Multiagent Graph Coloring: Pareto Efficiency, Fairness and Individual Rationality

Yaad Blum Jeffrey S. Rosenschein
School of Computer Science and Engineering
The Hebrew University of Jerusalem
Jerusalem, Israel
{yaadb, jeff}@cs.huji.ac.il

Abstract

We consider a multiagent extension of single-agent graph coloring. Multiple agents hold disjoint autonomous subgraphs of a global graph, and every color used by the agents in coloring the graph has associated cost. In this multiagent graph coloring scenario, we seek a minimum legal coloring of the global graph's vertices, such that the coloring is also Pareto efficient, socially fair, and individual rational. We analyze complexity of individual-rational solutions in special graph classes where classical coloring algorithms are known. Multiagent graph coloring has application to a wide variety of multiagent coordination problems, including multiagent scheduling.

Introduction

Analysis of resource allocation scenarios is an important concern in multiagent systems (MAS). Some of these interactions can be modeled using graph theory, which studies the properties of connections between objects and offers a generic theoretical framework for analyzing agent relationships. In this paper, we investigate graph theoretic issues from an economics perspective, considering multiple rational agents autonomously coloring the vertices of a graph.

We choose to focus on the computationally hard problem of graph coloring primarily due to the inherent trade-offs that lie at the heart of the problem, trade-offs that serve as effective representatives of inter-agent constraints. In addition, the nature of the coloring problem captures the strong effect that local changes, made in relatively small subgraphs, might have in significantly altering global constraints—thus requiring coordination mechanisms to maintain legality of the coloring. Finally, the vast literature on graph coloring provides a good foundation for our analysis.

The overall problem is to find a way of coloring vertices of a graph, using a minimal number of colors, such that no two adjacent vertices share the same color. In this paper, we are mainly interested in the economic behavior arising from input graphs, when more than one agent is coloring them.

In a single-agent graph coloring problem, colors are nothing more than markers for keeping track of adjacency or incidence; for the multiagent case, however, every agent will hold a preference relation over the color set. We will need to find conditions that ensure value efficiency of the algorithm's assignment solutions. Also, we would like to know which structures lead cost-minimizing agents to cooperate in coloring. In which cases would agents be *forced* to compete for color resources, as players in a zero-sum game?

The objective properties that might be of interest in our solution, as they are in general negotiation mechanisms (Rosenschein and Zlotkin 1994), include Individual Rationality, Pareto Optimality, Symmetry and Fairness, and Simplicity. We investigate the characteristics of combined graph structures, assembled under every agent's individual constraints, that can be optimally colored both globally and locally, so that each agent may remain at its original level of resource investment. This superimposes on the problem an economic criterion for agreements that are mutually beneficial; every assignment should take into account the color preferences of agents, and strive to fulfill them to their joint possible limit. In the following sections, we will prove the existence of graph classes that make low computational demands on coloring agents, and describe an algorithm to work with these graph classes.

Throughout this paper, we assume a certain benevolence among the agents. Specifically, agents report their preferences and subgraph structure to some central decision maker, and they do not employ manipulative/untruthful strategies. Additionally, we will assume no utility transfer among the agents.

As a specific example, consider the following multi-agent extension to a timetabling scenario (Roberts and Tesman 2004). The civil administration of Lilliput is considering efficient timetabling mechanisms in order to schedule regular weekly committee meetings of three state offices: Housing, Treasury, and Environment. Each office had chosen to divide its work into four regional districts according to cardinal directions, which we denote respectively as $H_{\{n,w,s,e\}}$, $T_{\{n,w,s,e\}}$, $E_{\{n,w,s,e\}}$. Despite this separation, there are weeks when agenda topics lie at a jurisdictional border, and according to Lilliputian law, must include at least 1 member from each related group. For example, when environmental issues concerning the *southwest* territories are discussed at the weekly E_s meeting, one member from the western district must attend this event, in addition to his regular E_w session. Thus environment committee meetings for southern and western districts cannot take place simultaneously—on that particular week. On other occasions, certain subjects may call for members of different offices.

The graph in Figure 1 represents constraints that the Lilliputian administration faces on a typical week, where a vertex is drawn for every weekly meeting, and two vertices are adjacent via an undirected edge if and only if the two meetings have a common member (for example, the conflict described above is represented by the E_w–E_s edge).

