# Asynchronous Coordination Under Preferences and Constraints 

Armando Castañeda, Pierre Fraigniaud, Eli Gafni, Sergio Rajsbaum, Matthieu Roy

## - To cite this version:

Armando Castañeda, Pierre Fraigniaud, Eli Gafni, Sergio Rajsbaum, Matthieu Roy. Asynchronous Coordination Under Preferences and Constraints. 23rd International Colloquium on Structural Information and Communication Complexity, Jul 2016, Helsinki, Finland. hal-01341710

## HAL Id: hal-01341710 <br> https://hal.science/hal-01341710

Submitted on 4 Jul 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Asynchronous Coordination Under Preferences and Constraints* 

Armando Castañeda ${ }^{1}$, Pierre Fraigniaud ${ }^{2}$, Eli Gafni ${ }^{3}$, Sergio Rajsbaum ${ }^{1}$, and Matthieu Roy ${ }^{4}$

1 Instituto de Matemáticas, UNAM, Mexico.<br>2 CNRS and University Paris Diderot, France.<br>3 Computer Science Department, UCLA, USA.<br>4 LAAS-CNRS, Université de Toulouse, CNRS, Toulouse, France.


#### Abstract

Adaptive renaming can be viewed as a coordination task involving a set of asynchronous agents, each aiming at grabbing a single resource out of a set of resources totally ordered by their desirability. Similarly, musical chairs is also defined as a coordination task involving a set of asynchronous agents, each aiming at picking one of a set of available resources, where every agent comes with an a priori preference for some resource. We foresee instances in which some combinations of resources are allowed, while others are disallowed. We model these constraints, i.e., the restrictions on the ability to use some combinations of resources, as an undirected graph whose nodes represent the resources, and an edge between two resources indicates that these two resources cannot be used simultaneously. In other words, the sets of resources that are allowed are those which form independent sets in the graph. E.g., renaming and musical chairs are specific cases where the graph is stable (i.e., it the empty graph containing no edges).

As for musical chairs, we assume that each agent comes with an a priori preference for some resource. If an agent's preference is not in conflict with the preferences of the other agents, then this preference can be grabbed by the agent. Otherwise, the agents must coordinate to resolve their conflicts, and potentially choose non preferred resources. We investigate the following problem: given a graph, what is the maximum number of agents that can be accommodated subject to non-altruistic behaviors of early arriving agents?

We entirely solve this problem under the restriction that agents which cannot grab their preferred resources must then choose a resource among the nodes of a predefined independent set. However, the general case, where agents which cannot grab their preferred resource are then free to choose any resource, is shown to be far more complex. In particular, just for cyclic constraints, the problem is surprisingly difficult. Indeed, we show that, intriguingly, the natural algorithm inspired from optimal solutions to adaptive renaming or musical chairs is sub-optimal for cycles, but proven to be at most 1 to the optimal. The main message of this paper is that finding optimal solutions to the coordination with constraints and preferences task requires to design "dynamic" algorithms, that is, algorithms of a completely different nature than the "static" algorithms used for, e.g., renaming.


1998 ACM Subject Classification D.1.3 Concurrent Programming (Distributed programming); F.2.2 Nonnumerical Algorithms and Problems.

Keywords and phrases Adaptative renaming, resource allocation, asynchronous computing.
Digital Object Identifier 10.4230/LIPIcs.xxx.yyy.p

[^0]
© Armando Castañeda, Pierre Fraigniaud, Eli Gafni, Sergio Rajsbaum, and Matthieu Roy; licensed under Creative Commons License CC-BY

## 1 Introduction

### 1.1 Context and objective

In distributed computing, several tasks have their adaptive versions in which the quality of the solution must depend only on the number of processes that participate in a given execution, and not on the total number of processes that could be involved in this task. A typical example of an adaptive task is adaptive renaming [4]. In renaming, each process is aiming at acquiring a name taken from a small range of integers $[1, r]$, under the constraint that all acquired names must be pairwise distinct. The quality of a renaming algorithm is judged based on the range $r$ of names, the smaller the better. In adaptive renaming, $r$ must depend only on the number $k$ of participating processes. In the asynchronous setting with crash-prone processes and read/write registers, the optimal value for the range is known to be $r=2 k-1[5,13]$.

Interestingly, adaptive renaming can also be viewed as a task by interpreting the integers $1, \ldots, r$ as a total order on the names, where name $i$ is preferred to name $j$ whenever $i<j$. Hence, adaptive renaming can be viewed as an abstraction of the problem in which asynchronous agents are competing for resources totally ordered by their desirability. In other words, adaptive renaming is an abstraction of a problem of coordination between agents under preferences. Coordination between agents under preferences has been recently investigated in $[2,3]$ where the musical chairs game has been formally defined and solved. In this game, a set of players (modeling the agents) must coordinate so that each player eventually picks one of the available chairs (modeling the resources). Each player initially comes with an a priori preference for one chair. In absence of conflict with other players, the player can pick the desired chair, otherwise the conflicting players must coordinate so that they pick different chairs. It is proved that the smallest number $r$ of chairs for which musical chairs with $k$ players has a solution is $r=2 k-1$.

We foresee that neither adaptive renaming nor musical chairs fully capture typical scenarios of agents competing for resources. Indeed, both tasks only capture scenarios in which the constraint is that no two agents can acquire the same resource. In practice, resources may not be independent, and the literature on scheduling, partitioning, resource allocation, etc. (see, e.g., $[6,7,11,15,16]$ ) provide several examples of problems in which resources are inter-dependent, causing some resource $a$ not being allowed to be used simultaneously with resource $b$. That is, using one resource disables others. In this paper, we consider the case in which constraints are modeled as an indirected graph whose nodes are resources, and every edge $\{a, b\}$ indicates that resources $a$ and $b$ cannot be both simultaneously acquired, i.e., acquiring a node disables all its neighbors. In other words, the sets of resources that are allowed are those which form independent sets in the graphs. In this framework, renaming as well as musical chairs correspond to the case where the graph of constraints is a stable one (i.e., a graph with no edges). We thus address an extension of renaming and musical chairs, targeting an abstraction of a problem of coordination between agents under constraints.

Our objective is to understand the power and limitation of coordination between agents competing for interdependent resources. We are focussing on a scenario inspired from musical chairs in which a resource is a priori assigned to each agent, and the agents have to coordinate between them so that to eventually acquire pairwise non conflicting resources. In particular, if the initial assignment forms an independent set, then the agents do not have to do anything. Alternatively, if they are initially assigned conflicting resources, then they have to spread out and coordinate themselves so that they eventually acquire a set of resources that form an independent set. In other words, each agent comes with an a priori preference
for some resource - these preferences for the resources do not need to be different. If an agent's preference is not in conflict with the preference of another agent, then it can grab its preference. Otherwise, this agent must choose another resource.

