On the Minimum Collisions Assignment Problem in Interdependent Networked Systems

Maria Diamanti*, Nikolaos Fryganiotis*, Symeon Papavassiliou*, Christos Pelekis* and Eirini Eleni Tsiropoulou†
* School of Electrical and Computer Engineering, National Technical University of Athens, Athens, Greece
*{mdiamanti@netmode, el16444@mail, pelekis@netmode, papavass@mail}.ntua.gr
† Dept. of Electrical and Computer Engineering, University of New Mexico, Albuquerque, NM, USA
†eirini@unm.edu

2022 IEEE Symposium on Computers and Communications (ISCC) | 978-1-6654-9792-3/22/$31.00 ©2022 IEEE | DOI: 10.1109/ISCC55528.2022.9912969

***Abstract*—The Minimum Collisions Assignment in an interdependent networked system is the problem of assigning a finite set of resources over the nodes of the network, such that the number of collisions, i.e., the number of interdependent nodes receiving the same resource, is minimized. It has been shown in the literature that, when the number of resources is larger than the maximum degree of the underlying graph, there exists a randomized algorithm which converges, with high probability, to an assignment of resources having zero collisions. In this work we investigate the case of a resource-constrained networked system, where the number of resources is less than or equal to the maximum degree of the underlying graph. We provide and analyze a distributed, randomized, algorithm that converges in a logarithmic number rounds to an assignment of resources over the network for which every node has at most a certain number of collisions.**

***Index Terms*—Graph coloring, Defective coloring, Games on graphs, Greedy algorithms, resource allocation.**

I. Introduction

In this article we shall be concerned with the *minimum collisions assignment* (henceforth MCA) problem in interdependent networked systems. Such a system is typically represented in the form of a finite graph, whose vertices and edges correspond to the networked nodes and their in-between dependencies, respectively, and which will be referred to as the *underlying graph* of the system. Given an interdependent networked system, the MCA problem is the problem of assigning a finite set of resources to the nodes of the network in such a way that the number of *collisions*, i.e., the number of pairs of interdependent nodes that are assigned the same resource, is minimized. This is a graph coloring problem: instead of assigning resources to the nodes of the network, one may equivalently consider the problem of assigning colors to the vertices of the underlying graph, and the MCA problem is then equivalent to the problem of assigning a given number of colors over the vertices of the graph in such a way that the number of monochromatic edges is minimum. The problem is known to be NP-hard (see [1]).

An example of a MCA problem in interdependent networked systems is the, classical, channel assignment problem in wireless networks under frequency and/or time domains. Several graph-coloring algorithms have been proposed over the years tailored to WiFi channel assignment [2], Resource Block allocation in small cell and Device-to-Device communication networks [1], [3], and time scheduling in Wireless Sensor Networks [4]. Under the aforementioned channel assignment problem, the different wireless links are regarded as vertices and the wireless channels as colors, and various graph coloring variants are examined. More recently, the problem of content placement/caching has been, also, considered in wireless networks under the prism of graph coloring. In such a setting, a specific number of files, comprising the content, needs to be assigned to the different base stations for caching in order to optimise their retrieval by potential users and reduce the transmission time through the backhaul [5].

Let us remark that the majority of the literature appears to focus on MCA problems for which the number of available colors (resources) is larger than the maximum degree of the underlying graph, and several algorithms have been proposed, both centralized and distributed, that result in colorings having zero collisions. However, it appears that instances of the MCA problem with "few", i.e., less than or equal to the maximum degree of the underlying graph, colors (e.g., channels or resource blocks) are much less investigated and in this paper we aim towards filling this gap. Respecting the need for decentralized resource management, we capitalize on each networked node's local-awareness with respect to its neighbors' allocated resources and propose a distributed approach to the MCA problem in the case of there being "few" available colors. In the following we refer to the potential entities/users of an interdependent networked system as players, and to the set of resources as colors.

II. Basic Definitions & Problem Formulation

All graphs considered in this text are finite, without loops, and undirected. Throughout the text, given a positive integer k, we denote by $[k]$ the set $\{1,\dots,k\}$ and, given a finite set F, we denote by $|F|$ its cardinality. Given a graph $G=(V,E)$ and a vertex $v\in V$, we let $\mathcal{N}_G(v)=\{u\in V:(u,v)\in E\}$ be the *neighborhood* of v in G. The *degree* of v equals $\deg_G(v)=|\mathcal{N}_G(v)|$ and the maximum degree of vertices in G is denoted Δ_G. A *k-coloring* of G is a function $\chi:V\to[k]$. Given a k-coloring χ of a graph $G=(V,E)$, and a subset $A\subset V$, we denote by $\chi(A):=\bigcup_{v\in A}\{\chi(v)\}$ the set of colors

This research work was supported by the Hellenic Foundation for Research and Innovation (H.F.R.I.) under the "1st Call for H.F.R.I. Research Projects to support Faculty members and Researchers and the procurement of high-cost research equipment grant" (Project Number: HFRI-FM17-2436).

of the vertices in A. Moreover, the *collision number* of χ is defined as $\mathcal{C}_G(\chi) = |\{e = (u, v) \in E : \chi(u) = \chi(v)\}|$, and the *collision number of a vertex* $v \in V$ under a k-coloring χ of G is defined as $\mathcal{C}_G(v; \chi) = |\{u \in \mathcal{N}_G(v) : \chi(u) = \chi(v)\}|$. In other words, $\mathcal{C}_G(\chi)$ equals the cardinality of the set consisting of all *monochromatic edges* of G under the coloring χ, and $\mathcal{C}_G(v; \chi)$ is the number of neighbors of v that receive the same color as v.