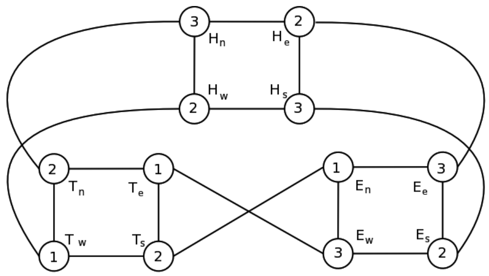

Figure 1: Lilliput administration constraints; a **Moebius Ladder**, painted with inefficient colors.

A coloring of the graph is equivalent to a schedule. All meetings colored ‘Red’ can occur at the same time (say, on day 1); meetings colored ‘Blue’ will occur on day 2, etc. Each office maintains an agent in the global graph, representing their particular constraints. Assuming, for this example, that all committees prefer to finish their work as early in the week as possible, the allocation described in Figure 1 is not Pareto optimal. By switching to color 1 for H_n, H_s, we can obtain a legal globally-optimal coloring, which provides better utility for the Housing agent without harming the Treasury and Environment agents’ values. Moreover, each agent wants an individual-rational coloring, where the number of allocated resources does not exceed its inherent constraint level. Yet the structure described does not allow for this property to hold for all agents, since in every legal 3-coloring of the global graph, there will be an office assigned more than its necessary 2 colors.

The paper is organized as follows. We first describe relevant work in the areas of multiagent systems and graph coloring. We then provide formal definitions for the informal goals stated above. Afterward, we describe conditions for a coloring to be Pareto Optimal, followed by guidelines for allocating fair cost value. We describe the computational hardness of deciding individual rational partitions by investigating global cycle graphs and split graphs. We then conclude and suggest directions for future work.

Related Work

The minimum coloring problem is known to be NP-hard. The decision version asking whether a given graph is k-colorable ($k \geq 3$), was among the first NP-Complete problems identified (Karp 1972). The best-known approximation algorithm for the coloring problem is given in (Halldórsson 1993), and guarantees an exponential approximation ratio. Furthermore, a strong inapproximability result is provided in (Lund and Yannakakis 1994), where it was proved that there is a constant $\epsilon > 0$ such that no polynomial-time approximation algorithm for graph coloring can achieve a ratio n^ϵ unless $P = NP$. With these hardness and inapproximability results, it is reasonable to expect intractability aspects in the multiagent scenario as well.

In the *precoloring extension* problem, a subset of vertices $W \subset V$ is preassigned colors and the goal is to extend the coloring of these vertices to the whole graph, using a minimal number of colors (Biró, Hujter, and Tuza 1992). This coloring variant is NP-Complete for interval graphs. Although we may apply a precoloring algorithm, by first assigning optimal colors to one agent and extending this coloring to the other agents, this context is too restrictive for our purposes, since it does not take into account the different possible colorings that agents’ subgraphs permit.

The *list coloring problem* relates to preference in coloring the vertices of a graph. In this variant, a graph coloring should be found such that the color assigned to a vertex is chosen from a prescribed list of acceptable colors. One can devise lists in a way that cannot be satisfied. In such a case, one may allow some vertices to accept colors not on their original list. A simple way to think of this is to allow some vertices x to expand their list $L(x)$ by adding an additional color (from the available colors). The result (Mahadev and Roberts 2003) can be interpreted as saying that there are G, L such that almost every vertex has to accept a color not on its original list. In light of this evidence, instead of giving every agent the freedom to dictate its acceptable colors, we choose to outfit every agent with a preference relation, and design the system to optimize according to this ordering.

Multiagent graph coloring was studied as a heuristic approach to the classical graph coloring problem, in the context of swarm intelligence. In (Costa and Hertz 1997) each ant’s role is to color the graph in some constructive way, choosing the lowest possible color according to an ordering of the vertices built from a degree-of-saturation parameter.

Graph coloring was used as a benchmark in (Yokoo et al. 1998), where a general framework for solving distributed constraint satisfaction problems is offered. Variables and constraints are distributed among automated cooperative communicating agents. Every node corresponds to an agent, and each agent tries to determine its color so that neighbors do not have the same color.

The current paper looks at graph coloring from an economic perspective; agents may hold more than one variable (node), and coordination is maintained centrally under the assumption of self-motivated, but truthful, agents.