The coordination task between agents under preferences and constraints is thus defined as follows. Given an $n$-node graph $G=(V, E)$ modeling the constraints between the resources, an input is a multiset $M$ of $k$ elements in $V$ representing the preferences of $k$ processes $p_{1}, \ldots, p_{k}$ modeling the $k$ agents. Outputs are independent sets $I=\left\{u_{1}, \ldots, u_{k}\right\}$ in $G$, of size $k$ representing the fact that process $p_{i}$ acquires $u_{i}$, for $i=1, \ldots, k$. The literature on renaming [5] and musical chair [3] taught us that, in an asynchronous system in which the processes are subject to crash failures, the task is not solvable for $k$ larger than some bound, even for the stable graph $G$ (the value of the bound on $k$ for the stable graph is roughly half the number of nodes of the graph). We are interested in the impact of the constraints on this bound. That is, given a graph $G$, we are interested in the largest $k$ for which the coordination with constraints and preferences task in $G$ is solvable for every preference multiset $M$ of size at most $k$. We focus on asynchronous systems in which an arbitrarily large number of processes are subject to crash failures. Each process has its own private registers, and the processes communicate via read/write accesses to a shared memory.

### 1.2 Our results

We first focus on the problem for the $n$-node path $P_{n}$ because it enables to prove a lower bound on the size of Hamiltonian graphs for which the coordination with constraints and preferences task is solvable. Interestingly, this lower bound is almost twice as large as the $2 k-1$ bound without constraints resulting from renaming or musical chairs. Specifically, we establish the following:

- Theorem 1. Let $k$ be a positive integer. The smallest integer $n$ for which the coordination with constraints and preferences task in $P_{n}$ is solvable for $k$ processes satisfies $n=4 k-3$. As a consequence, if the coordination with constraints and preferences task in an n-node Hamiltonian graph $G$ is solvable for $k$ processes then $n \geq 4 k-3$.

The lower bound on $n$ is based on a reduction to impossibility results for musical chairs, i.e., renaming with initial preferences. The upper bound on $n$ comes from a wait-free algorithm, inspired from an optimal adaptive renaming algorithm, whose main lines are: (1) fix a maximum independent set $I$ in $P_{n},(2)$ index the vertices of $I$ from 1 to $2 k-1$, and (3) run an optimal (adaptive) renaming algorithm on these indexes.

From this preliminary result on $P_{n}$, one may think that solving the coordination with constraints and preferences task in a graph $G$ boils down to classical renaming once a maximum independent set in $G$ is fixed. We show that this is not the case. In fact, even for an instance as simple as the $n$-node ring $C_{n}$, the problem becomes highly non trivial.

- Theorem 2. Let $k$ be a positive integer. The smallest $n$ for which the coordination with constraints and preferences task in $C_{n}$ is solvable for $k$ processes satisfies $4 k-3 \leq n \leq 4 k-2$.

The lower bound is a consequence of Theorem 1 since $C_{n}$ is Hamiltonian. A quite intriguing fact is that the wait-free algorithm derived from an adaptation of an optimal algorithm for classical renaming run on a maximum independent set of $C_{n}$ does not match the lower bound, and is off by an additive factor +1 . In fact, we prove that the true answer is probably the lower bound $4 k-3$, which is shown to be tight for $k=2$ and 3 agents, using ad hoc algorithms that are radically different from renaming algorithm.

We believe that the difference of 1 between the lower and upper bounds for $C_{n}$ is certainly not anecdotal, but is the witness of a profound phenomenon that is not yet understood, with potential impact on classical renaming and musical chairs. The main outcome of this paper is probably the observation that "static" algorithms, i.e., algorithms based on fixed precomputed positions in the graph of constraints, might be sub-optimal by allocating less resources than the optimal. Our optimal ad hoc algorithms for coordinating two or three processes in the ring are not static, and the set of allocated resources output by these algorithms can form any independent set. The design of optimal "dynamic" (i.e., non static) algorithms for solving the coordination with constraints and preferences task appears to be a challenge, even in the specific case of the cycle $C_{n}$.

The enormous difficulty for asynchronous crash-prone processes to coordinate under constraints and preferences, even in graphs with arbitrarily large independent sets, is also illustrated by the case of the complete bipartite graph $K_{x, y}$ with $n=x+y$ nodes. We show that, although $K_{x, y}$ has very large independent sets (of size at least $\min \{x, y\}$ ), processes cannot coordinate at all in this graph.

- Theorem 3. Let $x, y$ be positive integers. Coordination with constraints and preferences in the complete bipartite graph $K_{x, y}$ is unsolvable for more than one process.

Finally, on the positive side, given any graph $G$, we can design an static algorithm ALG solving the coordination with constraints and preferences task in $G$. ALG is based on the novel notion of $k$-admissible independent sets, which may have its interest on its own: given $G=(V, E)$, an independent set $I$ of $G$ is $k$-admissible if for every $W \subseteq V$ of size at most $k-1$, we have $|I \backslash N[W]| \geq|I \cap W|+1$ where $N[W]$ denotes the set of nodes at distance at most 1 from a node in $W$. We prove that among static algorithms, ALG is optimal, which completely closes the problem for static algorithms.

- Theorem 4. Let $G$ be a graph, and $k$ be a positive integer. Let $I$ be a $k$-admissible independent set in $G$. Then, ALG instantiated with I solves the coordination with constraints and preferences task in $G$ with $k$ processes. Moreover, if $G$ has no $(k+1)$-admissible independent set, then no static algorithms can solve the coordination with constraints and preferences task in $G$ with more than $k$ processes.


### 1.3 Related Work

Since its introduction, the renaming problem has been extensively studied (see for example [5, $10,19]$ ). It was initially introduced as a non-adaptive problem in which processes just need to pick distinct output names in the space $[1, \ldots, M]$, where $M$ is on function only on the total number of processes that might participate [4]. Several algorithms were proposed (e.g. [4, 9, 14]), and, as far as we know, all those initial algorithms are adaptive. Then the adaptive version of renaming was coined. The study of lower bounds for renaming have inspired new developments in topology techniques (see [17] for a detailed description). As explained above, the variant of renaming that is closest to this paper is the adaptive renaming version. This is the version we use to solve the coordination with constraints and preferences task on a graph with no edges.

