Collaborative Procrastination

Aris Anagnostopoulos
Sapienza University of Rome, Italy
http://aris.me
aris@diag.uniroma1.it

Aristides Gionis
KTH Royal Institute of Technology, Stockholm, Sweden
https://www.kth.se/profile/argioni
argioni@kth.se

Nikos Parotsidis
University of Copenhagen, Denmark
https://sites.google.com/view/nikosparotsidis
nipa@di.ku.dk

Abstract

The problem of inconsistent planning in decision making, which leads to undesirable effects such as procrastination, has been studied in the behavioral-economics literature, and more recently in the context of computational behavioral models. Individuals, however, do not function in isolation, and successful projects most often rely on team work. Team performance does not depend only on the skills of the individual team members, but also on other collective factors, such as team spirit and cohesion. It is not an uncommon situation (for instance, experienced by the authors while working on this paper) that a hard-working individual has the capacity to give a good example to her team-mates and motivate them to work harder.

In this paper we adopt the model of Kleinberg and Oren (EC'14) on time-inconsistent planning, and extend it to account for the influence of procrastination within the members of a team. Our first contribution is to model collaborative work so that the relative progress of the team members, with respect to their respective subtasks, motivates (or discourages) them to work harder. We compare the total cost of completing a team project when the team members communicate with each other about their progress, with the corresponding cost when they work in isolation. Our main result is a tight bound on the ratio of these two costs, under mild assumptions. We also show that communication can either increase or decrease the total cost.

We also consider the problem of assigning subtasks to team members, with the objective of minimizing the negative effects of collaborative procrastination. We show that whereas a simple problem of forming teams of two members can be solved in polynomial time, the problem of assigning n tasks to n agents is **NP**-hard.

2012 ACM Subject Classification Applied computing → Economics; Applied computing → Sociology

Keywords and phrases time-inconsistent planning, computational behavioral science, collaborative work, collaborative environments

Digital Object Identifier 10.4230/LIPIcs.FUN.2021.2

Funding *Aris Anagnostopoulos*: Supported by the ERC Advanced Grant 788893 AMDROMA, the EC H2020RIA project "SoBigData++" (871042), and the MIUR PRIN project ALGADIMAR "Algorithms, Games, and Digital Markets".
Aristides Gionis: Supported by three Academy of Finland projects (286211, 313927, 317085), the EC H2020RIA project "SoBigData++" (871042), and the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation.
Nikos Parotsidis: Supported by Grant Number 16582, Basic Algorithms Research Copenhagen (BARC), from the VILLUM Foundation.

1 Introduction

> Procrastination has taught me how to do 30 minutes of work in 8 hours and 8 hours of work in 30 minutes.
>
> – anonymous internet user

The synthesis of teams is a fundamental activity within organizations. The importance of teams in the production of knowledge is increasing. For instance, in the context of scientific research, teams typically produce more frequently cited research than single individuals [15]. Furthermore, it has been observed that simply putting together the best individuals does not necessarily create a great team [11], as there are aspects characterizing effective group members and successful collaborations that are not evident in an individual's performance. When forming teams, it is necessary to take many aspects of the team into account, such as diversity, learning, and cohesion. Although all the aforementioned characteristics play an important role in the performance of a team, they fail to characterize how team dynamics evolve when individuals tend to procrastinate.

In many cases, a project is divided in subtasks, which are assigned to the members of a team. Such a division facilitates cooperation and takes advantage of the different skillsets of the team members. In such situations, typically the final outcome of the project depends on the completion of the subtasks that are assigned to the team members. In this work we assume that the individual members of the team work independently on the subtasks that have been assigned to them, they are aware of the progress that has been made by their teammates, and they do not help each other by working on others' tasks, lacking either the expertise or the incentive to do so. This is a reasonable assumption, especially in the case that each subtask requires different skillsets.

To motivate our setting, consider the following example.

Example. A software company gets assigned a project and the project manager gathers a team of engineers to form a team and work on the project. The project has a number of different subtasks and the project manager recruits one person with the required skills for each subtask (e.g., back-end development, data analytics, user interface). Success in the project depends on completing all subtasks; if one subtask fails, the whole project fails. The team holds regular meetings, sets milestones, discusses problems that occur, and reports progress made in the different subtasks. The collective progress affects the motivation and performance of the team members. An engineer who would normally be motivated to work and would rarely procrastinate might feel unmotivated if the others do not make progress on their subtasks. Conversely, if everyone makes good progress, an engineer who is prone to procrastination might fear that the project will fail because of him and he would put his best effort to keep up with the team. □

In this example it is clear that progress by motivated individuals may help to motivate others. At the same time, motivated individuals can be discouraged to make further progress if they realize that their reward will be unfairly proportional to their effort. Similarly for procrastinating individuals, their "free-ride" attitude may discourage other team members, or they can get motivated by realizing that they are the ones who keep the team behind. Therefore, the overall process is governed by complex dynamics.

Motivated by this discussion, the questions that we study in this paper are the following:

Q1: How can we model interactions among members of a team who work on the same project, such as to capture the dynamics for motivating (or demotivating) each other?

Q2: What are the effects of such a model of interaction among team members, and to what extent the performance of the team can be sped up or slowed down?

Q3: Can we assign optimally team members to the subtasks of a project such as to take advantage of the interactions among the team members and to minimize the total cost of completing a project?

The time-inconsistent planning model. To model procrastination of individual team members we use the time-inconsistent planning model [1, 13, 14]. Here we adopt the formulation introduced by Kleinberg and Oren [8]. We refer to individual team members as agents. According to this model, the progress of an agent for a particular task is represented as a single-source–single-sink directed acyclic graph. The graph simulates a discrete-time process. Each vertex in the graph represents the current state in the project and the progress made so far. An agent being at vertex u at time t picks an edge (u, v) going out of u and moves to vertex v at time $t + 1$. The source vertex represents the start of the project, and the sink vertex the completion of the project. Edge weights model the effort required to move along the edges, and agents are assumed that they try to minimize their total effort to complete the project. An agent with no bias will move from start to completion by following a shortest path from source to sink. To simulate procrastination, the model assumes a present-time bias, where agents perceive the cost at present time higher than what it is in reality. In particular, at any given time, the weights of the outgoing edges from the current vertex are multiplied by a factor $b \geq 1$. The agent calculates the shortest path to the sink using the inflated weights for the next-step edges. This leads agents to bias their choice of next-step edges towards low-cost edges, and as a result they follow paths whose total cost is larger than the cost of the shortest path.

The proposed model. To model collaborative procrastination, and provide an answer to **Q1**, we extend the time-inconsistent planning model, to account for interaction among team members. In particular, we assume that the overall project is divided in subtasks, each team member is assigned to one of the subtasks, and each subtask is represented by a single-source–single-sink directed acyclic graph, which is used to model the actions and progress of the assigned agent to the subtask. We assume that each agent i has a present-time bias b_i. In addition to the original model, we assume that an agent i takes steps towards completing her subtask, the fraction $q_i(t) \in [0, 1]$ capturing the progress made up to time t. The fraction $q_i(t)$ is known to agent i, as well as to all other agents. Given two agents i and j, the difference $q_i(t) - q_j(t)$ expresses the difference in their progress in their respective subtasks, at time t. If $q_i(t) - q_j(t) > 0$ agent i is *ahead* in her subtask and she may feel discouraged by the fact that j has not worked as hard. Conversely, agent j is *behind* and he may feel motivated to catch up. To capture the dynamics of this interaction, we propose to introduce a multiplicative factor $\gamma_i^{q_i(t)-q_j(t)}$ in the present-time bias factor of i, for some $\gamma_i \geq 1$. The effect of our model is that, in addition to the personal present-time bias factor b_i, which captures the tendency of i for procrastination, agents further slow down if they have done more progress than their peers, or speed up if they have done less progress.

Our results. We define formally the *collaborative-procrastination* model, outlined above. To answer research question **Q2** we consider the total cost required by the team to complete a task when they interact and their behavior follows the collaborative procrastination model, and we compare this with the cost that would be required if each agent was working independently on their subtasks. We focus on grid graphs, where at each state of the subtask an agent has two options: make progress or procrastinate. This family of graphs reportedly captures the worst-case task graphs that exhibit the less efficient planning by an agent.

To avoid pathological cases, we consider that the subtask graphs satisfy certain natural assumptions, as proposed by Gravin et al. [7] . Namely, we consider a *bounded distance* property, where in all subtask graphs the optimal path to complete the subtask from any state is never worse than the optimal path from the initial state, and a *monotonicity* property, which ensures that in each subtask graph the cost to complete the subtask does not increase over time. We express our results with respect to the size (n) of the subtask graph, and the number of agents (k) in the team.

Our main result is to show that, assuming the bounded-distance and monotonicity properties, the total cost paid by all agents in the collaboration model, compared to the total cost paid by all agents when they work in isolation, cannot increase more than a factor of $\Theta(n)$. Furthermore, we provide an example, which indicates that this bound is tight.

It is also possible that collaboration helps the overall team performance. We show that, under an additional (mild) assumption for subtask graphs, namely, that procrastinating is less costly at the current step than taking an action towards completing the task, our collaborative model can lead to speeding up the time to complete the overall task by a factor of $\Theta(n)$.

Finally, we turn to our research question **Q3**, for assigning team members to subtasks such as to minimize the total cost of completing the project. We consider a simple version of the problem when the subtask graphs are fixed for all agents, and each agent is characterized by their own present-time bias parameter, and interaction parameters with other agents. We show that even this simple version of the problem is **NP**-hard. We leave as an open problem the design of an efficient approximation algorithm.