Definitions

A *vertex coloring* of a graph G is a function $c : V(G) \rightarrow L$ such that $c(v) \neq c(w)$ whenever v and w are adjacent. The elements of $L \subseteq \{l_1, \ldots, l_n\}$ are the available colors. The smallest integer k such that G can be colored using k colors is the *chromatic number* of G and is denoted by $\chi(G)$.

In multiagent graph coloring, we will use the following notations and definitions. A set of agents $\mathcal{A}g := \{A_1, \ldots, A_r\}$ induce a partition of a *joint global graph* $\widehat{G}$ into disjoint agent subgraphs G_i, where cross-agent edges are allowed. We mark the vertex set of agent A_i as $V_i := V(G_i)$, and by $\chi_i := \chi(G_i)$, $\alpha(G_i)$ her subgraph's chromatic number and size of a maximal independent set, respectively. The entire set of vertices is $\widehat{V} := V(\widehat{G})$, and $\widehat{\chi} := \chi(\widehat{G})$ marks the chromatic number of the global graph. For $X \subseteq \widehat{V}$, we denote $c(X) = \{c(x) \mid x \in X\}$. Thus a *globally optimal coloring* is a legal coloring c of $\widehat{G}$ such that $|c(\widehat{V})| = \widehat{\chi}$.

Definition. *A globally optimal coloring c is* individual rational *if:*

$$\forall A_i \in \mathcal{A}g \quad |c(V_i)| = \chi_i$$

In some global graphs, such as completely connected graphs (denoted as K_n in graph theory), a *chromatic sum equality* $\chi(\widehat{G}) = \chi_1 + \ldots + \chi_r$ guarantees an individual-rational coloring. However, inequality may hold, and consequently a combination of the individual assignments need not be globally optimal. Moreover, there are structures, such as Figure 1, where no legal mapping exists exhibiting these two properties together.

Definition. *Given a partition $\mathcal{G} = \{G_1, \ldots, G_r\}$ of a global graph $\widehat{G}$, we say that $\mathcal{G}$ is an* individual rational partition *if there is an individual rational coloring c such that $|c(\widehat{V})| = \chi(\widehat{G})$. If no such coloring exists, then $\mathcal{G}$ is called an* irrational partition.

Thus, individual rational partitions enable separate agents to behave in an individual rational way, such that global performance does not decrease. In contrast, in irrational partitions (as in Figure 1), if global efficiency is the decisive measure, there will be an agent forced to pursue a wasteful strategy at its own expense.

Definition. *We say an irrational partition $\mathcal{G}$* lacks individuality of degree d *when,*

$$d := \min_{c:|c(\widehat{V})|=\widehat{\chi}} \max_{A_i \in \mathcal{A}g} \{|c(V_i)| - \chi_i\}$$

Some agents prefer Blues while others like Reds, so in order to describe efficient *value* allocations we will need additional terminology. Every agent has a linear ordering $\prec_i$ on the set of colors L. More formally, $\prec_i$ is a bijection $\prec_i$: $\{1, \ldots, |L|\} \rightarrow L$, and we refer to the j-th favorite color of A_i using $\prec_i (j)$. Since the coloring problem is of a minimization nature, we discuss loss (instead of utility) functions. Accordingly, the cost for agent i using a single color s is $\mathcal{L}_i(s) :=\prec_i^{-1} (s)$, while its total cost according to c is $\mathcal{L}_i^c := \sum_{s \in c(V_i)} \mathcal{L}_i(s)$. Notice that the loss function is indifferent to multiple uses of the same color. This is justified from a scheduling perspective, where all jobs painted using the same color are processed in parallel. We refer to the above valuation form as a *private preference model*, in contrast to a *public preference model* where all agents have the same color preference. In the latter context, it would be convenient to denote L simply as $\{1, \ldots, n\}$.

A coloring function c will be called *Pareto optimal* if there does not exist another coloring f that benefits one of the agents without raising the costs of others. More formally:

Definition. *A coloring c of a global graph $\widehat{G}$ will be called* Pareto optimal *if every other coloring f of $\widehat{G}$ satisfies the following condition:*

$$\exists A_x \quad \mathcal{L}_x^f < \mathcal{L}_x^c \Rightarrow \exists A_y \quad \mathcal{L}_y^f > \mathcal{L}_y^c$$

Furthermore, we would like our allocation to be distributed fairly among the agents. Of course, we need to take into account the conflict level each agent contributed to the joint structure. Denote the minimum loss for agent A_i over all possible optimal colorings of $\widehat{G}$ as $\mathcal{L}_i^*$. The *cost of coordination* for agent A_i using a global coloring c is denoted $\mathcal{D}_i^c$ and signifies the difference $\mathcal{D}_i^c = \mathcal{L}_i^c - \mathcal{L}_i^*$. Correspondingly, by setting the average agent coordination cost $\mu_c = \frac{1}{|\mathcal{A}g|} \sum_{A_i \in \mathcal{A}g} \mathcal{D}_i^c$, a socially fair division would minimize the loss variance among participating agents.