A k-coloring χ of $G = (V, E)$ is called *s-colliding* if $\mathcal{C}_G(\chi) \leq s$; it is called *$d$-defective* if $\mathcal{C}_G(v; \chi) \leq d$ holds true for all $v \in V$. A 0-colliding coloring is referred to as a *proper coloring* in the literature. As already mentioned in the introduction, the MCA problem is a graph coloring problem which is equivalent to the following.

Problem 1 (MCA Problem)**.** *Given a graph $G = (V, E)$ and a positive integer k, determine*

$$\mathcal{C}_k(G) := \min_{\chi} \mathcal{C}_G(\chi),$$

where the minimum is over all k-colorings of G.

III. Related Work

In the classical *graph coloring problem* (GCP) the objective is to find the minimum positive integer k for which a given graph G on n vertices admits a 0-colliding k-coloring. This minimum value of k is referred to as the *chromatic number* of G, and is denoted $chr(G)$. The MCA problem is, in some sense, dual to the GCP. In the setting of the MCA problem, the parameter k is fixed and the objective is to find a k-coloring of a given graph G on n vertices that has minimum collision number, among all k-colorings of G. Observe that when $k \geq chr(G)$ then the GCP implies that the MCA problem admits a 0-colliding k-coloring. In particular, it is well known that when $k \geq \Delta_G + 1$ the graph G admits a 0-colliding k-coloring, and one can obtain such a coloring using a variety of algorithms, both centralized and distributed. It is also well known that, when $k \geq \Delta_G + 1$, a 0-colliding coloring of G can be found in linear time by a centralized algorithm. However, the problem becomes more delicate when the algorithm is required to be distributed. In this work we shall be interested in distributed algorithms for the MCA problem.

Distributed algorithms for GCP are provided in [6] and [7]. Both algorithms result in a 0-colliding coloring of a graph in $O(\log(n))$ rounds, when the number of available colors is at least $\Delta + 1$, but require that each player communicates to her neighbors whether she has any conflict or not. A purely game-theoretic distributed algorithm for the GCP is provided in [8]; the algorithm does not require any cooperation/communication between the players and yields a 0-colliding coloring of a graph in $O(\log(n))$ rounds, when the number of colors available is at least $\Delta_G + 2$. The algorithm from [8] has been improved in [1] to a distributed algorithm that yields a 0-colliding coloring of a graph in at most $O(\Delta_G \cdot \log(n))$ rounds, when the number of available colors is at least $\Delta_G + 1$ colors, but requires communication among neighbors. Another improvement of the algorithm in [8], which likewise assumes no cooperation/communication among neighbors and which yields a 0-colliding coloring in $O(\log(n))$ rounds when the number of available colors is at least $\Delta_G + 1$, can be found in [9]. In other words, when $k \geq \Delta_G + 1$, it holds $\mathcal{C}_G(k) = 0$, for any graph G.

In this article we focus on instances of the MCA problem for which $k \leq \Delta_G$ which, to the best of our knowledge, appear to be less investigated. One approach to the problem is to allow the possibility of leaving some vertices uncolored, and thus employ *incomplete* 0-colliding k-colorings of the underlying graph (see [10] and references therein). Another approach is based on *dispersion games* (see [11] and [12]), but only apply to instances of the MCA problem for which the underlying graph is complete. Our approach is based on defective colorings (see [13]), and builds upon ideas from [8]. In particular, in [8] the authors define the *network coloring game*, which is played on a graph G, and study the dynamics of the game when the players adopt a particular greedy, randomized, strategy. It is shown in [8] that the dynamics of the network coloring game under the aforementioned greedy strategy converge to a Nash equilibrium that gives rise to a 0-colliding k-coloring of G, provided that $k \geq \Delta_G + 2$. In this article, following a similar approach, we introduce the *defective coloring game* and study its dynamics when a particular greedy, randomized, strategy is adopted by the players. Our main result states that the dynamics of the game under this greedy strategy converge to a Nash equilibrium that gives rise to a defective coloring of the underlying graph which, in turn, provides an upper bound on the number of collisions.

IV. Main Result: Defective Coloring Game

In this section we introduce the *defective coloring game*, which may be seen as a variant of the *network coloring game* introduced in [8]. We fix a graph $G = (V, E)$ having $n = |V|$ vertices and maximum degree Δ_G, as well as two integers k, d such that $k \in \{2, \ldots, \Delta_G\}$ and $d \in [\Delta_G - 1]$. The defective coloring game on the graph G, denoted $DCG(G; k, d)$, is defined as follows.