2 Model

Our model builds on the graph-theoretic planning model that was introduced by Kleinberg and Oren [8] for a single agent. According to that model, a task is represented by a directed acyclic graph $G = (V, E)$, where each vertex represents a possible state of the task at a specific time point. In this paper we work with a specific family of task graphs that have a grid structure. Specifically, we identify every vertex v_t^ℓ by its time step t and an index $\ell \in \{0, 1, \ldots, \ell_{\max}\}$ indicating the progress that has been made so far towards completion of the task. We assume that no agent is failing her task, even at the expense of heavy cost by the last-minute work. Hence, the vertex set of the task graph consists of all vertices v_t^ℓ with $\ell \geq t$ and $t_{max} - t \geq \ell_{max} - \ell$. See Figure 1 for two examples of graphs.

There is a distinguished start vertex $\sigma = v_0^0$ and a target vertex $\tau = v_{t_{\max}}^{\ell_{\max}}$ that represent the starting point and the completion of the task, respectively. The edge set E contains the following directed edges: (1) an edge $(v_t^\ell, v_{t+1}^{\ell+1})$ for each vertex v_t^ℓ such that for $\ell = 0, \ldots, \ell_{\max} - 1$, $t = \ell, \ldots, t_{\max} - 1$, (2) an edge (v_t^ℓ, v_{t+1}^ℓ) for each vertex v_t^ℓ such that $t_{max} - t \geq \ell_{max} - \ell$ (this condition ensures that the agent does not procrastinate when there is no time for procrastination). For an easy interpretation of the notation, we denote each edge of type (1) by $e_{\nearrow}(v_t^\ell)$ and each edge of type (2) by $e_{\rightarrow}(v_t^\ell)$, and they represent progression and procrastination of the agent at state v_t^ℓ, respectively. Each edge $e = (u, v) \in E$ has a cost $c(e)$, which represents the *effort* to go from state u to state v.

We note that the family of grid graphs is not a compromise. Kleinberg and Oren [8] showed that all graphs that exploit the worst-case behavior of a time inconsistent agent on general directed acyclic graphs, contain as a minor a graph that is trivially simulated by a grid

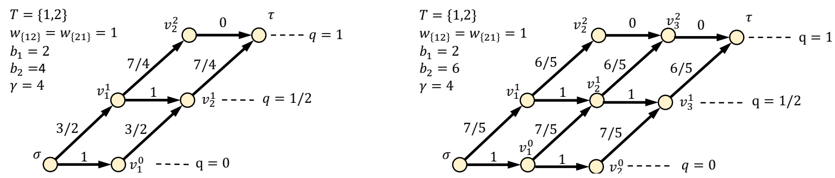

Figure 1 On the left an example task graph where the overall cost of completing the tasks increases compared to the case where the individuals are not aware of each other. On the right, a task graph where the overall cost to complete the tasks decreases.

graph. Specifically, Kleinberg and Oren [8] show that every task graph that forces an agent to follow a path that has exponentially larger cost compared to the optimum path, contains as a minor the graph that has $n+2$ states $\sigma = v_0, v_1, \ldots, v_n, \tau$ and edges (v_i, v_{i+1}) and (v_i, τ) for all $0 \le i \le n$. We construct a grid with states $\sigma = v_0^0, v_1^0, \ldots, v_n^0, v_1^1, v_2^1, \ldots, v_{n+1}^1 = \tau$, and edges (v_i^0, v_{i+1}^0) for all $0 \le i \le n-1$ with cost equal to the cost of the edges (v_i, v_{i+1}) in the worst-case graph; edges (v_i^0, v_{i+1}^1) for $1 \le i \le n$ with cost equal to the cost of the edges (v_i, τ) in the worst-case graph; and edges (v_i^1, v_{i+1}^1) for $0 \le i \le n$ with cost 0. The grid graph essentially splits the state τ of the worst-case graph and replaces it with a path of cost 0 in all of its edges. This path represents completion of the task, as the remaining path to τ is zero.

Given a task graph, present-time biased agents act according to their interpretation of the most effective sequence of actions. Notice that the (objectively) optimum sequence of actions by the agent corresponds to the shortest path in the task graph from σ to τ. An agent who follows the shortest path from a state executes the best actions and minimizes her overall cost. However, according to the quasi-hyperbolic–discounting model [12] the agent misinterprets the cost of her next actions: the costs of all actions in the next step are amplified by a multiplicative factor b. In other words, at state v_t^ℓ the agent perceives the overall effort to accomplish the task as $b \cdot c(e_\nearrow(v_t^\ell)) + d(v_{t+1}^{\ell+1}, \tau)$, if she chooses to make progress, and as $b \cdot c(e_\rightarrow(v_t^\ell)) + d(v_{t+1}^\ell, \tau)$, if she chooses to postpone actions to future time steps. Subsequently, the agent picks the action that minimizes the perceived cost. Throughout the paper we assume that if the perceived cost of making progress equals the perceived cost to procrastinate, the agent chooses to make progress.

In our model we assume that the members of a team T are assigned individual tasks graphs. Each agent (team member) performs on his own task graph. The present-time bias of each agent is affected by two factors, the *personal bias* and the *social bias*. The personal bias $b_i \ge 1$ depends solely on the agent and it does not change throughout the process. As the agents proceed by performing the tasks assigned to them they interact with each other and learn their progress. The unnormalized progress of an agent $i \in T$ at time t is denoted by $r_i(t) \in \{0, \ldots, \ell_{\max}\}$, where $\ell_{\max}$ is the maximum progress level on i's task graph. We define also the (normalized) progress $q_i(t) \in [0,1]$ of agent i as $q_i(t) = r_i(t)/\ell_{\max}$. This allows us to compare the progress of agents with different task graphs and different number of progress levels.

Being aware of the progress made by the other members of the team might affect the motivation of an agent. We assume that agents exert to each other an amount of social influence, which is denoted by a weight $w_{ij} \in [0,1]$, for each pair of agents i and j (in general $w_{ij} \ne w_{ji}$). We assume $w_{ii} = 0$, for all agents $i \in T$. We are now ready to define the social-bias factor of our model.

▶ **Definition 1.** *The social bias of an agent $i \in T$ at time t is denoted by $\Gamma_i^T(t)$ and defined as*

$$\Gamma_i^T(t) = \gamma_i^{\sum_{j\in T} w_{ij}(q_i(t)-q_j(t))},$$

where $\gamma_i \geq 1$ is a social-bias parameter, $w_{ij} \in [0,1]$ is the social influence between agents i and j, and $q_j(t)$ is the (normalized) progress level of each agent $j \in T$ at time t.

Obviously one can consider different functions $\Gamma_i^T(t)$, but in this paper we specialize to this particular form. Notice that our model is a generalization of the model by Kleinberg and Oren because $T = \{i\}$ implies that $\Gamma_i^T(t) = 1$. The simplest, nontrivial, case for our model is when there are only two agents, that is, $T = \{i, j\}$. In that case $\Gamma_i^{\{i,j\}}(t) = \gamma_i^{q_i(t)-q_j(t)}$. Often we assume that we have $\gamma_i = \gamma_j$ for all $i, j \in T$ and in this case we just use $\gamma_i = \gamma$. Whenever $\Gamma_i^T(t) < 1$ we say that agent i is motivated because of the influence of j, and when $\Gamma_i^T(t) > 1$ we say that agent i is discouraged. The following property follows from our model.

▶ **Property 2.** *Consider a team T and an agent $i \in T$, with $\gamma_i \geq 1$ and $w_{ij} \geq 0$, for all $j \in T$. If $q_j(t) \geq q_i(t)$ for all $j \in T$ then $\Gamma_i^T(t) \leq 1$. Similarly, if $q_j(t) \leq q_i(t)$ for all $j \in T$ then $\Gamma_i^T(t) \geq 1$.*

The *present-time bias* of an agent $i \in T$ is defined as $B_i^T(t) = \max\{b_i\Gamma_i^T(t), 1\}$. The present-time bias affects the *perceived* cost of a path for an agent. A *path* p is a sequence of nodes $p = \langle v_1, \dots, v_k\rangle$ such that (v_j, v_{j+1}) is an edge of the graph for all $1 \leq j \leq k-1$. The *cost* of p is $\sum_{j=1}^{k-1} c(v_j, v_{j+1})$. Consider a path $p = \langle v_1, \dots, v_k\rangle$, where $v_1 = v_t^\ell$ is the current state of an agent $i \in T$. Then the *perceived cost* for an agent i for the path p is

$$B_i^T(t) \cdot c(v_1, v_2) + d(v_2, \tau).$$

An agent, who aims to complete the assigned task G, follows a path from σ to τ. Note that each such path is of the form $p = \langle \sigma = v_0^0, v_1^{\ell_1}, \dots, v_{t_{\max}}^{\ell_{\max}} = \tau\rangle$. We call such a path *progress path* on task G. Given two progress paths $p = \langle \sigma = v_0^0, v_1^{\ell_1}, v_2^{\ell_2}, \dots, v_{t_{\max}}^{\ell_{\max}} = \tau\rangle$ and $p' = \langle \sigma = v_0^0, v_1^{\ell_1'}, v_2^{\ell_2'}, \dots, v_{t_{\max}}^{\ell_{\max}} = \tau\rangle$ on the same task G, we say that p *is above* path p', and write $p \succeq_G p'$, if for each $t = 1, \dots, t_{\max}$ we have that $\ell_t \geq \ell_t'$. We define $\preceq_G$ (i.e., *below*) analogously. Let $p^T(i)$ be the progress path of agent $i \in T$. We use $p(i)$ to denote $p^{\{i\}}(i)$.