Musical chairs [2, 3] is a coordination problem on a stable graph in which each process starts with a initial vertex (chair) and processes are required to decide distinct vertices (chairs). The problem is studied in a model where the only communication between processes is an indication when two processes propose the same vertex. It has been shown that $k$ processes can solve the problem only if the stable graph has at least $2 k-1$ vertices. It has been also shown that musical chairs and adaptive renaming are equivalent problems.

Interestingly, the coordination with constraints and preferences is also related to mobile computing [12], where mobile entities (modeling physical robots, software agents, insects, etc.) must cooperate in order to solve tasks such as rendezvous, gathering, exploration, patrolling, etc. In particular, in the asynchronous look-compute-move model of mobile computation, the "look" operation is very similar to a "snapshot" operation in shared memory, and the "move" operation is very similar to the "write" operation. The major differences between the wait-free model of distributed computation and the look-compute-move model of mobile computation are (1) the presence of failures in the former, and (2) the fact that agents are moving in an anonymous graph in the latter.

## 2 Model and examples

### 2.1 Computational model

We consider the standard asynchronous wait-free read/write model with $k$ processes, $p_{1}, \ldots, p_{k}$ [5, 19]. Processes are asynchronous, communicate by writing and reading from a reliable shared memory, and any set of processes may crash. We assume, without loss of generality, that processes can read the whole shared memory in a single atomic snapshot [1].

Problems in the wait-free model are usually defined as tasks, where processes get input values, and decide after a bounded number of operations on output values, such that the decided values represent a valid configuration associated to the initial values of the execution for this task.

Without loss of generality, we can assume that algorithms solving tasks are in normal form, that is, are of the form of a loop consisting of (1) writing to the shared memory the local state of the process, (2) taking a snapshot, and (3) performing some local computation. The loop is executed until the process returns an output (i.e., decides).

### 2.2 Coordination with Constraints and Preferences (CCP)

The task coordination with constraints and preferences (or CCP for short) is instantiated by a fixed $n$-node graph $G=(V, E)$. The graph is modeling the constraints. It is supposed to be simple, i.e., without loops and multiple edges, but does not need to be connected. Each process $p_{i}$ gets as input one vertex $u \in V$, called its initial private preference, and must eventually decide on a vertex $v \in V$. It is required that the decided vertices form an independent set of $G$, that is, no two processes decide the same vertex, and no two decided vertices belong to the same edge. It is also required that if the initial preferences form an independent set, then each process must decide its initial preference (enforcing the fact that processes cannot discard their preferences).

We are interested in computing, for every $n$-node graph $G$, the largest $k$ such that CCP in $G$ is wait-free solvable for $k$ processes. Note that an algorithm solving CCP in $G$ is given the full description of $G$ a priori. Hence, there are no issues such as, e.g., breaking symmetry between nodes of $G$, even if $G$ is vertex-transitive. (In particular, the nodes of $G$ might be given labels from 1 to $n$, a priori, in a specific order which may facilitate the task for the processes).

### 2.3 Examples and Basic Observations

CCP is trivially solvable for one process in every graph, by selecting its initial preference as output vertex. Also, CCP is trivially not solvable in $G$ for $k$ processes if $k$ exceeds the size of a maximum independent set. In fact, CCP is not solvable in $G$ for $k$ processes if $k$
exceeds the size of the smallest maximal independent set. Indeed, let $I=\left\{u_{1}, \ldots, u_{\ell}\right\}$ be a smallest maximal independent set in $G$, and assume that $\ell$ processes $p_{1}, \ldots, p_{\ell}$ are given preferences in $I\left(u_{i}\right.$ to $p_{i}$ for $\left.i=1, \ldots, \ell\right)$. In a wait-free execution in which only those $\ell$ processes participate, they must decide $I$. If another process $p_{\ell+1}$ "wakes up" after the $\ell$ processes have decided, there is no more room for $p_{\ell+1}$ to acquire a vertex, because $I$ is maximal. This holds even if there exists another independent set $I^{\prime}$ larger than $I$, since the first $\ell$ processes have already terminated.

The following result is a direct consequence of [3] as Musical Chairs is exactly our problem on the stable graph.

- Proposition 5. Let $k$ be a positive integer. The smallest integer $n$ for which the coordination with constraints and preferences task in the n-node stable graph is solvable for $k$ processes satisfies $n=2 k-1$.

Also, we have the following observation.

- Proposition 6. Let $G=(V, E)$ be a graph, and $G^{\prime}=\left(V, E^{\prime}\right)$ with $E^{\prime} \subseteq E$ be a subgraph of $G$. If the coordination with constraints and preferences task is solvable for $k$ processes in $G$, then it is solvable for $k$ processes in $G^{\prime}$.

As a consequence of the above two propositions, we get a general lower bound on the size of graphs in which the coordination with constraints and preferences task is solvable for $k$ processes.

- Corollary 7. Let $G$ be an n-node graph. If the coordination with constraints and preferences task in $G$ is solvable for $k$ processes then $2 k-1 \leq n$.


## 3 The Case of Cyclic Constraints

Our first results concern simple non-trivial sets of constraints, namely the cases of the $n$-node paths and cycles, respectively denoted $P_{n}$ and $C_{n}$. The case of the path is entirely solved by the following results, which establish Theorem 1:

- Proposition 8. Let $k$ be a positive integer. The smallest integer $n$ for which the coordination with constraints and preferences task in the n-node path is solvable for $k$ processes satisfies $n=4 k-3$.

Proof. Let us assume, for the purpose of contradiction, that there is an algorithm $\mathcal{A}$ solving CCP in the $n$-node path for $k$ processes with $n=4 k-4$. Such a path has a maximum matching $M$ of size $2 k-2$. $\mathcal{A}$ guarantees that, for every edge of $M$, at most one process acquires an extremity of that edge. $\mathcal{A}$ can be used as a subroutine to solve CCP on a stable graph of size $2 k-2$, as we show below.

We assume that the $n$-path is oriented from left to right and hence for each edge in $M$ there is a left vertex and a right vertex (thus, in the path, not two left (right) vertices are adjacent). Also, each vertex $v$ of the stable graph of size $2 k-2$ is mapped to a unique edge $f(v)$ in $M$. To solve CCP on the stable graph, each process $p_{i}$ with initial preference $v$, invokes $\mathcal{A}$ with the left vertex of $f(v)$ and decides $f^{-1}(e)$, where $e$ is the edge in $M$ containing the vertex that $\mathcal{A}$ outputs to $p_{i}$ The resulting algorithm $\mathcal{A}$ solves CCP on the stable graph of size $2 k-2$ because if processes start with distinct vertices, then all of them invoke $\mathcal{A}$ with left vertices and hence, each process decide its initial preference. If processes start with conflicting initial preferences, then $\mathcal{A}$ outputs vertices that belong to distinct edges in $M$. This is a contradiction with Proposition 5 because $M$ is of size $2 k-2$.