The players of $DCG(G; k, d)$ are the vertices of G, and participate in a game that is played over a number of rounds. In every round all players simultaneously and individually choose a color from their set of available colors, which is assumed to be the set $[k]$. Thus, after round t, the choices of the players give rise to a k-coloring of G, which is denoted χ_t. The players of $DCG(G; k, d)$ only have local information of the graph: they can only observe the colors chosen by their neighbors, and are not allowed to communicate or cooperate with one another. A player $v \in V$ is said to be *happy* after round t if her collision number under χ_t is at most d; i.e., when $\mathcal{C}_G(v; \chi_t) \leq d$. Otherwise the player is *unhappy*. The payoff to a player in $DCG(G; k, d)$ is 1 when she is happy, and 0 when she is unhappy, and a configuration of colors for which every player receives payoff 1 is a *Nash equilibrium* of the $DCG(G; k, d)$, in the sense that no player has the incentive to unilaterally change strategy under such a configuration. Observe that, when the players have chosen colors that constitute a Nash equilibrium, the corresponding k-coloring of the graph is d-defective.

In this work we define a *symmetric strategy* for the players in $DCG(G;k,d)$ (i.e., a strategy that is the same for all players) and show that it achieves convergence to Nash equilibrium after a finite number of rounds. In order to formally define this strategy we need to introduce some notation. Let $\chi_t(v)$ denote the color chosen by player $v \in V$ after round t, and let $\chi_t(\mathcal{N}(v))$ be the set of colors chosen by the neighbors of v after round t.

Greedy strategy. Suppose that $k = \Delta_G - s$ and $d = s+2$, for some fixed $s \in \{0,1,\ldots,\Delta_G - 3\}$. Assume further that each player in $DCG(G;k,d)$ adopts the following strategy: if a player, say v, is happy after a certain round, say t, then she sticks to her choice in all subsequent rounds, i.e., $\chi_s(v) = \chi_t(v)$, for all $s > t$. If she is unhappy then in the next round she *changes* color, and chooses uniformly at random a color from the set $[k] \setminus \chi_t(\mathcal{N}(v))$.

In other words, under the Greedy strategy, a player who is unhappy after a certain round, say t, chooses in the next round a color uniformly at random from the set consisting of those colors that are *not* chosen by her neighbors after round t. The corresponding algorithm is summarized in Algorithm 1.

Remark 1. *Notice that, when the players of $DCG(G;k,d)$, with $k = \Delta_G - s$ and $d = s+2$, adopt the Greedy strategy, a player who is happy after a certain round remains happy in all subsequent rounds. Furthermore, if player v is unhappy after round t, then it holds $|[k] \setminus \chi_t(\mathcal{N}(v))| \geq 2$. In particular, an unhappy player has always at least two available colors to choose from in the next round.*

Now suppose that all players in $DCG(G;k,d)$ adopt the Greedy strategy. For each $v \in V$, let τ_v be the first round after which player v is happy. Notice that τ_v is a random variable and that $\tau := \max_{v \in V} \tau_v$ is the first round after which the game reaches a Nash equilibrium. Our main result states that the expected value of τ is of logarithmic order, and reads as follows. Recall that, given two random variables X, Y assuming non-negative values, the random variable X is said to be stochastically smaller than the random variable Y, denoted $X \leq_{st} Y$, when it holds $\mathbb{P}(X > t) \leq \mathbb{P}(Y > t)$, for all t. In particular, $X \leq_{st} Y$ implies that $\mathbb{E}(X) \leq \mathbb{E}(Y)$.

Theorem 1. *Let G be a graph on n vertices and maximum degree $\Delta_G \geq 3$. Let k, d be fixed positive integers such that $k = \Delta_G - s$ and $d = s+2$, for some $s \in \{0,1,\ldots,\Delta_G - 3\}$. Suppose that each player in $DCG(G;k,d)$ adopts the Greedy strategy. Let τ be the first round after which all players are happy. Then, for any starting assignment of colors to the vertices, it holds that τ is stochastically smaller than a random variable T which satisfies*

$$\mathbb{E}(T) \leq \frac{4}{\mu}(1 + \log(n)) \quad \text{and} \quad Var(T) \leq \frac{16 \cdot n}{\mu^2},$$

where $\mu = -\log\left(1 - \frac{1-1/4^{d-1}}{1-1/5^{d-1}}\right)$.

Algorithm 1 Greedy strategy algorithm.

Initialize the graph $G = (V,E)$ of the interdependent networked system, having maximum degree Δ_G, and a positive integer $s \in \{0,1,\ldots,\Delta_G - 3\}$.
$d \leftarrow s+2$
$k \leftarrow \Delta_G - s$
for each $v \in V$ **do**
  Choose c_0 randomly from the set $\mathcal{A}_0(v) := \{1,\ldots,k\}$.
  $t \leftarrow 0$
  while $\mathcal{C}_G(v;\chi_t) \geq d+1$ **do**
    $t \leftarrow t+1$
    $\mathcal{A}_{t+1}(v) \leftarrow [k] \setminus \chi_t(\mathcal{N}(v))$
    Choose a color c_{t+1} uniformly at random from the set $\mathcal{A}_{t+1}(v)$.
  end while
  Return c_{t+1}.
end for

In other words, when s does not depend on n and the players in the defective coloring game adopt the Greedy strategy, the game reaches a d-defective k-coloring of the graph in $O(\log(n))$ expected number of rounds. Now the number of monochromatic edges in such a coloring provides an upper bound on the quantity $\mathcal{C}_G(k)$, in Problem 1. In particular, Theorem 1 yields the following.