Example 1.[1] We demonstrate our model with two examples in the case of teams with two members. In the first example, in Figure 1 (left), the total cost of the progress paths followed by the two agents increases when the two agents interact compared to the case where the two agents do not have knowledge of each other's progress. The two agents operate on the same task graph. Agent 1 has personal bias $b_1 = 2$ and agent 2 has $b_2 = 4$. If agent 1 would operate independently of agent 2, it would start from state $\sigma = v_0^0$, and would evaluate the options of following either edge $e_\nearrow(v_0^0)$ (i.e., to make progress) or edge $e_\rightarrow(v_0^0)$ (i.e., procrastinate). The perceived cost of following edge $e_\nearrow(v_0^0)$ is $\Gamma_1(0) \cdot b_1 \cdot c(e_\nearrow(v_0^0)) + d(v_1^1, \tau) = 19/4$, whereas the perceived cost of following the edge $e_\rightarrow(v_0^0)$ is $\Gamma_1(0) \cdot b_1 \cdot c(e_\rightarrow(v_0^0)) + d(v_1^0, \tau) = 21/4$. Therefore, at state σ agent 1 would follow edge $e_\nearrow(v_0^0)$. At state v_1^1 the agent would again proceed based on the perceived cost of following either edge $e_\nearrow(v_1^1)$ or $e_\rightarrow(v_1^1)$. The perceived cost of following edge $e_\nearrow(v_1^1)$ is $\Gamma_1(1) \cdot b_1 \cdot c(e_\nearrow(v_1^1)) + d(v_2^2, \tau) = 14/4$, whereas the perceived cost of following the edge $e_\rightarrow(v_1^1)$ is $\Gamma_1(1) \cdot b_1 \cdot c(e_\rightarrow(v_1^1)) + d(v_2^1, \tau) = 15/4$. Therefore, at

[1] In the appendix we provide the examples with the calculations performed explicitly.

state v_1^1 agent 1 would follow edge $e_\nearrow(v_1^1)$. For the last edge, the agent has only the option to follow the edge $e_\rightarrow(v_2^2)$ with cost 0 to reach τ. Hence, the cost of the progress path of agent 1 when operating individually would be 13/4.

Similarly, if agent 2 operates independently of agent 1, the perceived cost of following edge $e_\nearrow(v_0^0)$ is $\Gamma_2(0) \cdot b_2 \cdot c(e_\nearrow(v_0^0)) + d(v_1^1, \tau) = 31/4$, whereas the perceived cost of following the edge $e_\rightarrow(v_0^0)$ is $\Gamma_2(0) \cdot b_2 \cdot c(e_\rightarrow(v_0^0)) + d(v_1^0, \tau) = 29/4$. Therefore, at state v_0^0 agent 2 would follow edge $e_\rightarrow(v_0^0)$. At state v_1^0 the agent has no other options that to follow the edges $e_\nearrow(v_1^0)$ and then the edge $e_\nearrow(v_2^1)$ to reach τ. Hence, the cost of the progress path of agent 2 when operating individually would be 17/4.

Now we analyze the behavior of agents 1 and 2 when they collaborate on the same project and they both have to perform the same task. At time step $t = 0$, we have that the social bias is $\Gamma_1^{\{1,2\}}(0) = \gamma^{q_1(0)-q_2(0)} = 4^0 = 1$. Analogously, we have that $\Gamma_2^{\{1,2\}}(0) = 1$. Therefore, the choice of each agent at time step $t = 0$ is the same as when they perform independently as their personal bias remains unchanged. That is, agent 1 follows the edge $e_\nearrow(v_0^0)$ making progress 1 at time $t = 1$ and agent 2 follows the edge $e_\rightarrow(v_0^0)$ making no progress at time $t = 1$. At time step $t = 1$, agent 1 evaluates the options of following edge $e_\nearrow(v_1^1)$ or $e_\rightarrow(v_1^1)$. Notice that now the social bias of agent 1 is $\Gamma_1^{\{1,2\}}(1) = \gamma^{q_1(1)-q_2(1)} = 4^{1/2} = 2$. Hence, the perceived cost of following edge $e_\nearrow(v_1^1)$ is $\Gamma_1^{\{1,2\}}(1) \cdot b_1 \cdot c(e_\nearrow(v_1^1)) + d(v_2^2, \tau) = 28/4$, whereas the perceived cost of following the edge $e_\rightarrow(v_1^1)$ is $\Gamma_1^{\{1,2\}}(1) \cdot b_1 \cdot c(e_\rightarrow(v_1^1)) + d(v_2^1, \tau) = 23/4$. Therefore, at state v_1^1 agent 1 would follow edge $e_\rightarrow(v_1^1)$. For the last edge, agent 1 has only one option, that is, to follow the edge $e_\nearrow(v_2^1)$ to reach τ. Hence, the cost of the progress path of agent 1 is 17/4, compared to the cost 13/4 of the progress path that it would follow independently. The progress path of agent 2 does not change when operating with agent 1, as after the first choice to follow edge $e_\rightarrow(v_0^0)$ there are not alternative paths that agent 2 could follow. In conclusion, the total cost of the two agents when operating together is 34/4 compared to the 30/4 when operating independently.

Example 2. We now proceed with an example where the collaboration of two agents leads to a decrease to the total cost of their progress paths. Consider the case where two agents 1, 2 with personal biases $b_1 = 2, b_2 = 6$ operate on the task graph in Figure 1 (right). It can be verified that the progress path of agent 1 is $p(i) = \langle \sigma = v_0^0, v_1^1, v_2^2, v_3^2, v_4^2 = \tau \rangle$ and the progress path of agent 2 is $p(2) = \langle \sigma = v_0^0, v_1^0, v_2^0, v_3^1, v_4^2 = \tau \rangle$. Therefore, the total cost of the $p(1)$ and $p(2)$ is 36/5.

Now we consider the case where the two agents interact with each other. Similarly to the first example, at time step $t = 0$ the social bias is $\Gamma_1^{\{1,2\}}(0) = \Gamma_2^{\{1,2\}}(0) = 1$ and hence the choices at time $t = 0$ of agents 1, 2 are the same as in the case where they operate independently. That is, $q_1(1) = 1/2$ and $q_2(1) = 0$. At time step $t = 1$, the social bias of agent 1 is $\Gamma_1^{\{1,2\}}(1) = \gamma^{q_1(1)-q_2(1)} = 4^{1/2} = 2$. According to our model, the perceived cost of following edge $e_\nearrow(v_1^1)$ by agent 1 is $\Gamma_1^{\{1,2\}}(1) \cdot b_1 \cdot c(e_\nearrow(v_1^1)) + d(v_2^2, \tau) = 24/5$, whereas the perceived cost of following the edge $e_\rightarrow(v_1^1)$ is $\Gamma_1^{\{1,2\}}(1) \cdot b_1 \cdot c(e_\rightarrow(v_1^1)) + d(v_2^1, \tau) = 26/5$. Therefore, at state v_1^1 agent 1 follows edge $e_\nearrow(v_1^1)$. We now review the decision of agent 2 at time $t = 1$, whose social bias is $\Gamma_2^{\{1,2\}}(1) = \gamma^{q_2(1)-q_1(1)} = 4^{-1/2} = 1/2$. Hence, agent 2 at time $t = 1$ perceives cost $\Gamma_2^{\{1,2\}}(1) \cdot b_2 \cdot c(e_\nearrow(v_1^0)) + d(v_2^1, \tau) = 27/5$ for following edge $e_\nearrow(v_1^0)$, and cost $\Gamma_2^{\{1,2\}}(1) \cdot b_1 \cdot c(e_\rightarrow(v_1^0)) + d(v_2^0, \tau) = 28/5$ for following edge $e_\rightarrow(v_1^0)$. Therefore, at state v_1^0 agent 2 follows edge $e_\nearrow(v_1^0)$ to reach state v_2^1. At time $t = 2$, we have $q_1(2) = 1, q_2(2) = 1/2$. Agent 1 has no options other than to follow the path $\langle v_2^2, v_3^2, \tau \rangle$ to reach τ. The social bias of agent 2 at time $t = 2$ is $\Gamma_2^{\{1,2\}}(2) = \gamma^{q_2(2)-q_1(2)} = 1/2$. Hence, the perceived cost of agent 2 at

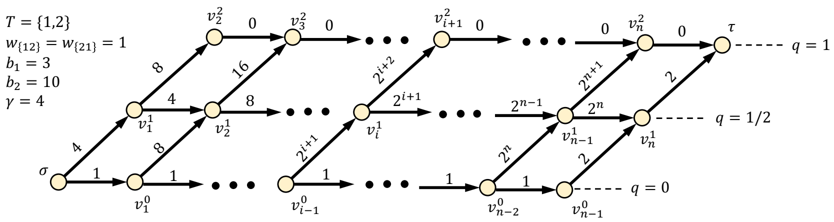

Figure 2 An example where the interaction of two agents increases the total cost of the two progress paths exponentially, even when they operate on the same task graph. When the agents operate independently, they follow the progress paths $p(1) = \langle \sigma, v_1^1, v_2^2, v_3^2, \dots, \tau \rangle$, $p(2) = \langle \sigma, v_1^0, \dots, v_{n-1}^0, v_n^1, \tau \rangle$, with total cost $\Theta(n)$. When the agents collaborate, they follow the progress paths $p^{\{1,2\}}(1) = \langle \sigma, v_1^1, v_2^1, \dots, v_n^1, \tau \rangle$, $p^{\{1,2\}}(2) = \langle \sigma, v_1^0, \dots, v_{n-1}^0, v_n^1, \tau \rangle$, with total cost $\Theta(2^n)$.

time $t = 2$ in the case of following the edge $e_{\nearrow}(v_2^1)$ is $\Gamma_2^{\{1,2\}}(2) \cdot b_2 \cdot c(e_{\nearrow}(v_2^1)) + d(v_3^2, \tau) = 18/5$, and in the case of following the edge $e_{\rightarrow}(v_2^1)$ is $\Gamma_2^{\{1,2\}}(2) \cdot b_2 \cdot c(e_{\rightarrow}(v_2^1)) + d(v_3^1, \tau) = 21/5$. Therefore, at state v_2^1 agent 2 follows edge $e_{\nearrow}(v_2^1)$. Finally, at time $t = 3$, both agents have no other option than to follow edge (v_3^2, τ) to reach τ. In conclusion, agent 1 follows the progress path $p^{\{1,2\}}(1) = \langle \sigma = v_0^0, v_1^1, v_2^2, v_3^2, v_4^2 = \tau \rangle$ and agent 2 follows that progress path $p^{\{1,2\}}(2) = \langle \sigma = v_0^0, v_1^0, v_2^1, v_3^2, v_4^2 = \tau \rangle$, with total cost 31/5. That is, if agents 1, 2 collaborate they decrease the total cost.