We now describe an algorithm solving CCP in the $n$-node path for $k$ processes with $n=4 k-3$. The algorithm is based on a maximum independent set $I$ of size $2 k-1$ in $P_{4 k-3}$. That is, the nodes of $P_{4 k-3}$ are labeled off line as $\left(v_{1}, v_{2}, v_{3}, \ldots, v_{4 k-3}\right)$, and we define $I=\left\{v_{1}, v_{3}, v_{5}, \ldots, v_{4 k-3}\right\}$. Essentially, the algorithm runs the textbook renaming algorithm of [5] on $I$, adapted to handle initial preferences. Indeed, selecting a node $w \notin I$ may block two positions in $I$ (the two neighboring nodes of $w$ ). Nevertheless, there is still enough room to perform renaming, and hence to solve CCP. Indeed, let $N[w]$ be the closed neighborhood of $w$ in $P_{n}$, i.e., the at most three nodes at distance at most 1 from $w$, and, for a set of nodes $W$, let $N[W]=\cup_{w \in W} N[w]$. In classical renaming, if $W$ is the multiset of currently chosen names, there remain at least $2 k-1-|W| \geq k$ available names to choose from. In the path, if $W$ is the multiset of currently chosen nodes, there are only $|I \backslash N[W]|$ available nodes in $I$ to choose from, and this number of available nodes can be less than $k$. However, the crucial observation is that $|I \backslash N[W]|\left|>|I \cap W|\right.$ in the path $P_{4 k-3}$, for any set of nodes $W$ of size at most $k-1$. Hence, there are more free nodes in $I$ than occupied nodes in $I$, and thus the idea is to perform the ranking of the renaming algorithm only on processes sitting on the occupied nodes of $I$. Since $|I \backslash N[W]||>|I \cap W|$, this ranking is valid, that is, systematically provides a position in $I \backslash N[W]$. Termination follows from classical arguments by assuming, by way of contradiction, that some processes do not terminate, and then by considering the process $p$ with lowest ID that does not terminate. Eventually, the rank $r$ of $p$ will remain forever the same, and no other processes that do not terminate will conflict with the $r$ th node in the subset of nodes in $I$ that are not conflicting with terminated processes. At this point, process $p$ terminates.

As a consequence of this result combined with Proposition 6, we get a general lower bound on the size of Hamiltonian graphs in which the coordination with constraints and preferences task is solvable for $k$ processes. Interestingly, this bound is roughly twice as big as the bound for arbitrary graphs (cf. Corollary 7).

- Corollary 9. Let $G$ be an n-node Hamiltonian graph. If the coordination with constraints and preferences task in $G$ is solvable for $k$ processes then $4 k-3 \leq n$.

The case of $P_{n}$ has attracted our interest for it enables deriving bounds for Hamiltonian graphs. The case of the cycle $C_{n}$ may seem to behave quite similarly as $P_{n}$. Surprisingly, this is not the case, as the wraparound constraint yields an interesting phenomenon, namely "static" solutions inspired from renaming algorithms such as the ones for the stable graph and the path are not anymore optimal in term of number of processes, and are off by an additive factor +1 from the optimal. More precisely, we show the following, which establishes Theorem 2:

- Proposition 10. Let $k$ be a positive integer. The smallest integer $n$ for which the coordination with constraints and preferences task in $C_{n}$ is solvable for $k$ processes satisfies $4 k-3 \leq n \leq 4 k-2$.

Proof. The lower bound follows directly from Corollary 9. The upper bound is directly derived from the algorithm in the proof of Proposition 8 by fixing a maximum independent set of size $2 k-1$ in the cycle $C_{4 k-2}$. The correctness of the algorithm follows from the same arguments as for the path $P_{4 k-3}$.

Interestingly, the lower bound $4 k-3$ is most probably the right answer, and not the upper bound $4 k-2$. At least, this is the case for small numbers of processes:


Figure $1 C C P$ algorithm for two processes in $C_{5}$ : rules when two processes are executing

- Proposition 11. The smallest integer $n$ for which the coordination with constraints and preferences task in $C_{n}$ is solvable for $k$ processes satisfies $n=4 k-3$ for $k=2$ and $k=3$.

Proof. The nodes of $C_{4 k-3}$ are sequentially labeled offline as $v_{1}, v_{2}, \ldots, v_{4 k-3}$. This labeling induces a clockwise direction (increasing labels) and a counterclockwise direction.

The algorithm for two processes in $C_{5}$ is depicted on Figure 1, which represents the snapshot of a process, and the action to take (represented as arrows) based on this snapshot when 2 processes participate. There are three cases, depending on whether the two processes are currently occupying nodes at distance 0,1 , or 2 . (Of course, if the snapshot reveals that the process is alone, then it decides the node that it currently occupies, i.e., its preferred node). If the snapshot reveals that the two processes occupy the same node, then the action depends on the ID: going clockwise for the process with smallest ID, and counterclockwise otherwise. If the snapshot reveals that the two processes occupy two neighboring nodes, then the action is: going away from the other node. Finally, if the snapshot reveals that the two processes occupy two nodes at distance 2 , then the action is to decide the currently occupied node. One can check that this asynchronous algorithm terminates, and wait-free solves CCP.

In the appendix, we provide a similar algorithm for three processes in $C_{9}$.
From these two cases, we conjecture that the smallest cycle $C_{n}$ enabling to solve CCP is $n=4 k-3$, for all $k \geq 2$. If this is correct, it means that optimality requires processes to coordinate in a more complex way than they do for renaming, in order to spread out optimally in the graph of constraints, and eventually occupy a large number of nodes. The independent set they will eventually agree on cannot be decided a priori, but the processes must agree on line in order to eventually decide an independent set that fits their initial preferences, the constraints and the uncertainty resulting for asynchrony and failures.

## 4 A Generic Algorithm

The algorithms inspired by the original algorithm of [5] used in the proofs of Theorems 1 and 2 to establish the upper bounds on the smallest integer $n$ for which the coordination with constraints and preferences task in the $n$-node path and in the $n$-node cycle are static in the sense that they are aiming at deciding within a fixed independent set. More precisely:

- Definition 12. Let $G$ be a graph, and $I$ be an independent set in $G$. An algorithm $\mathcal{A}$ solving the coordination with constraints and preferences task in $G$ is static with respect to $I$ if, for every execution of $\mathcal{A}$, and for every process $p$, if $p$ does not decide its initial input, then it decides a vertex in $I$.