Corollary 1. *Let G be a graph on $n = |V|$ vertices and maximum degree $\Delta_G \geq 3$. Suppose that $k = \Delta_G - s$, for some $s \in \{0,1,\ldots,\Delta_G - 3\}$. Then it holds $\mathcal{C}_k(G) \leq \frac{n(s+2)}{2}$.*

Proof. Consider the defective-coloring game $DCG(G;k,d)$, where $d = s+2$. From Theorem 1 we know that, when the players in $DCG(G;k,d)$ adopt the Greedy strategy, the game reaches a Nash equilibrium in $O(\log(n))$ expected number of rounds. Let χ be the k-coloring corresponding to a Nash equilibrium. In such an equilibrium point, all vertices are happy and the collision number of each vertex $v \in V$ is at most d. Let G_χ be the graph induced by the monochromatic edges of G under χ. Then $\mathcal{C}_G(\chi)$ equals the number of edges in G_χ, and the result follows from the degree-sum formula in G_χ. □

V. Proof of Theorem 1

In this section we prove Theorem 1. We fix a starting assignment of colors to the vertices, and we assume that each player in $DCG(G;k,d)$ adopts the Greedy strategy. Recall that we assume that $k = \Delta_G - s$ and $d = s+2$, for some $s \in \{0,1,\ldots,\Delta_G - 3\}$. We begin with a result that provides a lower bound on the probability that a player, who is unhappy after round a certain round, receives "enough" available colors in the next round. This will require some additional piece of notation.

Recall that $\chi_t(v)$ denotes the color chosen by player v after round t, and that $\chi_t(\mathcal{N}(v))$ is the set of colors chosen by its neighbors after round t. For each $t \geq 1$, let $\mathcal{H}_t$ be the set of happy players after round t, and $\mathcal{U}_t = V \setminus \mathcal{H}_t$ the set of unhappy players after round t; hence $v \in \mathcal{U}_t$ means that $\mathcal{C}_G(v;\chi_t) \geq d+1$. Given $v \in \mathcal{U}_t$, let $\mathcal{A}_t(v) := [k] \setminus \chi_t(\mathcal{N}(v))$ be the set of colors *available* to v after round t; hence, under the Greedy strategy, player v chooses in the next round a color uniformly at random from the set $\mathcal{A}_t(v)$. Let $p_t(v) = \frac{1}{|\mathcal{A}_t(v)|}$ be the probability with which the unhappy player v chooses a color in the next round. For $v \in \mathcal{H}_t$, set $\mathcal{A}_t(v) = \{\chi_t(v)\}$ and $p_t(v) = 1$. Similarly, given a vertex $v \in V$, let $\mathcal{H}_t(v) := \mathcal{H}_t \cap \mathcal{N}(v)$ denote the set of happy neighbors of v after round t, and let $\mathcal{F}_t(v) := \chi_t(\mathcal{H}_t(v))$ be the set of colors chosen by the happy neighbors of v after round t. Let also $\mathcal{U}_t(v) = \mathcal{N}(v) \setminus \mathcal{H}_t(v)$ denote the set consisting of the unhappy neighbors of v after round t. Observe that every color from the set $[k] \setminus \mathcal{F}_t(v)$ has a positive chance of not being chosen by the unhappy heighbors, and therefore has a positive chance of belonging to the set $\mathcal{A}_{t+1}(v)$. Finally, let $f_t(v) = |\mathcal{F}_t(v)|$, and observe that, since happy players do not change their color, the sequence $\{f_t(v)\}_{t\geq 1}$ is non-decreasing. In particular, this implies that the number of colors available to v after round $t+1$ as well as after round $t+2$ is less than or equal to $k - f_t(v)$. The next lemma establishes a lower estimate on the probability that the number of colors available to player $v \in \mathcal{U}_t$ after round $t+1$ is at least $\frac{k-f_t(v)}{5^{d-1}}$.

Lemma 1. *For each $t \geq 1$ and each $v \in \mathcal{U}_t$, it holds*

$$\mathbb{P}\left(|\mathcal{A}_{t+1}(v)| \geq \frac{k - f_t(v)}{5^{d-1}}\right) \geq 1 - \frac{1 - 1/4^{d-1}}{1 - 1/5^{d-1}}\,.$$

Proof. Let $v \in \mathcal{U}_t$ be fixed and, for simplicity in notation, let us set $f := f_t(v)$. Since v is unhappy, there are at least $d+1$ vertices $u \in \mathcal{N}(v)$ such that $\chi_t(u) = \chi_t(v)$. We now proceed with estimating $\mathbb{E}(|\mathcal{A}_{t+1}(v)|)$ from below; the proof is then completed by applying Markov's inequality. As mentioned already, any color from the set $[k] \setminus \mathcal{F}_t(v)$ has a positive chance of belonging to $\mathcal{A}_{t+1}(v)$. In particular, color $i \in [k] \setminus \mathcal{F}_t(v)$ belongs to $\mathcal{A}_{t+1}(v)$ if it is not chosen by any neighbor $u \in \mathcal{U}_t(v)$ for which $i \in \mathcal{A}_t(u)$; this happens with probability $\prod_{\{u\in\mathcal{U}_t(v): i \in \mathcal{A}_t(u)\}}(1 - p_t(u))$. Therefore, denoting $E := \mathbb{E}(|\mathcal{A}_{t+1}(v)|)$, the arithmetic-geometric means inequality implies that