▶ **Corollary 3.** *The total cost of the progress paths of a team can either decrease or increase (or, of course, remain the same) compared to the total cost of the progress paths of the team members when they operate in isolation (i.e., with no communication) on the same tasks.*

3 Limitations and Further Assumptions

We now show that the vanilla version of our model can lead to unnatural phenomena in the interaction of the agents in a team. We construct examples having two interacting agents. Guided by these extreme behaviors we make a set of reasonable assumptions that eliminate those unnatural phenomena. Similar assumptions have been made previously for the behavior of individual agents in absence of a team. More specifically, Gravin et al. [7] showed that the progress path of a time-inconsistent agent can have exponentially larger cost compared to the optimal progress path on a task graph. Here, we extend their example to show that there can be an exponential increase to the total cost of the progress paths of two agents, compared to the case where they operate individually. Our example is depicted in Figure 2. We note that unlike Gravin et al. [7] , where the cost of an agent can be exponentially larger compared to the optimal progress path, which was never an option of the agent, in our case the increase in the total cost is compared to the progress path in the case where the agents operate individually.

The main reason behind the exponential increase in the total cost of two agents is that the optimal cost of completing the task can increase at a future state in the progress path of an agent. In many scenarios, this is an unnatural phenomenon as it implies that the required effort to complete a task increases exponentially over time. Gravin et al. [7] introduce the following two assumptions on the task graph that eliminate such pathological instances.

▶ **Property 4** (Bounded-distance property)**.** *Let G be a task graph. For every vertex $v \in V(G)$, it is $d(v, \tau) \leq d(\sigma, \tau)$.*

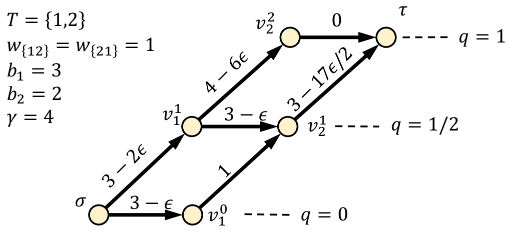

Figure 3 An example where the motivator affects the procrastinator to procrastinate further.

The bounded-distance property allows only task graphs in which the optimal path to complete the task from any state is never worse than the optimal progress path from the initial state. This is a natural assumption in a plethora of real-world tasks. For instance, this includes the tasks in which starting over is always a free and feasible option. Gravin et al. [7] show that for task graph with the bounded-distance property the cost of an agent increases by at most a factor of $\mathcal{O}(n)$, compared to the optimal progress path.

▶ **Property 5** (Monotone-distance property). *For every transition from a vertex u to a vertex v, where $u, v \in V(G)$, it holds that $d(v, \tau) \leq d(u, \tau)$.*

The monotone-distance property of task graphs implies that the cost to complete the task does not increase over time. Notice that the monotone-distance property implies the bounded-distance property: any graph with the monotone-distance property also has the bounded-distance property. Gravin et al. [7] show that if the task graph has the monotone-distance property, and the present-time bias of the agent is drawn from a restricted distribution, then the cost of the progress path compared to the optimal progress path is bounded by a factor much smaller than n.

In Section 4 we study the behavior of agents in task graphs that obey Properties 4 and 5. More specifically, we show that the total cost of progress paths by all agents cannot increase more than a factor of $\mathcal{O}(n)$, compared to the cost of the progress paths in the case where the agents operate individually. We further show that this bound is tight.

In our model, we consider the interaction of two or more agents, which introduces further pathological scenarios in the behavior of the agents. For instance, consider the example in Figure 3, where both agents in a team $T = \{1, 2\}$ operate on the same task graph. Agent 1 has a higher bias than agent 2, and follows the optimal progress path (that is, the path $p(1) = \langle \sigma = v_0^0, v_1^0, v_2^1, v_3^2 = \tau \rangle$ with cost $7 - 9.5\epsilon$), while the agent with lower bias follows a progress path with larger cost (that is, the path $p(2) = \langle \sigma = v_0^0, v_1^1, v_2^2, v_3^3 = \tau \rangle$ with cost $7 - 8\epsilon$). This phenomenon is unnatural as in this example the motivated individual (i.e., the agent with smaller personal bias) follows a progress path with larger cost compared to the procrastinating individual (i.e., the agent with larger personal bias). Moreover, when the two agents in Figure 3 interact, the motivated individual causes the procrastinating individual to further procrastinate (follow a progress path with larger cost). To eliminate such behaviors, we introduce an additional assumption on the task graph. Our assumption is that from any state the action leading to progress costs more than the action of postponing the progress.

▶ **Property 6.** *Given a task graph $G = (V, E)$ it holds that $c(e_{\nearrow}(v_t^\ell)) \geq c(e_{\rightarrow}(v_t^\ell))$, for all $v_t^\ell \in V$.*

4 Team Behavior on Task Graphs

We now study the behavior of time-inconsistent agents in teams under our model. We begin by bounding the change in the total cost of all progress paths compared to the cost in the case where the agents operate individually. The objective is to bound the maximum loss on the total effort made by the team when the agents communicate their progress, compared to the case where the agents operate individually.[2]

▶ **Lemma 7.** *Let $T = \{i_1, \ldots, i_k\}$ be a team of agents, where $w_{ij} = 1$, for all $i, j \in T$, operating on task graphs $G_1, \ldots, G_k$, where all task graphs should be completed in $t_{\max} = n$ time steps and all task graphs have Properties 4 and 5. The scenario in which the agents collaborate can lead to total cost of their progress paths that is larger than the case where they operate individually by a factor $\Omega(n)$.*

Gravin et al. [7] showed that the cost of an agent on a graph with Properties 4 and 5 cannot exceed n times the cost of the shortest path. Their proof suffices to prove the following lemma.

▶ **Lemma 8.** *Let $T = \{i_1, \ldots, i_k\}$ be a team of agents, where $w_{ij} = 1$, for all $i, j \in T$, operating on task graphs $G_1, \ldots, G_k$, where all task graphs have $t_{\max} = n$ time steps and all task graphs have Properties 4 and 5. Collaboration can increase by at most a factor of n the total cost spent by the agents to accomplish the assignment compared to the total cost of the agents operating individually.*

We now provide a lower bound on the speedup that the collaboration in a team can achieve.

▶ **Lemma 9.** *Let $T = \{i_1, \ldots, i_k\}$ be a team of agents, where $w_{ij} = 1$, for all $i, j \in T$, operating on task graphs $G_1, \ldots, G_k$, where all task graphs should be completed in $t_{\max} = n$ time steps and all task graphs obey Properties 4, 5 and 6. The total cost may decrease by a factor of $\Omega(n)$ due to collaboration.*

Agents operating on identical task graphs. We now compare the progress paths of the agents and the way they relate to each other, in the case where all agents perform on the same task graph. We begin with the following lemma that states that the order of the progress paths of the agents is the same as the reverse order of their personal biases.

▶ **Lemma 10.** *Consider two agent $i, j \in T$ operating on the same task graph G. If $b_i \geq b_j$, then $p(i) \preceq_G p(j)$.*

Next we relate all progress paths when the agents collaborate with respect to the progress paths of the agents with the maximum and minimum personal biases. That is, throughout the process of collaboration in a team, no agent does more (resp., less) progress than the most (resp., least) motivated agent does independently, at any time. The lemma suggests that all progress paths in the case where the agents collaborate are between the progress paths of the most motivated and the least motivated agents when operating individually. We call this the *envelope property*.

▶ **Lemma 11** (Envelope property). *Consider a team with agents $T = \{i_1, \ldots, i_k\}$, with $b_{i_1} \geq \cdots \geq b_{i_k}$, operating over the same task graph G. For each $i \in T$ we have that $p^T(i) \succeq_G p(i_1)$ and $p^T(i) \preceq_G p(i_k)$.*

[2] The proof of this and further results appear in the appendix.

We now have developed a better understanding on the interaction between agents of a team operating on the same task graph. Next, we use these results to further bound the ratio of the total cost of progress paths when the team collaborates compared to the cost of operating individually. We observe that this setting still allows examples in which the cost can increase by a factor of $\Omega(n)$ (as in Figure 4, used to prove Lemma 7). To cope with such extreme examples, we introduce the following restriction on the task graph on which the team operates.

▶ **Property 12.** *For every two vertices $v_i^j, v_i^l \in V(G)$, where $j \leq l$, it holds that $c(v_i^j, v_{i+1}^j) \geq c(v_i^l, v_{i+1}^l)$.*

Essentially, Property 12 implies that procrastinating at a specific time step cannot cost more if the agent made more progress compared to the case where the agent made less progress at the same time step. This is a reasonable restriction to the structure of the task graph. We acknowledge, however, that there exist scenarios where Property 12 is not natural. For instance, such a scenario appears in the case of lab experiments where the procrastination of an agent after starting the experiment might lead to a waste of the whole experiment (i.e., the resources), while postponing the starting time of the experiment simply delays the whole process.