Definition. *A coloring f of a global graph $\widehat{G}$ will be called* Socially Fair *if it satisfies the following equation:*[1]

$$f = \operatorname*{arg\,min}_{c:|c(\widehat{V})|=\widehat{\chi}} \left\{ \frac{1}{|\mathcal{A}g|} \sum_{A_i \in Ag} (\mathcal{D}_i^c - \mu_c)^2 \right\}$$

As a simple example, consider a global graph composed only of 2 adjacent vertices, each of which belongs to a different agent. There will always be a difference of 1 in the coordination cost between the agents under the public preference model, leading to an average coordination cost of $\frac{1}{2}$.

Pareto Efficient Colorings

In this section we will present general techniques for deciding whether a given assignment is Pareto optimal, and offer the means to find such an assignment. We will concentrate mainly on two-agent partitions; conditions for Pareto optimality in general graph partitions of arbitrary size are left for future work.

A first observation is that in the public preference model, any proper coloring of a global graph $\widehat{G} = K_n$ is Pareto optimal. In such scenarios, every color one agent gains is a loss for another, and thus the game is strictly competitive.Other global constructions are described in Lemma 1.

Lemma 1. *For a public value model, let c be a two-agent, individual rational $\widehat{\chi}$-coloring of $\widehat{G}$. The mapping c is Pareto optimal if and only if:*

1. $c(V_1) \cap c(V_2) = \{1, 2, \ldots, p\}$
2. $c(V_1) \,\Delta\, c(V_2) = \{p + 1, \ldots, \widehat{\chi}\}$

where $p := |c(V_1) \cap c(V_2)|$ denotes the number of colors used in both agents' graphs and Δ is the symmetric difference set operator.[2]

From an algorithmic perspective, the Lemma 1 instructs us how to construct a Pareto optimal coloring from an individual-rational, globally optimal one; examine the joint

[1]Other social welfare functions may be specified.

[2]Proofs are omitted due to lack of space.

set of colors and, if necessary, switch them to the first p numbers. Then spread the rest of the colors arbitrarily between the agents. This can be accomplished in $O(|\widehat{V}|)$ time.

In the private preference model we need to take into account different valuations that agents have for colors. Accordingly, it is not true that every coloring of the complete graph is Pareto optimal. As a simple example, consider K_2 where each node belongs to a different agent with the preference profiles 'Red' $\prec_1$ 'Yellow' and 'Yellow' $\prec_2$ 'Red'. Only the coloring that maps A_1's vertex to 'Red', and A_2's vertex to 'Yellow' is Pareto optimal.

To this end we define a *minimum* set for agents A_x and A_y. For a set of colors A and a subset $S \subseteq A$, we say that S is a minimum with respect to A, and denote $S = \mathcal{M}(A)$, if for every subset $T \subseteq A$ such that $|T| = |A|$, we have:

$$\sum_{t \in T} \mathcal{L}_x(t) < \sum_{s \in S} \mathcal{L}_x(s) \Rightarrow \sum_{t \in T} \mathcal{L}_y(t) > \sum_{s \in S} \mathcal{L}_y(s)$$

Lemma 2 describes necessary and sufficient conditions for Pareto optimality in the private preference model.

Lemma 2. *Let c be a two-agent, individual rational $\widehat{\chi}$-coloring of $\widehat{G}$. The mapping c is Pareto optimal if and only if the following conditions hold:*

1. $\mathcal{J} = \mathcal{M}(L)$
2. $l_i \in c(V_x) \setminus \mathcal{J}$ *and* $l_j \in L \setminus c(\widehat{V}) \Rightarrow l_i \prec_x l_j$
3. $l_i \in \mathcal{J}$ *and* $l_j \in c(V_x) \setminus \mathcal{J} \Rightarrow (l_i \prec_y l_j)$
4. $l_i \in c(V_1) \setminus \mathcal{J}$ *and* $l_j \in c(V_2) \setminus \mathcal{J} \Rightarrow (l_i \prec_1 l_j) \vee (l_j \prec_2 l_i)$

where $\mathcal{J} := c(V_1) \cap c(V_2)$ denotes the set of joint colors.