$$\begin{aligned}
E &= \sum_{i\in[k]\setminus\mathcal{F}_t(v)} \prod_{\{u\in\mathcal{U}_t(v): i\in\mathcal{A}_t(u)\}} (1 - p_t(u)) \\
&\geq (k-f) \cdot \left(\prod_{i\in[k]\setminus\mathcal{F}_t(v)} \prod_{\{u\in\mathcal{U}_t(v): i\in\mathcal{A}_t(u)\}} (1 - p_t(u)) \right)^{\frac{1}{k-f}} \\
&\geq (k-f) \cdot \left(\prod_{u\in\mathcal{U}_t(v)} \prod_{i\in\mathcal{A}_t(u)} (1 - p_t(u)) \right)^{\frac{1}{k-f}} \\
&= (k-f) \cdot \left(\prod_{u\in\mathcal{U}_t(v)} \left(1 - \frac{1}{|\mathcal{A}_t(u)|}\right)^{|\mathcal{A}_t(u)|} \right)^{\frac{1}{k-f}} .
\end{aligned}$$

Recall that $|\mathcal{A}_t(u)| \geq 2$, for every $u \in \mathcal{U}_t(v)$. Since the sequence $\{(1 - \frac{1}{m})^m\}_{m\geq 2}$ is non-decreasing, it holds $\left(1 - \frac{1}{|\mathcal{A}_t(u)|}\right)^{|\mathcal{A}_t(u)|} \geq \left(1 - \frac{1}{2}\right)^2 = \frac{1}{4}$. Summarizing the above, we have shown that $\mathbb{E}(|\mathcal{A}_{t+1}(v)|) \geq (k-f) \cdot \left(\frac{1}{4}\right)^{\frac{|\mathcal{U}_t(v)|}{k-f}}$.

Now, since $k = \Delta_G - d + 2$, it holds $|\mathcal{U}_t(v)| \leq \Delta_G - f = k + d - 2 - f$, and hence we have $\frac{|\mathcal{U}_t(v)|}{k-f} \leq d - 1$. This implies that $\left(\frac{1}{4}\right)^{\frac{|\mathcal{U}_t(v)|}{k-f}} \geq \left(\frac{1}{4}\right)^{d-1}$ and therefore it holds $\mathbb{E}(|\mathcal{A}_{t+1}(v)|) \geq \frac{k-f}{4^{d-1}}$. To finish the proof, let $X = k - f - |\mathcal{A}_{t+1}(v)|$ and apply the previous lower bound on $\mathbb{E}(|\mathcal{A}_{t+1}(v)|)$ together with Markov's inequality to deduce $\mathbb{P}\left(|\mathcal{A}_{t+1}(v)| < \frac{k-f}{5^{d-1}}\right) = \mathbb{P}\left(X > (k-f)\cdot(1 - \frac{1}{5^{d-1}})\right) < \frac{\mathbb{E}(X)}{(k-f)\cdot(1-\frac{1}{5^{d-1}})} \leq \frac{1-\frac{1}{4^{d-1}}}{1-\frac{1}{5^{d-1}}}$, as desired. □

The next lemma concerns a lower estimate on the probability that a player, who is unhappy after round t, becomes happy after two rounds.

Lemma 2. *It holds*

$$\mathbb{P}(v \in \mathcal{H}_{t+2} \,|\, v \in \mathcal{U}_t) \geq \left(\frac{1}{4}\right)^{5^{d-1}\cdot(d-1)} \cdot \left(1 - \frac{1 - 1/4^{d-1}}{1 - 1/5^{d-1}}\right)\,.$$

Proof. Let $v \in \mathcal{U}_t$. Then, conditional on $\mathcal{A}_{t+1}(v)$ and $v \in \mathcal{U}_{t+1}$, the probability that player v is happy after round $t+2$ is the average of the probabilities that a fixed color from $\mathcal{A}_{t+1}(v)$ is chosen by at most d unhappy neighbors of v. Now the probability that a fixed color $i \in \mathcal{A}_{t+1}(v)$ is chosen by at most d players $u \in \mathcal{U}_{t+1}(v)$ is greater than or equal to the probability that color i is not chosen by any player from $\mathcal{U}_{t+1}(v)$; the latter probability being equal to $\prod_{\{u\in\mathcal{U}_{t+1}(v): i\in\mathcal{A}_{t+1}(u)\}}(1 - p_{t+1}(u))$. Therefore, conditional on $\mathcal{A}_{t+1}(v)$ and $v \in \mathcal{U}_{t+1}$, the probability that player v is happy after round $t+2$ is at least