▶ **Lemma 13.** *Let T be a team of k agents operating on the same task graph, which has Properties 4, 5, 6, and 12. The total cost of all progress paths when the agents collaborate is at most k times higher than the sum of cost of progress paths when the agent operate independently.*

5 Assignment Problems

Until now we studied the scenario where the assignment of task graphs to agents is given in advance. Another natural scenario is when a given task can be assigned to more than one agents with similar skills. Can we then determine the best assignment so as to minimize the cost due to procrastination? A simple special case is when a project consists of $n/2$ identical tasks and there exist n agents which should be grouped into $n/2$ two-member teams, such that the total cost payed by all agents is minimized.

▶ **Problem 14.** *Assume that we are given $n/2$ copies of the same task graph G, n agents $1, \ldots, n$ with personal biases $b_1, \ldots, b_n$ and social-influence weights w_{ij} for all $1 \leq i, j \leq n, i \neq j$. The goal is to partition the n agents into $n/2$ two-member teams, such that when the two agents of each team work on the common task specified by G, the total cost over all agents is minimized.*

Problem 14 has a simple solution. For each of the $\binom{n}{2}$ pairs of agents, we can compute the cost of the two agents collaborating together on the given task. We obtain a complete weighted graph where each node represents an agent and each edge weights correspond to the total cost of the two agents paired as a team. The problem then reduces to finding a minimum-weight matching.

Assume that we have a single team of n agents, one project consisting of n tasks, each one having its own task graph, and we need to assign one agent to each of the tasks, so as to minimize the total cost of finishing the project. Without much loss of generality we assume that each team member can perform all tasks. The optimal assignment problem takes the following form.

▶ **Problem 15** (OptimalGroupAssignment). *Consider n task graphs $G_1, \ldots, G_n$, n agents $1, \ldots, n$ with personal biases $b_1, \ldots, b_n$ and social-influence weights w_{ij} for all $1 \leq i, j \leq n, i \neq j$. The task is to assign one agent to each task, such that the total cost of completing all the tasks in the collaborative-procrastination model is minimized.*

We next prove that this problem is hard.

▶ **Theorem 16.** *The OptimalGroupAssignment problem is* **NP**-*hard.*

6 Related Work

Some of the first studies in economics attempting to formulate time-inconsistent planning behavior was the work of Strotz [14] and Pollak [13]. The theory of time-inconsistency developed to what is called *quasi-hyperbolic discounting* Laibson [12], Frederick et al. [5]. The theory provides a natural way to model the decision of an agent to procrastinate, using the notion of *present-time bias* – the tendency to view costs and benefits that are incurred at the present moment to be more salient than those incurred in the future. Kleinberg and Oren [8] propose a graph-theoretic model, in which dependencies among actions are represented by a directed acyclic graph, and a time-inconsistent agent follows a path through this graph based on the agent's biased evaluation of the actions at each time step. Kleinberg and Oren [8] characterize the worst-case procrastination ratio, and they consider the problem of reducing the procrastination cost by deleting nodes and/or edges from an underlying graph.

Gravin et al. [7] consider the case where the present-time bias of the agent is drawn at each time step at random, from a distribution $\mathcal{F}$, They characterize the worst possible cost of a path chosen by an agent compared to the cost of the optimal path, and under reasonable assumptions they provide bounds for this ratio. Kleinberg et al. [9] model the behavior of sophisticated agents – agents who are aware of their tendency to procrastinate and they plan in advance. Their study includes tight upper bounds on the procrastination relatively to the optimal path in a task graph. Kleinberg et al. [10] consider the interaction of multiple biases on an agent's behavior: they study the interaction of present-time bias factor and sunk-cost bias factor – the tendency to incorporate costs incurred in the past into ones plans for the future, even when these past costs are no longer relevant to optimal planning. Moreover, based again on the model of Kleinberg and Oren [8], several studies consider optimization problems where the objective is to minimize the cost of the path followed by an agent [2, 4, 3].

Gans and Landry [6] consider the interaction of teams with two present-time biased agents who collaborate to accomplish a common goal. They assume that both agents can accomplish all subtasks, and that the agents can either be sophisticated or be naïve – in the sense that they either know their present-time bias factor or not. The objective of each agent is to complete the task with the minimum possible effort from their side. The model of Gans and Landry [6] is different from ours as progress can be done by any agent, and at each time step there is no distinction with respect to which agent achieved the progress in the previous step.

7 Conclusion and Open Problems

In this paper we extended the model of Kleinberg and Oren [8] on time-inconsistent planning into settings where individuals are members of a team and the decision on whether to perform a task or postpone it depends on the progress of the other team members. Our model incorporates phenomena that are encountered in real life: participating in a team can motivate (or demotivate) individuals compared to when they work individually. In the proposed setting we showed how different assumptions allow to deduce the extent that participation to a team may increase or decrease performance. We also showed that our model can be used to define matching and team-formation problems, when the goal is to form teams that keep the members motivated.

Whereas our model captures some elements of how agents in teams may collaborate, there are many other modeling choices. Often, the load of one member who has not progressed may be transferred to other team members; this may lead to free-riding phenomena, and a game-theoretic approach may be suitable to model such settings. Note that Lemma 11 implies that the effort of a member who participates on a team cannot exceed the one of the most efficient member; in particular, it implies that the most efficient member cannot improve by participating in a team. Often this is not the case: for instance, one can attempt to model *competition* between team members, which may lead to more efficient performance.

References

1. George A Akerlof. Procrastination and obedience. *The American Economic Review*, 81(2):1–19, 1991.
2. Susanne Albers and Dennis Kraft. Motivating time-inconsistent agents: A computational approach. In *Web and Internet Economics*, pages 309–323, 2016.
3. Susanne Albers and Dennis Kraft. On the Value of Penalties in Time-Inconsistent Planning. In *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, pages 10:1–10:12, 2017.
4. Susanne Albers and Dennis Kraft. The price of uncertainty in present-biased planning. In *Web and Internet Economics*, pages 325–339, 2017.
5. Shane Frederick, George Loewenstein, and Ted O'donoghue. Time discounting and time preference: A critical review. *Journal of economic literature*, 40(2):351–401, 2002.
6. Joshua S Gans and Peter Landry. Procrastination in teams. Technical report, National Bureau of Economic Research, 2016.
7. Nick Gravin, Nicole Immorlica, Brendan Lucier, and Emmanouil Pountourakis. Procrastination with variable present bias. In *Proc. of the 2016 ACM Conference on Economics and Computation*, EC '16. ACM, 2016.
8. Jon Kleinberg and Sigal Oren. Time-inconsistent planning: a computational problem in behavioral economics. In *Proc. of the fifteenth ACM conference on Economics and computation*, EC '14, pages 547–564. ACM, 2014.
9. Jon Kleinberg, Sigal Oren, and Manish Raghavan. Planning problems for sophisticated agents with present bias. In *Proc. of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 343–360, New York, NY, USA, 2016. ACM. doi:10.1145/2940716.2940764.
10. Jon Kleinberg, Sigal Oren, and Manish Raghavan. Planning with multiple biases. *arXiv preprint arXiv:1706.01062*, 2017.
11. Jon Kleinberg and Maithra Raghu. Team performance with test scores. In *Proc. of the Sixteenth ACM Conference on Economics and Computation*, pages 511–528. ACM, 2015.
12. David Laibson. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478, 1997.
13. Robert A Pollak. Consistent planning. *The Review of Economic Studies*, 35(2):201–208, 1968.
14. Robert Henry Strotz. Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, 23(3):165–180, 1955.
15. Stefan Wuchty, Benjamin F Jones, and Brian Uzzi. The increasing dominance of teams in production of knowledge. *Science*, 316(5827):1036–1039, 2007.

A Appendix

Examples of Section 2 with Calculations

In this section we provide the calculations used for the examples in Section 2.

Example 1. We demonstrate our model with two examples in the case of teams with two members. In the first example, in Figure 1 (left), the total cost of the progress paths followed by the two agents increases when the two agents interact compared to the case where the two agents do not have knowledge of each other's progress. The two agents operate on the same task graph. Agent 1 has personal bias $b_1 = 2$ and agent 2 has $b_2 = 4$. If agent 1 would operate independently of agent 2, it would start from state $\sigma = v_0^0$, and would evaluate the options of following either edge $e_\nearrow(v_0^0)$ (i.e., to make progress) or edge $e_\rightarrow(v_0^0)$ (i.e., procrastinate). The perceived cost of following edge $e_\nearrow(v_0^0)$ is $\Gamma_1(0) \cdot b_1 \cdot c(e_\nearrow(v_0^0)) + d(v_1^1, \tau) = 1 \cdot 2 \cdot \frac{3}{2} + 7/4 + 0 = \frac{19}{4}$, whereas the perceived cost of following the edge $e_\rightarrow(v_0^0)$ is $\Gamma_1(0) \cdot b_1 \cdot c(e_\rightarrow(v_0^0)) + d(v_1^0, \tau) = 1 \cdot 2 \cdot 1 + \frac{3}{2} + \frac{7}{4} = \frac{21}{4}$. Therefore, at state σ agent 1 would follow edge $e_\nearrow(v_0^0)$. At state v_1^1 the agent would again proceed based on the perceived cost of following either edge $e_\nearrow(v_1^1)$ or $e_\rightarrow(v_1^1)$. The perceived cost of following edge $e_\nearrow(v_1^1)$ is $\Gamma_1(1) \cdot b_1 \cdot c(e_\nearrow(v_1^1)) + d(v_2^2, \tau) = 1 \cdot 2 \cdot \frac{7}{4} + 0 = \frac{14}{4}$, whereas the perceived cost of following the edge $e_\rightarrow(v_1^1)$ is $\Gamma_1(1) \cdot b_1 \cdot c(e_\rightarrow(v_1^1)) + d(v_2^1, \tau) = 1 \cdot 2 \cdot 1 + \frac{7}{4} = \frac{15}{4}$. Therefore, at state v_1^1 agent 1 would follow edge $e_\nearrow(v_1^1)$. For the last edge, the agent has only the option to follow the edge $e_\rightarrow(v_2^2)$ with cost 0 to reach τ. Hence, the cost of the progress path of agent 1 when operating individually would be $\frac{3}{2} + \frac{7}{4} + 0 = \frac{13}{4}$.