To establish Theorem 4, we present a generic static algorithm to solve CCP on every graph $G=(V, E)$, and prove that this algorithm is the best possible static algorithm in the sense that it maximizes the number $k$ of processes for which CCP is solvable in $G$.

The generic algorithm is instantiated with an a priori ordered independent set $I$. That is, $I=\left\{w_{1}, \ldots, w_{|I|}\right\}$, and this ordering of the nodes in $I$ is know a priori to every process. The processes proceed in a sequence of (asynchronous) rounds. At each round, every process $p_{i}$ proposes a current vertex denoted $c u r_{i}$ (the first proposal is the input $u_{i}$ of processes $p_{i}$ ). Then, $p_{i}$ checks whether there is a conflict with other proposals. In absence of conflict, $p_{i}$ decides its current proposal cur $_{i}$. If there is a conflict, $p_{i}$ computes a new proposal in $I$, and repeats. Hence, in particular, if a process sees no conflict in its initial proposal, then it stays there. Otherwise it will try a new proposal in $I$. The new proposal of $p_{i}$ is computed within the "free space" that is defined as the maximal subset of $I$ such that there is no conflict with other processes' proposals.

Algorithm 1 is the pseudocode of the generic algorithm. The algorithm uses a shared array view, accessed with write and snapshot operations, where each entry is initially $\perp$. For convenience, it is easier to consider the array view as a multiset of nodes. The local variable cur $_{i}$ stores the current proposal of process $p_{i}$. For a set $W \subseteq V$, we denote by $N[W]$ the closed neighborhood of $W$, that is, for $w \in V, N[w]=\{w\} \cup\{v \in V:\{v, w\} \in E\}$, and $N[W]=\cup_{w \in W} N[w]$.

```
Algorithm \(1 G=(V, E)\) is a graph, and \(I\) is an ordered independent set in \(G\). Code for \(p_{i}\).
function IndependentSet ( \(u_{i} \in V\) : initial preference of \(p_{i}\) )
    cur \(_{i} \leftarrow u_{i}\)
    loop
        write \(\left(\right.\) cur \(\left._{i}\right)\)
        snapshot memory to get view \(=\left\{\right.\) cur \(_{j_{1}}, \ldots\), cur \(\left._{j_{r}}\right\} \quad \triangleright\) multiset of \(r\) elements, for some \(r\)
        view \(^{\prime} \leftarrow\) view \(\backslash\left\{\right.\) cur \(\left._{i}\right\} \quad \triangleright\) remove one occurence of cur \(_{i}\) from view
        if view' \(^{\prime} \cap N\left[\right.\) cur \(\left._{i}\right]=\emptyset\) then \(\quad \triangleright\) check for conflicts
            return cur \(_{i} \quad \triangleright\) no conflict \(\Rightarrow\) decide cur \(_{i}\)
        else \(\quad \triangleright\) conflict detected \(\Rightarrow\) compute a new position
            free \(\leftarrow I \backslash N\left[\right.\) view \(\left.^{\prime}\right] \quad \triangleright\) rule out conflicting vertices from \(I\)
            \(\ell \leftarrow \mid\left\{s: \operatorname{cur}_{j_{s}} \in I\right.\) and \(\left.j_{s}<i\right\} \mid+1 \quad \triangleright\) ranking on the currently occupied vertices of \(I\)
            \(c^{\leftarrow} r_{i} \leftarrow \ell^{t h}\) element in free \(\quad \triangleright\) try the \(\ell^{t h}\) free node for the next round
        end if
    end loop
```

Note that if the initial preferences of participating processes are distinct and form an independent set, then Algorithm 1 guarantees that each process decides on its initial preference. Note also that if two processes decide on vertices $v_{1}, v_{2}$, then $v_{1} \neq v_{2}$ and $\left\{v_{1}, v_{2}\right\} \notin E$. However, for Algorithm 1 to function appropriately, we use it for a specific kind of independent sets, namely admissible independent sets, defined hereafter (see Figure 2 for an illustration):

- Definition 13. Let $k$ be a positive integer, and $G=(V, E)$ be a graph. An independent set $I$ of $G$ is $k$-admissible if for every $W \subseteq V$ of size at most $k-1$, we have $|I \backslash N[W]| \geq|I \cap W|+1$.

Notice that any $k$-admissible independent set $I$ satisfies $|I| \geq 2 k-1$ (instantiate Definition 13 with $W \subseteq I$ of size $k-1$ ). To establish Theorem 4, we first prove the following result.

- Proposition 14. Let $G$ be a graph, and $k$ be a positive integer. Let I be a $k$-admissible independent set in $G$. Then, Algorithm 1 instantiated with $I$ solves the coordination with constraints and preferences task in $G$ with $k$ processes.

Proof. We have seen that the safety conditions (i.e., respect of the preferences, and take decisions on an independent set) are satisfied. It just remains to show that the algorithm


Figure 2 To be $k$-admissible, $I$ must satisfy $|I \backslash N[W]| \geq|I \cap W|+1$ for every $W$ of size $\leq k-1$.
is valid (i.e., whenever a process detects a conflict on Line 6 , it is able to compute a new consistent preference), and terminate. We first show validity.

- Claim 1. For any process that is about to execute Line 11, $\mid$ free $\mid \geq \ell$.

Consider a process $p_{i}$ that is about to execute Line 11. Such a process $p_{i}$ must have detected a conflict Line 6. Let view be the snapshot of $p_{i}$ associated with this conflict, and $W=v i e w^{\prime}$. When $p_{i}$ is about to execute Line 11, we have free $=I \backslash N[W]$. As there are at most $k$ participating processes, it follows that $|W|<k$. Since $I$ is $k$-admissible, we have $|I \backslash N[W]| \geq|W \cap I|+1$, which implies $\mid$ free $\left|\geq|W \cap I|+1\right.$. Moreover, $p_{i}$ 's ranking computed in line 10 is at most $|W \cap I|+1$ because $W$ does not contain $p_{i}$ 's preference, i.e., one ignores the nodes not in $I$ when ranking. Hence $\mid$ free $\mid \geq \ell$, and thus the algorithm is valid.

- Claim 2. Algorithm 1 terminates.