$$\begin{aligned}
P &:= \frac{1}{|\mathcal{A}_{t+1}(v)|} \sum_{i\in\mathcal{A}_{t+1}(v)} \prod_{\{u\in\mathcal{U}_{t+1}(v): i\in\mathcal{A}_{t+1}(u)\}} (1 - p_{t+1}(u)) \\
&\geq \left(\prod_{i\in\mathcal{A}_{t+1}(v)} \prod_{\{u\in\mathcal{U}_{t+1}(v): i\in\mathcal{A}_{t+1}(u)\}} (1 - p_{t+1}(u)) \right)^{\frac{1}{|\mathcal{A}_{t+1}(v)|}} \\
&\geq \left(\prod_{u\in\mathcal{U}_{t+1}(v)} \prod_{i\in\mathcal{A}_{t+1}(u)} (1 - p_{t+1}(u)) \right)^{\frac{1}{|\mathcal{A}_{t+1}(v)|}} \\
&= \left(\prod_{u\in\mathcal{U}_{t+1}(v)} \left(1 - \frac{1}{|\mathcal{A}_{t+1}(u)|}\right)^{|\mathcal{A}_{t+1}(u)|} \right)^{\frac{1}{|\mathcal{A}_{t+1}(v)|}} ,
\end{aligned}$$

where the first estimate follows from the arithmetic-geometric means inequality. As in the proof of Lemma 1, it holds $\left(1 - \frac{1}{|\mathcal{A}_{t+1}(u)|}\right)^{|\mathcal{A}_{t+1}(u)|} \geq \frac{1}{4}$, and therefore we conclude that $P \geq \left(\frac{1}{4}\right)^{\frac{|\mathcal{U}_{t+1}(v)|}{|\mathcal{A}_{t+1}(v)|}}$. Now observe that $|\mathcal{U}_{t+1}(v)| \leq \Delta_G - |\mathcal{H}_{t+1}(v)| = k + d - 2 - |\mathcal{H}_{t+1}(v)| \leq k + d - 2 - f_t(v)$, which implies that, conditional on the event that $|\mathcal{A}_{t+1}(v)| \geq \frac{k-f_t(v)}{5^{d-1}}$, it holds $\frac{|\mathcal{U}_{t+1}(v)|}{|\mathcal{A}_{t+1}(v)|} \leq 5^{d-1} \cdot \frac{k+d-2-f_t(v)}{k-f_t(v)} \leq 5^{d-1} \cdot (d-1)$. Hence, conditional on the event that $|\mathcal{A}_{t+1}(v)| \geq \frac{k-f_t(v)}{5^{d-1}}$, we have $P \geq \left(\frac{1}{4}\right)^{5^{d-1}\cdot(d-1)}$, and the result follows from Lemma 1. $\square$

We now proceed with the proof of Theorem 1. Given $v \in V$, let τ_v be the first round after which player v is happy and set $\tau = \max_v \tau_v$ be the first round after which all players are happy. We want to establish an upper estimate on $\mathbb{E}(\tau)$. Observe that the random variables $\tau_v, v \in V$, are *not* mutually independent and therefore our bound on τ will be a worst-case estimate. To this end, we follow the approach from [14], and employ ideas from the theory of maximally dependent random variables. Given a real number $\mu > 0$, let X_μ denote an exponential random variable of parameter μ, and let $c_d := \left(\frac{1}{4}\right)^{5^{d-1}\cdot(d-1)} \cdot \left(1 - \frac{1-1/4^{d-1}}{1-1/5^{d-1}}\right)$, for $d \geq 2$.

Lemma 3. *For every $v \in V$, it holds $\tau_v \leq_{st} 4 \cdot X_\mu$, where $\mu = -\log(1 - c_d)$.*

Proof. We have to show that $\mathbb{P}(\tau_v > t) \leq \mathbb{P}(X_\mu > \frac{t}{4})$, for all t. From Lemma 2 we know that $\mathbb{P}(\tau_v > t+2 \,|\, \tau_v > t) = \mathbb{P}(v \in U_{t+2} \,|\, v \in U_t) \leq 1 - c_d$ holds true for every $t \geq 1$. Now note that when t is odd, say $t = 2m+1$, it holds

$$\mathbb{P}(\tau_v > t) \leq \prod_{i=1}^{m} \mathbb{P}(\tau_v > 2i+1 \,|\, \tau_v > 2i-1)$$
$$\leq (1-c_d)^m \leq (1-c_d)^{t/4} = \mathbb{P}\left(X_\mu > \frac{t}{4}\right),$$

as desired. If t is even, the proof is similar and is left to the reader. The result follows. $\square$

The proof of Theorem 1 is almost complete. Given two random variables X, Y, we denote the fact that they have the same distribution by $X \sim Y$.

Proof of Theorem 1. From Lemma 3 we know that for every $v \in V$ it holds $\tau_v \leq_{st} Z_v$, where $Z_v \sim 4 \cdot X_\mu$. Now the fact that $\tau_v \leq_{st} Z_v$ implies (see [15, Theorem 1.A.1]) that there exist random variables $\hat{\tau}_v, \hat{Z}_v$ such that $\hat{\tau}_v \sim \tau_v$, $\hat{Z}_v \sim Z_v$, and $\hat{\tau}_v \leq \hat{Z}_v$ with probability 1; hence $\max_v \hat{\tau}_v \leq \max_v \hat{Z}_v$ with probability 1. Since $\tau \sim \max_v \hat{\tau}_v$, it follows that $\tau \leq_{st} 4 \cdot M$, where M is the maximum of n exponential random variables, say $\{X_\mu^v\}_{v\in V}$, of parameter μ, and thus $\mathbb{E}(\tau) \leq 4 \cdot \mathbb{E}(M)$. It is therefore enough to upper bound $\mathbb{E}(M)$. To this end, we borrow ideas from [16]. Observe that for every real number a it holds $M \leq a + \sum_v \max\{X_\mu^v - a, 0\}$; hence it holds