Implementations of conditions 2-4 requires comparisons of agents preference which can accomplished in $O(|\widehat{V}|^2)$ time. To decide whether $\mathcal{J} = \mathcal{M}(L)$ we apply dynamic programing.

Define a function $\phi(i, j)$ as agent A_1's minimal cost for a subset of labels $S \subseteq \{l_1, \ldots, l_i\}$ such that $|S| = |\mathcal{J}|$ and $\sum_{s \in S} \mathcal{L}_2(s) \leq j$. Denoting $w := \sum_{s \in \mathcal{J}} \mathcal{L}_1(s)$ we shall compute $\phi(|L|, w)$. If $\phi(|L|, w) < \sum_{s \in \mathcal{J}} \mathcal{L}_2(s)$ then $\mathcal{J} \neq \mathcal{M}(L)$. Symmetricaly we can check wether A_2's cost can be lowered without raising the cost of A_1, and if both routines return false, we know that $\mathcal{J}$ is a minmum set, as required.

To evaluate $\phi(i, j)$ we maintain a table of dimensions $O(|L| \cdot w)$ and update values accoridng to the following rules. Let $p := |\mathcal{J}|$. For $0 \leq i < p$, set $\phi(i, j) = \infty$ and also $\phi(i, 0) = \infty$. In case $j < [\prec_2^{-1} (l_i)]$ then $\phi(i, j) = \phi(i-1, j)$ since the label l_i is too expensive. For $i = p$ we have $\phi(p, j) = p(p+1)/2$ in case $j \geq [\sum_{i=1}^{p} \prec_2^{-1} (i)]$ and otherwise $\phi(p, j) = \infty$. Finally, for $i > p$ we have $\phi(i, j) = \min\{\phi(i-1, j)\, , \phi(i, j - \prec_2^{-1} (l_i) + i)\}$. Because our input is a graph on $|\widehat{V}|$ nodes, and since $w = O(|\widehat{V}|^2)$ then the complexity of computing ϕ is bounded by $O(|\widehat{V}|^3)$.

Socially Fair Colorings

We now consider the variance of color value among the agents, while preserving Pareto optimality and assuming a given individual-rational allocation.

As a first observation, notice that if $\mathcal{G}$ is an individual rational partition of $\widehat{G}$, then for every agent A_i we have $\mathcal{L}_i^* = \frac{\chi_i(\chi_i+1)}{2}$. Next, we describe two-agent, socially fair colorings of the complete graph.

Lemma 3. *For a public preference model, the following is a socially-fair two-agent coloring of $\widehat{G} = K_n$, for $\chi_1 \geq \chi_2$.*

$$C(V_1) = \Big[\{2i-1\}_{i=1}^{\lfloor \frac{\chi_2}{2} \rfloor} \cup \{\chi_1 + 2i\}_{i=1}^{\lceil \frac{\chi_2}{2} \rceil} \cup \{\chi_2 + i\}_{i=1}^{\chi_1 - \chi_2} \Big] \setminus \{R\}$$

$$C(V_2) = \Big[\{2i\}_{i=1}^{\lfloor \frac{\chi_2}{2} \rfloor} \cup \{\chi_1 + 2i - 1\}_{i=1}^{\lfloor \frac{\chi_2}{2} \rfloor} \Big] \cup \{R\}$$

Where $R := \Big[\frac{\chi_1 + 2 \cdot \chi_2}{2} + Z \cdot (\chi_1 \ mod\ 2)\Big] \cdot \Big[\chi_2 \ mod\ 2\Big]$ *is a remainder term and Z is a random variable such that $Prob[Z = 0] = Prob[Z = 1] = \frac{1}{2}$.*

To interpret the equation above, first consider the situation in which χ_2 is even. Then the remainder term is cancelled, and calculation shows that $\mathcal{D}_1^c = \mathcal{D}_2^c = \frac{\chi_1 \cdot \chi_2}{2}$. Otherwise, the addition and removal of R is designed to equalize the coordination cost between the 2 agents. If both χ_1, χ_2 are odd, there must be an agent incurring an additional unit cost. To satisfy the symmetry property, we use Z to select the latter agent randomly. For example, consider K_8 divided into two equal agents; to achieve a total difference of zero, allocate odds to A_1 and evens to A_2 up to 4, and vice versa for greater values (which amounts to 18 each).