Assume, for the sake of contradiction, that there is a run $\alpha$ in which some processes take an infinite number of steps without deciding a vertex. In this execution, let $P$ be the set of all processes taking infinitely many steps, and let $p \in P$ be the process with minimum ID in $P$. Consider a suffix $\alpha^{\prime}$ of $\alpha$ in which (1) all processes that do not run forever have stopped, and (2) all processes in $P$ have already tried once to get a vertex in $I$, namely, they have executed Line 11 at least once. Note that such a suffix exists because every process that takes an infinite number of steps in $\alpha$ will see an infinite number of conflicts, and thus will eventually always execute Line 3 with its current vertex in $I$.

In $\alpha^{\prime}$, the rank of $p$ is fixed because, from there on, the set of processes that occupy vertices in $I$ is fixed. Let $r$ be the rank of $p$ among the processes of $\alpha^{\prime}$ with proposals in $I$, and let good be the set of vertices in $I$ that do not conflict with preferences of stopped processes in $\alpha^{\prime}$. Eventually, there are no processes in $P$ that are proposing one of the first $r-1$ elements of good. Indeed, the rank of every process in $P$ is at least $r$, and when a process $q \in P$ takes a snapshot, the set free that it computes satisfies free $\subseteq$ good. Thus, $q$ proposes the $x$-th element in free, where $x \geq r$. Hence, this element cannot be one of the first $r-1$ elements in good.

It follows that, in $\alpha^{\prime}$, as soon as all running processes have seen each other, and have written at least twice, when these processes compute free, the first $r$ elements are all elements of good, and only $p$ will try to get the $r$-th element in free. When it does, it detects no conflict. Thus $p$ can terminate the algorithm on Line 3 , which yields to a contradiction.

We now show the second part of Theorem 4, that is, Algorithm 1 is optimal on the number of processes, among static CCP algorithms. This is established thanks to the following result.

- Proposition 15. Let $G$ be a graph, $k$ be a positive integer, and assume that $G$ has no $(k+1)$-admissible independent set. Then, no static algorithm can solve the coordination with constraints and preferences task in $G$ with more than $k$ processes.

The proof is based on the following claim, which is an interesting consequence of the Wait-free Computability Theorem of [18]. A concise proof is presented in appendix.

- Claim 3. If there is a $k$-process CCP algorithm for $G$, then there exists a $k$-process CCP algorithm for $G$ in which whenever a process sees only itself in its first snapshot, it immediately decides.

Proof. For contradiction, assume there is a static CCP algorithm $A$ on $G$ for $k+1$ processes, with respect to a non $(k+1)$-admissible independent set $I$. By Claim 5 , we can assume that if a process sees only itself in its first snapshot, then it decides its initial value (recall that $A$ can be assumed to be in normal form). Since $I$ is not $(k+1)$-admissible, there exists a set $W \subseteq V$ with at most $k$ vertices such that $|I \backslash N[W]|<|W \cap I|+1$.

- Claim 4. $W \cap I \neq \emptyset$.

Proof of Claim 4: Suppose $W \cap I$ is empty. It follows that $|W \cap I|+1=1$, and thus $|I \backslash N[W]|=0$, i.e., $I \backslash N[W]$ is empty as well. Thus $I \cap N[W]=I$. Now, let us consider an execution $\alpha$ of $A$ in which processes $p_{1}, \ldots, p_{|W|}$ start with preferences for distinct vertices in $W$, and just write their input in shared memory, i.e., they only perform one write operation in shared memory. Then, consider any extension $\beta$ of $\alpha$ in which a process $q$ starts with a preference to a vertex in $I$, runs alone (the $p_{i}$ 's from $\alpha$ do not take steps), and decides. Note that such a process exists since $|W| \leq k$ and $A$ is for $k+1$ processes. Since $A$ is static with respect to $I, q$ decides a vertex $v \in I$. Let $v^{\prime} \in W$ that is adjacent to $v$, whose existence is guaranteed by $I \cap N[W]=I$, and let $q^{\prime}$ be the process in $\alpha$ with input $v^{\prime}$. Let $\alpha^{\prime}$ be a reordering of $\alpha$ in which $q^{\prime}$ executes solo, decides, then all other processes in $\alpha$ do their write operation. $q^{\prime}$ sees only itself in its first snapshot, and will thus decide its input, from Claim 5. As $q$ cannot distinguish $\alpha$ from $\alpha^{\prime}, \beta$ is also a valid extension of $\alpha^{\prime}$. In $\alpha^{\prime} \beta$, process $q$ decides $v$. This is a contradiction, since $q^{\prime}$ decides $v^{\prime}$ in this execution, which is in conflict with $v$. Thus, $W \cap I \neq \emptyset$, which completes the proof of Claim 4.

Let $\rho=|W \cap I|$. By Claim 4, $\rho>0$. We have $|I \backslash N[W]|<|W \cap I|+1=\rho+1$, and thus $|(I \backslash N[W]) \cup(W \cap I)|<2 \rho-1$. Let us consider an execution $\alpha$ of $A$ in which processes $p_{1}, \ldots, p_{|W \backslash I|}$ start with preferences for distinct vertices in $W \backslash I$, and they just write their input in the shared memory. Using the same indistinguishability argument as in the proof of Claim 4, we can show that that there is no extension of $\alpha$ in which a process $q \notin\left\{p_{1}, \ldots, p_{|W \backslash I|}\right\}$ decides a vertex in $I \cap N[W]$. Thus, in every extension of $\alpha$, such a process $q$ decides its initial preference vertex, or a vertex in $(I \backslash N[W]) \cup(W \cap I)$. Let us then consider processes $q_{1}, \ldots, q_{\rho+1}$, all distinct from $p_{1}, \ldots, p_{|W \backslash I|}$ (note that these processes exist because $|W| \leq k)$. In every extension of $\alpha$ in which $q_{1}, \ldots, q_{\rho+1}$ start with preferences in $(I \backslash N[W]) \cup(W \cap I)$, and these processes decide, they necessarily decide vertices in $(I \backslash N[W]) \cup(W \cap I)$. This implies that $|(I \backslash N[W]) \cup(W \cap I)| \geq \rho+1$.