Similarly, if agent 2 operates independently of agent 1, the perceived cost of following edge $e_\nearrow(v_0^0)$ is $\Gamma_2(0) \cdot b_2 \cdot c(e_\nearrow(v_0^0)) + d(v_1^1, \tau) = 1 \cdot 4 \cdot \frac{3}{2} + \frac{7}{4} + 0 = \frac{31}{4}$, whereas the perceived cost of following the edge $e_\rightarrow(v_0^0)$ is $\Gamma_2(0) \cdot b_2 \cdot c(e_\rightarrow(v_0^0)) + d(v_1^0, \tau) = 1 \cdot 4 \cdot 1 + \frac{3}{2} + \frac{7}{4} = \frac{29}{4}$. Therefore, at state v_0^0 agent 2 would follow edge $e_\rightarrow(v_0^0)$. At state v_1^0 the agent has no other options that to follow the edges $e_\nearrow(v_1^0)$ and then the edge $e_\nearrow(v_2^1)$ to reach τ. Hence, the cost of the progress path of agent 2 when operating individually would be $1 + \frac{3}{2} + \frac{7}{4} = \frac{17}{4}$.

Now we analyze the behavior of agents 1 and 2 when they collaborate on the same project and they both have to perform the same task. At time step $t = 0$, we have that the social bias is $\Gamma_1^{\{1,2\}}(0) = \gamma^{q_1(0) - q_2(0)} = 4^0 = 1$. Analogously, we have that $\Gamma_2^{\{1,2\}}(0) = 1$. Therefore, the choice of each agent at time step $t = 0$ is the same as when they perform independently as their personal bias remains unchanged. That is, agent 1 follows the edge $e_\nearrow(v_0^0)$ making progress 1 at time $t = 1$ and agent 2 follows the edge $e_\rightarrow(v_0^0)$ making no progress at time $t = 1$. At time step $t = 1$, agent 1 evaluates the options of following edge $e_\nearrow(v_1^1)$ or $e_\rightarrow(v_1^1)$. Notice that now the social bias of agent 1 is $\Gamma_1^{\{1,2\}}(1) = \gamma^{q_1(1) - q_2(1)} = 4^{1/2} = 2$. Hence, the perceived cost of following edge $e_\nearrow(v_1^1)$ is $\Gamma_1^{\{1,2\}}(1) \cdot b_1 \cdot c(e_\nearrow(v_1^1)) + d(v_2^2, \tau) = 2 \cdot 2 \cdot \frac{7}{4} + 0 = \frac{28}{4}$, whereas the perceived cost of following the edge $e_\rightarrow(v_1^1)$ is $\Gamma_1^{\{1,2\}}(1) \cdot b_1 \cdot c(e_\rightarrow(v_1^1)) + d(v_2^1, \tau) = 2 \cdot 2 \cdot 1 + \frac{7}{4} = \frac{23}{4}$. Therefore, at state v_1^1 agent 1 would follow edge $e_\rightarrow(v_1^1)$. For the last edge, agent 1 has only one option, that is, to follow the edge $e_\nearrow(v_2^1)$ to reach τ. Hence, the cost of the progress path of agent 1 is $\frac{3}{2} + 1 + \frac{7}{4} = \frac{17}{4}$, compared to the cost 13/4 of the progress path that it would follow independently. The progress path of agent 2 does not change when operating with agent 1, as after the first choice to follow edge $e_\rightarrow(v_0^0)$ there are not alternative paths that agent 2 could follow. In conclusion, the total cost of the two agents when operating together is 34/4 compared to the 30/4 when operating independently.

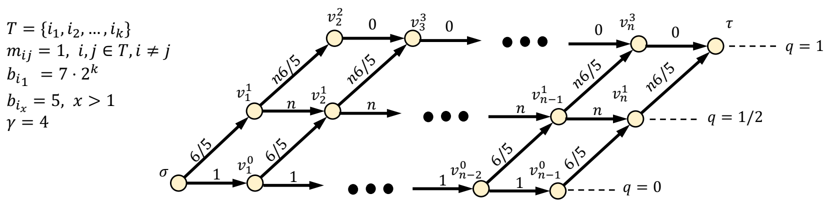

Figure 4 An example where the collaboration leads to an increase, by an $\mathcal{O}(n)$ factor, on the total cost.

Example 2. We now proceed with an example where the collaboration of two agents leads to a decrease to the total cost of their progress paths. Consider the case where two agents $1, 2$ with personal biases $b_1 = 2, b_2 = 6$ operate on the task graph in Figure 1 (right). It can be verified that the progress path of agent 1 is $p(i) = \langle \sigma = v_0^0, v_1^1, v_2^2, v_3^2, v_4^2 = \tau \rangle$ and the progress path of agent 2 is $p(2) = \langle \sigma = v_0^0, v_1^0, v_2^0, v_3^1, v_4^2 = \tau \rangle$. Therefore, the total cost of the $p(1)$ and $p(2)$ is $\frac{36}{5}$.

Now we consider the case where the two agents interact with each other. Similarly to the first example, at time step $t = 0$ the social bias is $\Gamma_1^{\{1,2\}}(0) = \Gamma_2^{\{1,2\}}(0) = 1$ and hence the choices at time $t = 0$ of agents $1, 2$ are the same as in the case where they operate independently. That is, $q_1(1) = 1/2$ and $q_2(1) = 0$. At time step $t = 1$, the social bias of agent 1 is $\Gamma_1^{\{1,2\}}(1) = \gamma^{q_1(1) - q_2(1)} = 4^{1/2} = 2$. According to our model, the perceived cost of following edge $e_{\nearrow}(v_1^1)$ by agent 1 is $\Gamma_1^{\{1,2\}}(1) \cdot b_1 \cdot c(e_{\nearrow}(v_1^1)) + d(v_2^2, \tau) = 2 \cdot 2 \cdot \frac{6}{5} + 0 + 0 = \frac{24}{5}$, whereas the perceived cost of following the edge $e_{\rightarrow}(v_1^1)$ is $\Gamma_1^{\{1,2\}}(1) \cdot b_1 \cdot c(e_{\rightarrow}(v_1^1)) + d(v_2^1, \tau) = 2 \cdot 2 \cdot 1 + \frac{6}{5} + 0 = \frac{26}{5}$. Therefore, at state v_1^1 agent 1 follows edge $e_{\nearrow}(v_1^1)$. We now review the decision of agent 2 at time $t = 1$, whose social bias is $\Gamma_2^{\{1,2\}}(1) = \gamma^{q_2(1) - q_1(1)} = 4^{-1/2} = 1/2$. Hence, agent 2 at time $t = 1$ perceives cost $\Gamma_2^{\{1,2\}}(1) \cdot b_2 \cdot c(e_{\nearrow}(v_1^0)) + d(v_2^1, \tau) = \frac{1}{2} \cdot 6 \cdot \frac{7}{5} + \frac{6}{5} + 0 = \frac{27}{5}$ for following edge $e_{\nearrow}(v_1^0)$, and cost $\Gamma_2^{\{1,2\}}(1) \cdot b_1 \cdot c(e_{\rightarrow}(v_1^0)) + d(v_2^0, \tau) = \frac{1}{2} \cdot 6 \cdot 1 + \frac{7}{5} + \frac{6}{5} = \frac{28}{5}$ for following edge $e_{\rightarrow}(v_1^0)$. Therefore, at state v_1^0 agent 2 follows edge $e_{\nearrow}(v_1^0)$ to reach state v_2^1. At time $t = 2$, we have $q_1(2) = 1, q_2(2) = 1/2$. Agent 1 has no options other than to follow the path $\langle v_2^2, v_3^2, \tau \rangle$ to reach τ. The social bias of agent 2 at time $t = 2$ is $\Gamma_2^{\{1,2\}}(2) = \gamma^{q_2(2) - q_1(2)} = 4^{-1/2} = 1/2$. Hence, the perceived cost of agent 2 at time $t = 2$ in the case of following the edge $e_{\nearrow}(v_2^1)$ is $\Gamma_2^{\{1,2\}}(2) \cdot b_2 \cdot c(e_{\nearrow}(v_2^1)) + d(v_3^2, \tau) = \frac{1}{2} \cdot 6 \cdot \frac{6}{5} + 0 = \frac{18}{5}$, and in the case of following the edge $e_{\rightarrow}(v_2^1)$ is $\Gamma_2^{\{1,2\}}(2) \cdot b_2 \cdot c(e_{\rightarrow}(v_2^1)) + d(v_3^1, \tau) = \frac{1}{2} \cdot 6 \cdot 1 + \frac{6}{5} = \frac{21}{5}$. Therefore, at state v_2^1 agent 2 follows edge $e_{\nearrow}(v_2^1)$. Finally, at time $t = 3$, both agents have no other option than to follow edge (v_3^2, τ) to reach τ. In conclusion, agent 1 follows the progress path $p^{\{1,2\}}(1) = \langle \sigma = v_0^0, v_1^1, v_2^2, v_3^2, v_4^2 = \tau \rangle$ and agent 2 follows that progress path $p^{\{1,2\}}(2) = \langle \sigma = v_0^0, v_1^0, v_2^1, v_3^2, v_4^2 = \tau \rangle$, with total cost $31/5$. That is, if agents $1, 2$ collaborate they decrease the total cost.

