2016.
- [22] Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro. *Distributed Computing by Oblivious Mobile Robots*. Morgan & Claypool, 2012.
- [23] Klaus-Tycho Förster, Thomas Luedi, Jochen Seidel, and Roger Wattenhofer. Local checkability, no strings attached. In *17th International Conference on Distributed Computing and Networking (ICDCN)*, page 21, 2016.
- [24] Pierre Fraigniaud. Distributed computational complexities: are you Volvo-addicted or Nascar-obsessed? In *29th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 171–172, 2010.
- [25] Pierre Fraigniaud. Locality in distributed graph algorithms. In *Encyclopedia of Algorithms*, pages 1143–1148. Springer, 2016.
- [26] Pierre Fraigniaud, Cyril Gavoille, David Ilcinkas, and Andrzej Pelc. Distributed computing with advice: information sensitivity of graph coloring. *Distributed Computing*, 21(6):395–403, 2009.
- [27] Pierre Fraigniaud, Mika Göös, Amos Korman, Merav Parter, and David Peleg. Randomized distributed decision. *Distributed Computing*, 27(6):419–434, 2014.

- [28] Pierre Fraigniaud, Mika Göös, Amos Korman, and Jukka Suomela. What can be decided locally without identifiers? In *ACM Symposium on Principles of Distributed Computing (PODC)*, pages 157–165, 2013.
- [29] Pierre Fraigniaud, Magnús M. Halldórsson, and Amos Korman. On the impact of identifiers on local decision. In *16th International Conference Principles of Distributed Systems (OPODIS)*, pages 224–238, 2012.
- [30] Pierre Fraigniaud, Juho Hirvonen, and Jukka Suomela. Node labels in local decision. In *22nd International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, pages 31–45, 2015.
- [31] Pierre Fraigniaud, David Ilcinkas, and Andrzej Pelc. Communication algorithms with advice. *J. Comput. Syst. Sci.*, 76(3-4):222–232, 2010.
- [32] Pierre Fraigniaud, Amos Korman, and Emmanuelle Lebhar. Local MST computation with short advice. *Theory Comput. Syst.*, 47(4):920–933, 2010.
- [33] Pierre Fraigniaud, Amos Korman, and David Peleg. Towards a complexity theory for local distributed computing. *J. ACM*, 60(5):35, 2013.
- [34] Pierre Fraigniaud and Andrzej Pelc. Decidability classes for mobile agents computing. In *10th Latin American Symposium on Theoretical Informatics (LATIN)*, pages 362–374, 2012.
- [35] Pierre Fraigniaud, Sergio Rajsbaum, Matthieu Roy, and Corentin Travers. The opinion number of set-agreement. In *18th International Conference on the Principles of Distributed Systems (OPODIS)*, pages 155–170, 2014.
- [36] Pierre Fraigniaud, Sergio Rajsbaum, and Corentin Travers. Locality and checkability in wait-free computing. *Distributed Computing*, 26(4):223–242, 2013.
- [37] Pierre Fraigniaud, Sergio Rajsbaum, and Corentin Travers. On the number of opinions needed for fault-tolerant run-time monitoring in distributed systems. In *5th International Conference on Runtime Verification (RV)*, pages 92–107, 2014.
- [38] Pierre Fraigniaud, Sergio Rajsbaum, and Corentin Travers. Minimizing the number of opinions for fault-tolerant distributed decision using well-quasi orderings. In *12th Latin American Symposium on Theoretical Informatics (LATIN)*, pages 497–508, 2016.
- [39] Pierre Fraigniaud, Sergio Rajsbaum, and Corentin Travers. Perfect failure detection with very few bits. Submitted, 2016.
- [40] Pierre Fraigniaud, Ivan Rapaport, Ville Salo, and Ioan Todinca. Distributed testing of excluded subgraphs. *CoRR*, abs/1605.03719, 2016.
- [41] Mika Göös and Jukka Suomela. Locally checkable proofs. In *30th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 159–168, 2011.
- [42] Gene Itkis and Leonid A. Levin. Fast and lean self-stabilizing asynchronous protocols. In *35th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 226–239, 1994.

- [43] Jarkko Kari, Martín Matamala, Ivan Rapaport, and Ville Salo. Solving the induced subgraph problem in the randomized multiparty simultaneous messages model. In *22nd International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, pages 370–384, 2015.
- [44] Liah Kor, Amos Korman, and David Peleg. Tight bounds for distributed minimum-weight spanning tree verification. *Theory Comput. Syst.*, 53(2):318–340, 2013.
- [45] Amos Korman and Shay Kutten. Distributed verification of minimum spanning trees. *Distributed Computing*, 20(4):253–266, 2007.
- [46] Amos Korman, Shay Kutten, and Toshimitsu Masuzawa. Fast and compact self-stabilizing verification, computation, and fault detection of an MST. *Distributed Computing*, 28(4):253–295, 2015.
- [47] Amos Korman, Shay Kutten, and David Peleg. Proof labeling schemes. *Distributed Computing*, 22(4):215–233, 2010.
- [48] Moni Naor and Larry J. Stockmeyer. What can be computed locally? *SIAM J. Comput.*, 24(6):1259–1277, 1995.
- [49] David Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, 2000.
- [50] Fabian Reiter. Distributed graph automata. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 192–201, 2015.
- [51] Stefan Schmid and Jukka Suomela. Exploiting locality in distributed SDN control. In *Proceedings of the Second ACM Workshop on Hot Topics in Software Defined Networking (HotSDN)*, pages 121–126, 2013.