Now, from $A$ and execution $\alpha$, we construct a $p$-process algorithm $A^{\prime}$ that solves CCP in the stable graph with too few vertices as follows. In $A^{\prime}$, the shared memory has the state in $\alpha$, and each process $q_{i}, 1 \leq i \leq \rho+1$, follows the same code as in $A$. To see that $A^{\prime}$ is correct, note that (1) if processes start with distinct vertices of $(I \backslash N[W]) \cup(W \cap I)$, then each process decides its input, since $A$ is correct, and (2) if processes start with inputs in conflict, then they decide distinct vertices in $(I \backslash N[W]) \cup(W \cap I)$, as shown before. Now, $A^{\prime}$ solves CCP for $\rho+1$ processes over the stable graph with vertex set $(I \backslash N[W]) \cup(W \cap I)$ (and no edges), which is impossible since this set has at most $2 \rho-2$ vertices, and, by Proposition 5 , CCP is unsolvable for so few vertices. Therefore, assuming the existence of $A$ yields a contradiction, which completes the proof.

The following is an interesting consequence of Proposition 15 to the case of CCP in cycles.

- Corollary 16. If $\mathcal{A}$ is a wait-free algorithm solving the coordination with constraints and preferences task in $C_{4 k-3}$ for $k$ processes then $\mathcal{A}$ cannot be static.

A natural guess is that CCP can be solved in $G$ for a number of processes that grows with the size of the smallest maximal independent set in $G$. Having this in mind, Theorem 4 may appear to be rather weak. Indeed, for instance, in complete bipartite graphs $K_{x, y}=$ $(X \cup Y, X \times Y)$, with $x=|X|$, and $y=|Y|$, there are no 2 -admissible independent sets. Thus our static algorithm, although optimal among static algorithms, cannot do better than solving CCP in $K_{x, y}$ for just one process! The truth is that our algorithm is not to be blamed. Indeed, the intuition that the number of processes that can be accommodated should grow with the size of the smallest maximal independent set is wrong, and any algorithm, not just static ones, cannot do better that solving CCP in $K_{x, y}$ for a single process only, as shown below (which establishes Theorem 3).

- Proposition 17. Let $x, y$ be positive integers. Coordination with constraints and preferences in $K_{x, y}$ is unsolvable for more than one process.

Proof. Recall that, in the s-set agreement task, each process proposes a value, and each correct process decides a proposed value so that the number of distinct decisions is at most $s$. We show that if algorithm $\mathcal{A}$ solves CCP in $K_{x, y}$ for $k$ processes (hence either $x$ or $y$ is at least $k$ ), then there is an algorithm $\mathcal{B}$ that solves $\left\lceil\frac{k}{2}\right\rceil$-set agreement for $k$ processes - yet, it is known that $s$-set agreement on $k$ processes is unsolvable for any $1 \leq s \leq k-1$ and $k \geq 2$ (see $[8,18,20]$ ), a contradiction. In $K_{x, y}=(X \cup Y, X \times Y), X$ and $Y$ are the unique two maximal independent sets. We say that processes $p_{1}, \ldots, p_{\left\lceil\frac{k}{2}\right\rceil}$ belong to $X$, and the remaining processes, $p_{\left\lceil\frac{k}{2}\right\rceil+1}, \ldots, p_{k}$, belong to $Y$, in the following sense: to the processes $p_{1}, \ldots, p_{\left\lceil\frac{k}{2}\right\rceil}$, we assign pairwise distinct vertices $v_{i} \in X$ as preferences, $i=1, \ldots,\left\lceil\frac{k}{2}\right\rceil$, and to the processes $p_{\left\lceil\frac{k}{2}\right\rceil+1}, \ldots, p_{k}$, we assign pairwise distinct vertices $v_{i} \in Y$ as preferences, $i=\left\lceil\frac{k}{2}\right\rceil+1, \ldots, k$. We use algorithm $\mathcal{A}$ solving CCP on this instance for construction an algorithm $\mathcal{B}$ solving $\left\lceil\frac{k}{2}\right\rceil$-set agreement.

A process $p_{i}$ that belongs to $X$ first announces its preference (considered as its proposal), and then invokes $\mathcal{A}$ using $v_{i} \in X$ as its input. If $p_{i}$ gets assigned to a vertex of $X$ in $\mathcal{A}$, then it decides $v_{i}$ in $\mathcal{B}$, otherwise it decides any vertex $v_{j} \in Y$ with $j \in\left\{\left\lceil\frac{k}{2}\right\rceil+1, \ldots, k\right\}$. A process $p_{i}$ that belongs to $Y$ proceeds similarly. That is, if $p_{i}$ gets assigned to a vertex of $Y$ in $\mathcal{A}$, then it decides its preference $v_{i}$ in $\mathcal{B}$, otherwise it decides any vertex $v_{j} \in X$ with $j \in\left\{1, \ldots,\left\lceil\frac{k}{2}\right\rceil\right\}$. Since $\mathcal{A}$ is correct, and since there are no independent sets in $K_{x, y}$ that include vertices of both $X$ and $Y$ simultaneously, it follows that if a process gets assigned to a vertex of $X$ (resp., $Y$ ) in $\mathcal{A}$, then every other process gets assigned to a vertex of $X$ (resp., $Y$ ) as well. Therefore, in every execution of $\mathcal{B}$, processes either decide the preferences of the processes that belongs to $X$, or the preferences of the processes that belongs to $Y$, hence at most $\left\lceil\frac{k}{2}\right\rceil$ proposals are decided in $\mathcal{B}$.

## 5 Conclusion

We have considered a generalization of renaming in graphs, in which deciding a node forbids others to use neighboring nodes. We proved a lower bound for Hamiltonian graphs, and provided optimal algorithms for 2 processes on a pentagon, and 3 processes on a nonagon. For the case where processes agree beforehand on a given maximal independent set, we designed optimal static algorithms for solving this problem. Static algorithms are however sub-optimal, as illustrated in the case of the rings. The design of optimal dynamic algorithms
for solving the coordination with preferences and constraints tasks in graphs remains an open problem, even for rings.