Proofs

Detailed proof of Lemma 7. See Figure 4, where all agents operate on the same task graph. For agent i_1 at state v_0^0 the perceived cost is

$$B_{i_1}^T(0) \cdot c(e_{\nearrow}(v_0^0)) + d(v_1^1, \tau) = b_{i_1} \cdot \Gamma_{i_1}^T(0) \cdot \frac{6}{5} + \frac{6n}{5}$$
$$= 7 \cdot 2^k \cdot 4^0 \cdot \frac{6}{5} + \frac{6n}{5} = \frac{42 \cdot 2^k + 6n}{5}$$

for following edge $e_{\nearrow}(v_0^0)$ and

$$B_{i_1}^T(0) \cdot c(e_{\rightarrow}(v_0^0)) + d(v_0^1, \tau) = b_{i_1} \cdot \Gamma_{i_1}^T(0) + \frac{6}{5} + \frac{6n}{5}$$
$$= 7 \cdot 2^k \cdot 4^0 + \frac{6}{5} + \frac{6n}{5} = \frac{35 \cdot 2^k + 6 + 6n}{5}$$

for following the edge $e_{\rightarrow}(v_0^0)$. Hence, agent i_1 follows the edge $e_{\rightarrow}(v_0^0)$.

For agent $i_x, x > 1$, at state v_0^0 the perceived cost is

$$B_{i_x}^T(0) \cdot c(e_{\nearrow}(v_0^0)) + d(v_1^1, \tau) = b_{i_x} \cdot \Gamma_{i_x}^T(0) \cdot \frac{6}{5} + \frac{6n}{5}$$
$$= 5 \cdot 4^0 \cdot \frac{6}{5} + \frac{6n}{5} = \frac{30 + 6n}{5}$$

for following edge $e_{\nearrow}(v_0^0)$ and

$$B_{i_x}^T(0) \cdot c(e_{\rightarrow}(v_0^0)) + d(v_0^1, \tau) = b_{i_x} \cdot \Gamma_{i_x}^T(0) + \frac{6}{5} + \frac{6n}{5}$$
$$= 5 \cdot 4^0 + \frac{6}{5} + \frac{6n}{5} = \frac{31 + 6n}{5}$$

for following the edge $e_{\rightarrow}(v_0^0)$. Hence, agent i_x chooses to follow edge $e_{\nearrow}(v_0^0)$.

At $t = 1$ agent i_1 is at state v_1^0 and all other agents are at state v_1^1.

For agent i_1 at state v_1^0 the perceived cost is

$$B_{i_1}^T(1) \cdot c(e_{\nearrow}(v_1^0)) + d(v_2^1, \tau) = b_{i_1} \cdot \Gamma_{i_1}^T(1) \cdot \frac{6}{5} + \frac{6n}{5}$$
$$= 7 \cdot 2^k \cdot 4^{-(k-1)/2} \cdot \frac{6}{5} + \frac{6n}{5} = \frac{84 + 6n}{5}$$

for following edge $e_{\nearrow}(v_1^0)$ and

$$B_{i_1}^T(1) \cdot c(e_{\rightarrow}(v_1^0)) + d(v_2^0, \tau) = b_{i_1} \cdot \Gamma_{i_1}^T(1) + \frac{6}{5} + \frac{6n}{5}$$
$$= 7 \cdot 2^k \cdot 4^{-(k-1)/2} + \frac{6}{5} + \frac{6n}{5} = \frac{76 + 6n}{5}$$

for following edge $e_{\nearrow}(v_1^0)$. Hence, agent i_1 follows edge $e_{\rightarrow}(v_1^0)$.

For agent $i_x, x > 1$, at state v_1^1 the perceived cost is

$$B_{i_x}^T(1) \cdot c(e_{\nearrow}(v_1^1)) + d(v_2^2, \tau) = b_{i_x} \cdot \Gamma_{i_x}^T(1) \cdot \frac{6n}{5}$$
$$= 5 \cdot 4^{1/2} \cdot \frac{6n}{5} = \frac{60n}{5}$$

for following edge $e_{\nearrow}(v_1^1)$ and

$$B_{i_x}^T(1) \cdot c(e_{\rightarrow}(v_1^1)) + d(v_2^1, \tau) = b_{i_x} \cdot \Gamma_{i_x}^T(1) \cdot n + \frac{6n}{5}$$
$$= 5 \cdot 4^{1/2} \cdot n + \frac{6n}{5} = \frac{56n}{5}$$

for following edge $e_{\rightarrow}(v_1^1)$. Hence, agent i_x chooses to follow edge $e_{\rightarrow}(v_1^1)$.

Notice that after time $t = 2$, all agents continue procrastinating as their perceived cost does not change. Eventually, the progress path $p^T(i_1)$ of agent i_1 costs $\frac{6}{5} + n \cdot (n-2) + \frac{6n}{5}$ and the progress path $p^T(i_x)$ for an agent $i_x, x > 1$, costs $n - 2 + \frac{6}{5} + \frac{6n}{5}$. Hence, the total cost is $\Omega(k \cdot n^2)$.

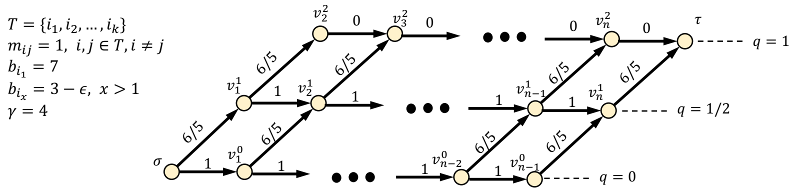

Figure 5 An example where collaboration leads to a decrease, by an $\Omega(n)$ factor, on the total cost.

Notice that in the case where $\Gamma_{i_z}^{T}(t) = 1$ for all $z \in T$, agent i_1 would still follow the same progress path, agent $i_x, x > 1$, would follow the edge $e_{\nearrow}(v_1^1)$ at state v_1^1 as its perceived cost would be:

$$B_{i_x}^{T}(1) \cdot c(e_{\nearrow}(v_1^1)) + d(v_2^2, \tau) = b_{i_x} \frac{6n}{5} = 5 \cdot \frac{6n}{5} = 6n$$

for following edge $e_{\nearrow}(v_1^1)$ and

$$B_{i_x}^{T}(1) \cdot c(e_{\rightarrow}(v_1^1)) + d(v_2^1, \tau) = b_{i_x}\ n\ +\ \frac{6n}{5} = 5 \cdot\ n + \frac{6n}{5} = \frac{31n}{5}$$

for following edge $e_{\rightarrow}(v_1^1)$. Hence, agent i_x would choose to follow edge $e_{\nearrow}(v_1^1)$. In this scenario, agent i_1 would pay cost $n - 2 + \frac{6}{5} + \frac{6n}{5}$ and $i_x, x > 1$, would pay cost $\frac{6}{5} + \frac{6n}{5}$. Hence, the total cost would be $\Omega(k \cdot n)$. Collectively, the total cost of progress paths can increase by a factor $\Omega(n)$. ◀

Proof of Lemma 8. Follows from Claim 5.1 from Gravin et al. [7] ◀

Proof of Lemma 9. See Figure 5. ◀

Proof of Lemma 10. In the case that $b_i = b_j$ the lemma is clearly true, assuming that i and j break ties consistently: the two agents will follow the exact same path. For the rest of the proof, and w.l.o.g., we assume that $b_i > b_j$. Note that for each t we have that $B_i^{\{i\}}(t) = b_i$ and $B_j^{\{j\}}(t) = b_j$. For the sake of leading to a contradiction assume that the statement of the lemma is false. Let v_t^ℓ be the first state that agent j went below agent i, formally, that $(v_t^\ell, v_{t+1}^{\ell+1}) \in p(i)$ and $(v_t^\ell, v_{t+1}^{\ell}) \in p(j)$. Let $p_i = \langle v_t^\ell = v_0, v_1^i, \ldots, v_k^i = \tau \rangle$ be the subpath of $p(i)$ from v_t^ℓ to τ and $p_j = \langle v_t^\ell = v_0, v_1^j, \ldots, v_k^j = \tau \rangle$ be the subpath of $p(j)$ from v_t^ℓ to τ. The perceived cost of the path p_i for i is $b_i \cdot c(v_0, v_1^i) + d(v_1^i, \tau)$. By the definition of $p(i)$ we have that

$$b_i \cdot c(v_0, v_1^i) + d(v_1^i, \tau) \leq b_i \cdot c(v_0, v_1^j) + d(v_1^j, \tau)$$

$$\Rightarrow b_i \cdot (c(v_0, v_1^i) - c(v_0, v_1^j)) + d(v_1^i, \tau) - d(v_1^j, \tau) \leq 0. \quad (1)$$

Similarly, for j we have that

$$b_j \cdot (c(v_0, v_1^j) - c(v_0, v_1^i)) + d(v_1^j, \tau) - d(v_1^i, \tau) \leq 0,$$

$$\Rightarrow b_j \cdot (c(v_0, v_1^j) - c(v_0, v_1^i)) - d(v_1^i, \tau) + d(v_1^j, \tau) \leq 0, \quad (2)$$