$$\mathbb{E}(M) \leq a + \sum_v \mathbb{E}(\max\{X_\mu^v - a, 0\})$$
$$= a + n\int_a^\infty (1 - F(x))\, dx,$$

where $F(\cdot)$ is the distribution function of X_μ^v. Now consider the function $h(a) = a + n\int_a^\infty (1-F(x))\, dx$, defined for real a, and notice that $h(\cdot)$ attains its minimum at $a_n := F^{-1}(1-\frac{1}{n})$. Since $F(x) = 1 - e^{\mu x}$, we deduce that

$$\mathbb{E}(M) \leq a_n + n\int_{a_n}^\infty e^{-\mu x} dx = \frac{1}{\mu}(1 + \log(n)),$$

as desired. Finally, the main result from [17] implies that $\mathrm{Var}(M) \leq n \cdot \mathrm{Var}(X_\mu) = \frac{n}{\mu^2}$. The result follows upon letting $T = 4 \cdot M$. $\square$

VI. Numerical Evaluation

To evaluate the operation and performance of the proposed strategy in terms of the resulting number of collisions at the Nash equilibrium and the required number of rounds until the convergence, as a function of the parameter s, different scenarios have been randomly generated and considered, using the Erdős–Rényi random graph model. Specifically, in the following, different numbers of vertices n are examined, while each edge is included in the corresponding graphs with a probability $p = 0.5$. The evaluation results have been averaged over 300 different random graph realizations.

Fig. 1a demonstrates the convergence of the greedy strategy in randomly generated graphs with $n = 20$ vertices and maximum degree $\Delta_G = 13$ for different values of $s \in \{0, 1, \ldots, 10\}$. At this point, it should be noted that the maximum degree of all generated graphs considered in this simulation setup was predetermined and set equal to $\Delta_G = 13$, which is an average maximum degree of graphs with $n = 20$ and $p = 0.5$, in order to yield the same range of values of s and enable the proper averaging of the obtained results. On the left vertical axis of Fig. 1a, the total number of collisions at the Nash equilibrium is depicted, while the number of rounds until the game reaches a Nash equilibrium is depicted on the right vertical axis. Based on the proposed strategy, as the value of s increases, the number of available colors k - calculated as $k = \Delta_G - s$ - decreases, whereas the vertices' threshold of acceptable collisions d between them and their neighbors - defined as $d = s + 2$ - increases. As a result of the former, an increase in the value of s yields a higher total number of collisions. On the contrary, an increase in the value of s results in a lower number of rounds, as the vertices' threshold of acceptable collisions d is looser.

Considering the same simulation setup with Fig. 1a, in Fig. 1b, we investigate the statistics of the resulting number of collisions per vertex under the different values of s. Given a value of s, e.g., $s = 0$, the illustrated boxplot depicts the distribution of the number of collisions per vertex that results from the averaging of the randomly generated graphs. The red horizontal line inside the box represents the mean number of collisions per vertex, the lines at the borders of the box characterize the lower and upper quartiles, while the lines outside the box denote the min and max number of collisions per vertex of the graph. Therefore, it appears that given a fixed s, the distribution of the number of collisions per vertex is quite uniform among the different vertices. On the other hand, considering different values of s, Fig. 1b reveals the increase in mean number of collisions per vertex as s increases, validating the outcome of Fig. 1a with respect to the total number of collisions of the whole graph. At this point, it is remarkable that our experiments suggest that, under the greedy strategy, the number of collisions per vertex at a Nash equilibrium is

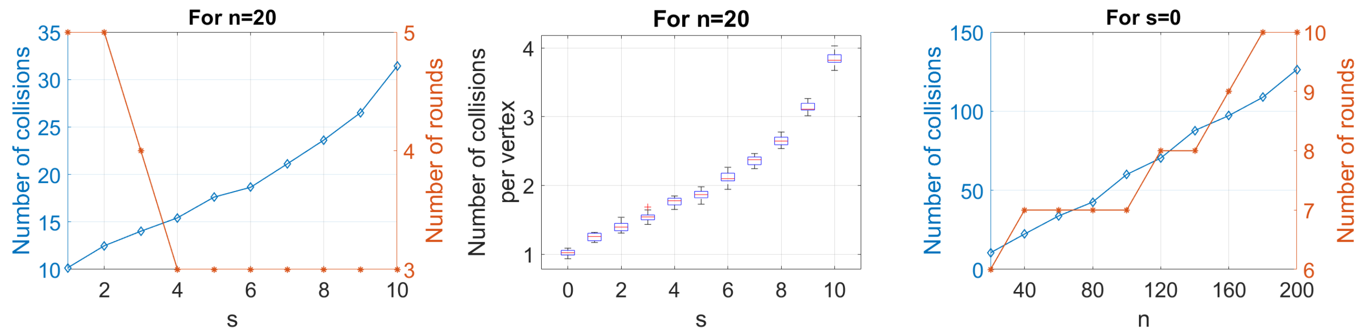

(a) (b) (c)

Fig. 1: Greedy strategy evaluation under randomly generated Erdős–Rényi graphs with probability 0.5.

much less than the bound provided in Corollary 1, which is a worst-case estimate.