## - References

1 Yehuda Afek, Hagit Attiya, Danny Dolev, Eli Gafni, Michael Merritt, and Nir Shavit. Atomic snapshots of shared memory. J. ACM, 40(4):873-890, September 1993.
2 Yehuda Afek, Yakov Babichenko, Uriel Feige, Eli Gafni, Nati Linial, and Benny Sudakov. Oblivious collaboration. In David Peleg, editor, Distributed Computing, volume 6950 of Lecture Notes in Computer Science, pages 489-504. Springer Berlin Heidelberg, 2011.
3 Yehuda Afek, Yakov Babichenko, Uriel Feige, Eli Gafni, Nati Linial, and Benny Sudakov. Musical chairs. SIAM Journal on Discrete Mathematics, 28(3):1578-1600, 2014.
4 H. Attiya, A. Bar-Noy, D. Dolev, D. Peleg, and R. Reischuk. Renaming in an asynchronous environment. Journal of the ACM, 37(3):524-548, 1990.
5 Hagit Attiya and Jennifer Welch. Distributed Computing Fundamentals, Simulations, and Advanced Topics, Second Edition. John Wiley and Sons, Inc., 2004.
6 Brenda S. Baker and Edward G. Coffman Jr. Mutual exclusion scheduling. Theor. Comput. Sci., 162(2):225-243, 1996.
7 Hans L. Bodlaender and Klaus Jansen. Restrictions of graph partition problems. part I. Theor. Comput. Sci., 148(1):93-109, 1995.
8 Elizabeth Borowsky and Eli Gafni. Generalized FLP impossibility result for $t$-resilient asynchronous computations. In STOC'93: Proceedings of the twenty-fifth annual ACM symposium on Theory of computing, pages 91-100, New York, NY, USA, 1993. ACM.
9 Elizabeth Borowsky and Eli Gafni. Immediate atomic snapshots and fast renaming. In Proceedings of the Twelfth Annual ACM Symposium on Principles of Distributed Computing, PODC '93, pages 41-51, New York, NY, USA, 1993. ACM.
10 Armando Castañeda, Sergio Rajsbaum, and Michel Raynal. The renaming problem in shared memory systems: An introduction. Comput. Sci. Rev., 5(3):229-251, August 2011.
11 Guy Even, Magnús M. Halldórsson, Lotem Kaplan, and Dana Ron. Scheduling with conflicts: online and offline algorithms. J. Scheduling, 12(2):199-224, 2009.
12 Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro. Distributed Computing by Oblivious Mobile Robots. Synthesis Lectures on Distributed Computing Theory. Morgan \& Claypool Publishers, 2012.
13 Eli Gafni, Achour Mostéfaoui, Michel Raynal, and Corentin Travers. From adaptive renaming to set agreement. Theoretical Computer Science, 410(14):1328-1335, 2009. Structural Information and Communication Complexity (SIROCCO 2007).
14 Eli Gafni and Sergio Rajsbaum. Recursion in distributed computing. In Shlomi Dolev, Jorge Cobb, Michael Fischer, and Moti Yung, editors, Stabilization, Safety, and Security of Distributed Systems, volume 6366 of Lecture Notes in Computer Science, pages 362-376. Springer Berlin Heidelberg, 2010.
15 M. R. Garey and Ronald L. Graham. Bounds for multiprocessor scheduling with resource constraints. SIAM J. Comput., 4(2):187-200, 1975.
16 Magnús M. Halldórsson, Guy Kortsarz, Andrzej Proskurowski, Ravit Salman, Hadas Shachnai, and Jan Arne Telle. Multicoloring trees. Inf. Comput., 180(2):113-129, 2003.
17 Maurice Herlihy, Dmitry Kozlov, and Sergio Rajsbaum. Distributed Computing Through Combinatorial Topology. Morgan Kaufmann, 2013.
18 Maurice Herlihy and Nir Shavit. The topological structure of asynchronous computability. J. ACM, 46(6):858-923, 1999.

19 Michel Raynal. Concurrent Programming: Algorithms, Principles, and Foundations. Springer, 2013.


Figure $3 C C P$ algorithm for three processes in $C_{9}$ (code for resolving conflicts with 3 processes).

20 Michael Saks and Fotios Zaharoglou. Wait-Free k-Set Agreement is Impossible: The Topology of Public Knowledge. SIAM J. Comput., 29(5):1449-1483, 2000.

## A Solving CCP for three processes in $C_{9}$

We extend the algorithm for 2 processes in $C_{5}$ (Figure 1) to three processes in $C_{9}$. First, notice that the 2-process algorithm is made of at most two phases: in a first phase, processes resolve repeated vertices (leftmost picture of Figure 1), and in a second phase, once processes are proposing distinct vertices, they propose non-conflicting vertices (center picture of Figure 1). Our 3-process algorithm follows the same idea. In the algorithm, if only two processes participate, they proceed as in the algorithm for $C_{5}$. If three processes participate, a process $p_{i}$ that gets a snapshot with three processes and a conflict with its preference proceeds as follows:

1. if there is no vertex next to $p_{i}$ 's current vertex that is non-conflicting (in any direction) then move to a middle non-conflicting vertex in the segment of non-conflicting vertices (see the top row in Figure 3)
2. else move clockwise or counterclockwise, depending on the conflict pattern, with the displacement depending on processes IDs' ranking, when at least another process is proposing $p_{i}$ 's vertex (see the bottom row in Figure 3).

The correctness of the algorithm follows from the observation that if three processes participate, then eventually they propose distinct vertices and from that point on, eventually, some process will consider itself as the "middle" process and is going to move far enough from the other two.

## B Proof of Claim 5

- Claim 5. If there is a $k$-process CCP algorithm for $G$, then there exists a $k$-process CCP algorithm for $G$ in which whenever a process sees only itself in its first snapshot, it immediately decides.

Proof. The Asynchronous Computability Theorem ACT [18] states that a task is wait-free read/write solvable if and only if there is a chromatic subdivision of the input complex of the task (the simplicial complex represents all possible input configurations) with a coloring that agrees with the specification of the task.

Assume there is a $k$-process CCP algorithm for $G$, then let $S$ be such a chromatic subdivision. Herlihy and Shavit proved that any chromatic subdivision can be approximated by an $M$-th standard chromatic subdivision [18], for some big enough $M$. This $M$-th standard chromatic subdivision is crucial because it directly implies a wait-free read/write protocol that solves the task that the coloring of $S$ respects.

Now, the particular algorithms in the claim change the class of subdivisions we are dealing with. These subdivisions are very similar to the $M$-th standard chromatic subdivision with the difference that that after the 1-st standard chromatic subdivision, the corners of the subdivision are not changed, while the rest of the simplexes are subdivided in the same way. Let us call these subdivisions $M$-solo chromatic subdivisions. Considering this class of subdivisions, in which the diameter of these subdivisions shrinks as $M$ grows, we can observe that, for a big enough $M$, say $M_{S}$, the $M_{S}$-solo chromatic subdivision approximates $S$. As explained above, this $M_{S}$-solo chromatic subdivision implies an algorithm that solves CCP for $k$ processes for G in which whenever a process sees only itself in its first snapshot, it immediately decides.


[^0]:    * The 1st and 4th authors are supported by UNAM-PAPIIT. The 4th author also received support from ECOS-CONACYT and LAISLA. The 2nd author received support from the ANR project DISPLEXITY, and from the INRIA project GANG. The 5th author is supported by CPSLab project H2020-ICT-644400.