For general graphs, applying the Pareto condition of Lemma 2 we can verify that $\mu_c = \frac{1}{2}[\chi_1 \cdot \chi_2 - \widehat{\chi}(\chi_1 + \chi_2)]$. Thus, in the minimization problem defining socially fair colorings, we need only consider the terms $\mathcal{D}_i^c$. Consequently, in order to construct a socially-fair, Pareto optimal, two-agent coloring from a given individual-rational allocation, we just need to color the graph $K_{\widehat{\chi}-p}$ and add the number p to each of the colors specified by the formulas of Lemma 3.

Solutions for the private preference case are left for future work.

Individual Rational Partitions of Cycles

We now begin our study of individual rational partitions. Our concern here is the number of colors used by each agent, and thus we can assume a public preference model.

We denote the simple chordless cycle on n vertices by C_n, and recall that $\chi(C_n) = 2$ if the number of vertices is even and $\chi(C_n) = 3$ if n is odd. In addition, to formulate the following lemmas, we will use the term *floating agent* to characterize A_i if $|V_i| = \alpha(G_i)$, i.e., it is an agent whose entire vertex set constitutes a stable set.

Even cycles have exactly two globally optimal legal colorings: assigning '1' to even nodes and '2' to odd nodes, and vice versa. To preserve the number of colors used in unicolored agent subgraphs, we must not allow agents to cross these two stable sets.

Lemma 4. *A partition $\mathcal{G}$ of cycle C_{2k} is individual rational iff $\mathcal{G}$ does not include floating agents with vertices of opposite parity.*

We remark that a similar result holds for trees; a partition of a tree is individual rational if and only if no floating agent exists whose vertex set includes vertices at even and odd levels of the tree. Both of these examples belong to the class of *uniquely colorable graphs*, which are graphs that have only one possible optimal coloring up to permutations of colors, and as such are extremely sensitive to distortions.

For odd chordless cycles, we have an extra color, which makes individual rational partitions more complicated. For example, consider the pentagon C_5; one can validate that no matter which partition we choose to apply, an individual rational 3-coloring will be available. More generally, if we restrict our discussion to two agents, then the following lemma holds.

Lemma 5. *A two-agent partition of an odd cycle is individual rational.*

In contrast, when considering three agents, Figure 2 illustrates a partition (of an odd cycle) that lacks individuality. Here, the principal idea suggests that, if a floating agent A_i has starting and terminating vertices in an even path whose internal vertices are composed only from A_j's vertices, then A_i and A_j do not share a common color. This is because, if we have a path of the form $P = v_0, v_1, \ldots, v_{2s}, v_{2s+1}$ such that $v_0, v_{2s+1} \in V_i$ while $v_1, \ldots, v_{2s} \in V_j$, then v_1, v_{2s} must have different colors, and both of them are connected to A_i. In such a scenario we would say that A_i is *blocking* A_j. Arguments along these lines lead to Lemma 6.

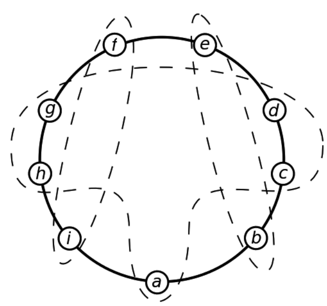

Figure 2: An irrational partition of C_9. $A_1 = \{a, c, d, g, h\}$, $A_2 = \{b, e\}$, $A_3 = \{f, i\}$. A_2 and A_3 are floating agents.

Lemma 6. *A three-agent partition $\mathcal{G}$ of the odd cycle C_{2k+1} is individual rational if and only if $\mathcal{G}$ does not contain two connected floating agents, each one blocking the third agent.*

Theorem 7 states that if we continue increasing the set $\mathcal{A}g$ by choosing finer partitions, we may introduce a high level of distortion to the global graph, and consequently make optimal individual-rational colorings hard to find.