In Fig. 1c, we perform a scalability analysis for an increasing number of vertices, while considering $s = 0$, which is the case that requires the maximum number of rounds until the convergence of the greedy strategy. Apparently, both the number of collisions (left vertical axis) and the number of rounds (right vertical axis) increase as the number of vertices increases. Nevertheless, even when considering $n = 200$, the required number of rounds until the Nash equilibrium is reached does not exceed the number of $t = 10$ rounds, confirming in this way numerically Theorem 1, and the ultimate purpose of this work.

VII. Conclusion and Future Work

This paper investigates the problem of assigning a finite set of resources, over the nodes of an interdependent network, so that the number of interdependent nodes receiving the same resource is minimized. We propose a game-theoretic modeling which results in a distributed randomized algorithm that converges, in a logarithmic number of rounds, to a Nash equilibrium of the game which, in turn, gives rise to a defective coloring of the underlying graph. Our findings are additionally supported with numerical simulations and evaluations. Several questions remain to be further investigated. Indicatively we highlight that the investigation and exploration of the tightness of the crude bound, provided by Corollary 1, on the number of collisions at a Nash equilibrium, is of high research and practical importance.

The present work lays the ground, from a theoretical perspective, for the application of our randomized algorithm to actual interdependent networked systems. A particularly interesting problem that has recently attracted the attention of the research community and that falls in the range of defective coloring games is the mitigation of the radar and communication frequency overlapping in integrated sensing and communication wireless networks, and we expect that we will report on this matter in the near future.

References

[1] M. Goonewardena, H. Akbari, W. Ajib, and H. Elbiaze, "On minimum-collisions assignment in heterogeneous self-organizing networks," in *2014 IEEE Global Communications Conference*, 2014, pp. 4665–4670.

[2] D. Orden, J. M. Gimenez-Guzman, I. Marsa-Maestre, and E. De la Hoz, "Spectrum graph coloring and applications to wi-fi channel assignment," *Symmetry*, vol. 10, no. 3, p. 65, 2018.

[3] T. Yang, R. Zhang, X. Cheng, and L. Yang, "Graph coloring based resource sharing (gcrs) scheme for d2d communications underlaying full-duplex cellular networks," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 8, pp. 7506–7517, 2017.

[4] S. C. Ergen and P. Varaiya, "Tdma scheduling algorithms for wireless sensor networks," *Wireless Networks*, vol. 16, pp. 985–997, 2010.

[5] M. Javedankherad, Z. Zeinalpour-Yazdi, and F. Ashtiani, "Content placement in cache networks using graph coloring," *IEEE Systems Journal*, vol. 14, no. 3, pp. 3129–3138, 2020.

[6] M. Luby, "Removing randomness in parallel computation without a processor penalty," in *[Proceedings 1988] 29th Annual Symposium on Foundations of Computer Science*, 1988, pp. 162–173.

[7] I. Finocchi, A. Panconesi, and R. Silvestri, "Experimental analysis of simple, distributed vertex coloring algorithms," in *13th Annual ACM-SIAM Symposium on Discrete Algorithms*. USA: Society for Industrial and Applied Mathematics, 2002, pp. 606–615.

[8] K. Chaudhuri, F. Chung Graham, and M. S. Jamall, "A network coloring game," in *International Workshop on Internet and Network Economics*, vol. 5385. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 522–530.

[9] N. Fryganiotis, S. Papavassiliou, and C. Pelekis, "A note on the network coloring game," 2021. [Online]. Available: arXiv:2106.00402

[10] K.-J. Wu, Y.-W. P. Hong, and J.-P. Sheu, "Coloring-based channel allocation for multiple coexisting wireless body area networks: A game-theoretic approach," *IEEE Transactions on Mobile Computing*, vol. 21, no. 1, pp. 63–75, 2022.

[11] S. Alpern and D. J. Reyniers, "Spatial dispersion as a dynamic coordination problem," *Theory and Desicion*, vol. 53, pp. 29–59, 2002.

[12] T. Grenager, R. Powers, and Y. Shoham, "Dispersion games: General definitions and some specific learning results," in *18th National Conference on Artificial Intelligence*. USA: American Association for Artificial Intelligence, 2002, pp. 398–403.

[13] C. Archetti, N. Bianchessi, A. Hertz, A. Colombet, and F. Gagnon, "Directed weighted improper coloring for cellular channel allocation," *Discrete Applied Mathematics*, vol. 182, pp. 46–60, 2015.

[14] C. Pelekis and M. Schauer, "Network coloring and colored coin games," in *Search Theory*. Springer New York, 2013, pp. 59–73.

[15] M. Shaked and J. G. Shanthikumar, *Stochastic orders*. Springer New York, 2007.

[16] T. Lai and H. Robbins, "Maximally dependent random variables," in *Proceedings of the National Academy of Sciences of the United States of America*, vol. 73, no. 2. National Academy of Sciences, 1976, pp. 286–288.

[17] T. Rychlik, "Extreme variances of order statistics in dependent samples," *Statistics & probability letters*, vol. 78, no. 12, pp. 1577–1582, 2008.