Theorem 7. *For $O(|\widehat{V}|)$ agents, determining whether a partition of C_{2k+1} is individual rational is NP-Complete.*

Surely the problem is in NP—a certificate is a coloring of C_{2k-1}. We can verify such a certificate (in polynomial time) by checking legality, making sure floating agents are unicolored while agents with edges are painted with exactly 2 colors, and confirm that a total of 3 colors were used. For completeness, we reduce the problem `3-COL` to finding such individual rational colorings. Given input graph $\langle G \rangle$, we construct a suitable global odd-cycle and partition $\mathcal{G}$. Let n denote the number of vertices in G and assume w.l.o.g. that G is connected. Expand every vertex $v_i \in V(G)$ to a path P_i of length $2 \times d(v_i)$. Connect paths P_i, P_{i+1} with the addition of an intermediate auxiliary vertex s_i. Finally, construct another path P_{n+1} of 1 or 2 vertices depending on the parity of n, and connect P_{n+1} to P_1, P_n in order to form an odd-cycle. Define $\mathcal{G} = \{G_1, \ldots, G_{2n}\}$; for an odd i, set G_i as the subgraph including the even nodes of $P_{(i+1)/2}$, along with an odd vertex in every path P_j corresponding to a vertex v_j such that $v_i v_j \in E$. For even indexes $i \leq (2n-2)$ set $V(G_i) = \{s_{i/2}\}$ while $G_{2n} = P_{n+1}$ (see Figure 3).

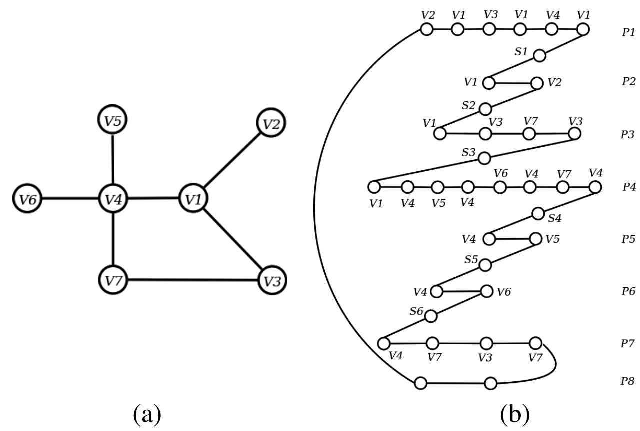

Figure 3: (a) an input to `3-COL`; (b) the corresponding global cycle. Identically marked vertices belong to the same agent.

Although recognizing individual rational partitions is a computationally hard problem, notice that for the odd cycle, by definition, 2 provides an upper bound to the partition's lack of individuality. In fact, from an approximation point of view, we can further lower this bound, independently from the number of participating agents.

Lemma 8. *If $\mathcal{G}$ is a partition of the odd cycle, then $\mathcal{G}$ lacks individuality of degree at most 1.*

From the discussion above, it may appear that the more agents we add, the more complicated our decision for rationality would be. However, bear in mind that for $|\mathcal{A}g| = |\widehat{V}|$, the partition $\mathcal{G}$ is always individual rational (in fact, by iteratively unifying consecutive agent segments, we can symmetrically show that every partition of a global odd cycle to $|\widehat{V}| - 1$ agents is individual rational). In other words, the complexity of recognizing individual rational partitions can be high, even in the restricted case of a global chordless cycle, and is not monotonic with respect to the parameter $|\mathcal{A}g|$.

Split Graphs & Individual Rational Partitions

An undirected graph $G = (V, E)$ is defined to be *split* if there is a partition $V = I \cup K$ of its vertex set into an independent set I and a completely connected graph K (Golumbic 2004). There is no restriction on the edges between vertices of I and K. Given only the degree sequence of a graph, it is possible to recognize in $O(|V|)$

time whether it is split (Hammer and Simeone 1977). In general, the partition $V = I \cup K$ will not be unique; neither will I (resp. K) necessarily be a maximal independent set (resp. clique). Yet from the work in (Hammer and Simeone 1977), we can always arrange, in $O(|V|)$ time, a canonical partition of a split graph, by ensuring that for every $v \in I$ the degree of v is smaller than $|K|$. Henceforth, we will assume an input given with this canonical partition.

Split graphs are perfect graphs, and thus their chromatic number $\chi(G)$ equals the size of the maximal clique in the graph denoted $\omega(G)$. Unlike cycles, split graphs can have a non-constant chromatic number and therefore may lack individuality of an unbounded degree. Figure 4 provides an extreme example for a partition of a split graph, which *fully* lacks individuality. The number of additional colors A_1 would incur, as a result of coordinating with A_2, equals $|V_1| - 1$, which is the highest possible extent.

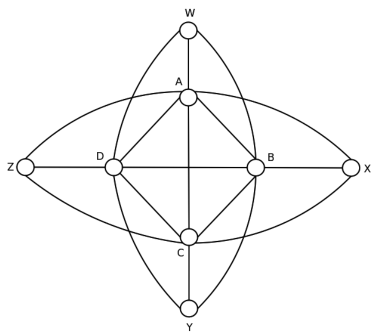

Figure 4: The Lilium, a split with an irrational partition, fully lacking individuality. $A_1 = \{A\text{–}D\}$, $A_2 = \{W\text{–}Z\}$.

We denote the set of vertices of the maximal independent set in $\widehat{G}$, that are also members of agent A_j's subgraph, as $I^{(j)} := V_j \cap I$. Similarly $K^{(j)} := V_j \cap K$ denotes A_j's part in the clique (not to be confused with a completely connected graph on j vertices).

A *border vertex* lies both in an agent's subclique and a sub-independent set. Thus, the set of border vertices is the *boundary of agent* A_j, which we denote as:

$$\mathcal{B}_j = \left\{ v \in I^{(j)} \mid K^{(j)} \subseteq N\left[\{v\}\right] \right\}$$

where $N(H)$ denotes H's set of neighbors and $N[H] = N(H) \cup H$. A vertex $z \in K$ is a *cooperative vertex* w.r.t. agent A_j if $z \notin N(\mathcal{B}_j)$. The main idea behind Theorem 9 is that A_j's boundary can be painted using the same color as A_j's cooperative vertex.

Theorem 9. *A partition of a global split graph is individual rational if and only if every agent has a cooperative vertex.*

As a corollary, we can find in $O(|V| + |E|)$ time an individual rational coloring of a split graph, if one exists.

Conclusions and Future Work

We have presented a framework for economically-oriented multiagent graph coloring, and investigated the dynamics arising from such definitions. Assuming an available individual-rational coloring, sufficient conditions were presented that ensure Pareto optimality, and a procedure was devised to allocate colors fairly among agents. Through the use of limited graph classes, we have seen that individual-rational colorings can be hard to find. However, we are able to quickly recognize individual rational partitions in split graphs. We plan to further investigate the notion of individual rational partitions in other graph classes. Although we have formulated our goals in the context of graph coloring, it possible to study individual-rational partitions in more general optimization problems, such as boolean maximum satisfiability. In parallel, we are working on classical graph coloring heuristics, based upon such decompositions.

References

Biró, M.; Hujter, M.; and Tuza, Z. 1992. Precoloring extension. I. Interval graphs. *Discrete Mathematics* 100(1-3):267–279.

Costa, D., and Hertz, A. 1997. Ants can colour graphs. *Journal of the Operational Research Society* 48:295–305.

Golumbic, M. C. 2004. *Algorithmic graph theory and perfect graphs*, volume 57 of *Annals of Discrete Mathematics*. Amsterdam: Elsevier Science B.V. Second Edition.

Halldórsson, M. M. 1993. A still better performance guarantee for approximate graph coloring. *Information Processing Letters* 45(1):19–23.

Hammer, P. L., and Simeone, B. 1977. The splittance of a graph. Technical Report CORR 77-39, Univ. of Waterloo, Dept. of Combinatorics and Optimization Res.

Karp, R. M. 1972. Reducibility among combinatorial problems. In *Complexity of Computer Computations*. NY: Plenum Press. 85–103.

Lund, C., and Yannakakis, M. 1994. On the hardness of approximating minimization problems. *Journal of the ACM* 41(5):960–981.

Mahadev, N. V. R., and Roberts, F. S. 2003. Consensus list colorings of graphs and physical mapping of DNA. In Janowitz, M. F.; Lapointe, F. J.; McMorris, F. R.; Mirkin, B.; and Roberts, F. S., eds., *Bioconsensus*, volume 61. AMS. 83–95.

Roberts, F. S., and Tesman, B. 2004. *Applied Combinatorics*. Prentice Hall, 2nd edition. chapter 1, 5–7.

Rosenschein, J. S., and Zlotkin, G. 1994. *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*. Cambridge, Massachusetts: MIT Press.

Yokoo, M.; Durfee, E. H.; Ishida, T.; and Kuwabara, K. 1998. The Distributed Constraint Satisfaction Problem: Formalization and Algorithms. *IEEE Transactions on Knowledge and Data Engineering* 10(5):673–685.