By the fact that $b_i > b_j$, and our assumption that when the perceived cost of making progress and procrastinating is equal then the agent chooses to make progress, it follows that at most one of the two inequalities can hold with equality. Summing the two inequalities, we obtain that

$$b_i \cdot (c(v_0, v_1^i) - c(v_0, v_1^j)) + b_j \cdot (c(v_0, v_1^j) - c(v_0, v_1^i)) < 0,$$

which, recalling that $v_0 = v_t^\ell$, that $v_1^i = v_{t+1}^{\ell+1}$, and that $v_1^j = v_{t+1}^\ell$, can be rewritten as

$$b_i \cdot (c(e_\nearrow(v_t^\ell)) - c(e_\rightarrow(v_t^\ell))) + b_j \cdot (c(e_\rightarrow(v_t^\ell)) - c(e_\nearrow(v_t^\ell))) < 0,$$

$$\Rightarrow (b_i - b_j) \cdot (c(e_\nearrow(v_t^\ell)) - c(e_\rightarrow(v_t^\ell))) < 0.$$

But this is a contradiction because we assumed that $b_i > b_j$ and by Property 6 it follows that $c(e_\nearrow(v_t^\ell)) \geq c(e_\rightarrow(v_t^\ell))$. ◀

Proof of Lemma 11. We prove it by contradiction. Let v_t^ℓ be the first node for which there exists an agent i for whom $(v_t^\ell, v_{t+1}^\ell) \in p^T(i)$ and $(v_t^\ell, v_{t+1}^{\ell+1}) \in p(i_1)$. Given that this is the first time that this happens, for each $j \in T$ we have that $q_j(t) \geq q_i(t)$. Note that by Property 2 we have that $\Gamma_i^T(t) \leq 1$, which implies $B_i^T(t) \leq \max\{b_i, 1\} \leq b_{i_1}$.

Arguing as in Lemma 10 we obtain that $(b_{i_1} - B_i^T(t))(c(e_\nearrow(v_t^\ell)) - c(e_\rightarrow(v_t^\ell))) < 0$, leading to a contradiction, which means that $p^T(i) \succeq_G p(i_1)$.

Repeating the argument, and observing that by Property 2 we have that $\Gamma_i^T(t) \geq 1$, giving $B_i^T(t) \geq \max\{b_i \, \Gamma_i^T(t), 1\} \geq b_{i_k}$, we obtain that $p^T(i) \preceq p(i_k)$. ◀

Proof of Lemma 13. First, we show that whenever, at some state v_t^ℓ agent i follows the edge $e_\nearrow(v_t^\ell)$, then i follows the shortest path from v_t^ℓ to σ. That is, $d(v_{t+1}^{\ell+1}, \tau) = d(v_t^\ell, \tau) - c(e_\nearrow(v_t^\ell))$. To prove our claim, we notice that

$$B_i^T(t) \cdot c(e_\nearrow(v_t^\ell)) + d(v_{t+1}^{\ell+1}, \tau) < B_i^T(t) \cdot c(e_\rightarrow(v_t^\ell)) + d(v_{t+1}^\ell, \tau).$$

$$\Rightarrow B_i^T(t) \cdot (c(e_\nearrow(v_t^\ell)) - c(e_\rightarrow(v_t^\ell))) < d(v_{t+1}^\ell, \tau) - d(v_{t+1}^{\ell+1}, \tau).$$

Since $B_i^T(t) \geq 1$ and $c(e_\nearrow(v_t^\ell)) > c(e_\rightarrow(v_t^\ell))$ by Property 6, we have

$$c(e_\nearrow(v_t^\ell)) - c(e_\rightarrow(v_t^\ell)) < d(v_{t+1}^\ell, \tau) - d(v_{t+1}^{\ell+1}, \tau),$$

which proves that $d(v_{t+1}^{\ell+1}, \tau) = d(v_t^\ell, \tau) - c(e_\nearrow(v_t^\ell))$. Hence, each time an agent follows an edge $e_\nearrow(v_t^\ell)$ from any state v_t^ℓ the remaining shortest path to τ decreases by $c(e_\nearrow(v_t^\ell))$ (by Property 5). Notice that this does not imply anything about the behavior of agent i at any other time $t' \neq t$. Our goal is to rely on Property 5 to guarantee that remaining shortest path cannot increase (independently of the followed path) and to bound the additional cost that each agent pays when not decreasing its distance to the target state. By our first claim, the progress path of agent i only increases (compared to the shortest path) when the agent procrastinates, that is, the agent follows an edge $e_\rightarrow(v_t^\ell)$ from some state v_t^ℓ.

Assume that i_{min} is the agent in T with the smallest personal bias $b_{i_{min}}$. Then, by Lemma 11, $p^T(j) \succeq p(i_{min})$, for all $j \in T, j \neq i_{min}$. Let $i \in T$ be any agent $i \neq i_{min}$. For any state v_t^ℓ such that i follows the edge $e_\rightarrow(v_t^\ell)$, agent i_{min} at time t at state $v_t^{\ell'}$ follows either the edge $e_\nearrow(v_t^{\ell'})$ or the edge $e_\rightarrow(v_t^{\ell'})$. By Properties 6 and 12 and the fact that $\ell' \leq \ell$ (since $p^T(i) \succeq p(i_{min})$) we have that $c(e_\nearrow(v_t^{\ell'})) \leq c(e_\rightarrow(v_t^{\ell'})) \leq c(e_\rightarrow(v_t^\ell))$. That is, each time agent i follows an edge $e_\rightarrow(v_t^\ell)$ from some state v_t^ℓ, agent i_{min} follows an edge with larger cost when operating independently. As the total increase in the progress path of agent i, compared to the shortest progress path, is bounded by the number of procrastination edges

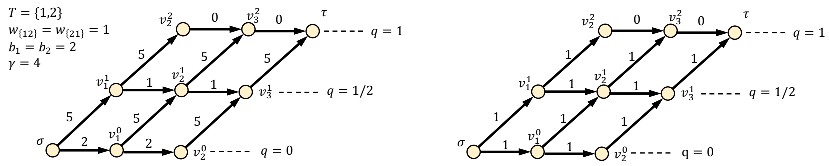

Figure 6 Graphs of type 1 (left) and type 2 (right) that are used in the proof of Theorem 16.

followed by agent i, it follows that the total increase is bounded by the cost of the progress path $p(i_{min})$. Hence, each agent adds at most cost equal to the cost of the progress path $p(i_{min})$. That is, the total cost increases by a factor of at most k. ◀

Proof of Theorem 16. We obtain a reduction from the SetCover problem. In the SetCover problem, we are given a universe of items U and a family of sets $C_1, C_2, \ldots, C_\ell$ and we are asked to find k sets such that each element from the universe is contained in at least one selected set. Our overall strategy is to construct an instance of the OptimalGroupAssignment problem, in polynomial time, from an instance of the SetCover problem. This means, if we can solve the OptimalGroupAssignment problem in polynomial time, then our reduction is a polynomial time algorithm for the SetCover problem, which is known to be **NP**-hard.

Assume that we are given an instance of the SetCover problem with universe U and sets $C_1, C_2, \ldots, C_\ell$. We construct an instance of the OptimalGroupAssignment problem as follows. We include k graphs of type 2 and $\ell + |U| - k$ graphs of type 1, as they are shown in Figure 6. For each element $u \in U$, we include an *element-agent* a_u. For each set $C_i \in \{C_1, C_2, \ldots, C_\ell\}$, we include a *set-agent* a_i. All agents have personal bias $b = 2$, and $\gamma = 2$. Finally, we set $w_{ui} = 1$ if $u \in C_i$, where $C_i \in \{C_1, C_2, \ldots, C_\ell\}$ and $u \in U$, and $w_{ui} = 0$ otherwise. Notice that any agent operating on a type-2 graph always pays cost 2 as it always follows the path $\langle s, v_1^1, v_2^1, v_3^2, t\rangle$. Therefore, the total cost only depends on the agents that operate on type 1 graphs. An agent operating individually on a type-1 graph follows the path $\langle s, v_1^0, v_2^0, v_3^1, t\rangle$, and therefore, pays 14. If for an agent i, that is operating on a type-1 graph, it holds that $w_{ji} = 1$ and agent j is operating on a type-2 graph, agent i follows the path $\langle s, v_1^0, v_2^1, v_3^1, t\rangle$ with cost 13. In the case where $w_{ji} = 0$, for all agents $j \neq i$, or $w_{ji} = 1$ but agent j follows a path containing v_1^0 (i.e., operates on a type 1 graph), it holds that i follows the path $\langle s, v_1^0, v_2^0, v_3^1, t\rangle$ with cost 14.

We show that for any instance of the SetCover problem there exist k sets covering all elements of the universe if and only if there is a solution to the OptimalGroupAssignment problem such that the total cost is $k \cdot 2 + |U| \cdot 13 + (\ell - k)14$. We begin with the first direction, that is, we show that if there exists a solution to the SetCover problem, then there exists an assignment in the OptimalGroupAssignment problem for which the total cost is $k \cdot 2 + |U| \cdot 13 + (\ell - k)14$. For each set C_i in the solution of the SetCover problem, we assign its corresponding set-agent a_i to a type-2 graph (they pay collectively $k \cdot 2$ cost). For the rest set-agents corresponding to sets C_j of the SetCover instance, it holds $w_{ij} = 0$, for all $i \neq j$, and therefore those $(\ell - k)$ agents collectively pay $(\ell - k) \cdot 14$ cost (i.e., they all follow the path $\langle s, v_1^0, v_2^0, v_3^1, t\rangle$ as we explained above). Finally, each element-agent u pays cost $|U| \cdot 13$ in total, as there exists a weight $w_{iu} = 1$, for some $i \neq u$ such that a_i operates on a type 2 graph (as there is a set C_i covering u). This proves the first direction.

Now we prove that if there exists an assignment in the OptimalGroupAssignment problem for which the total cost is $k \cdot 2 + |U| \cdot 13 + (\ell - k)14$ then there is a solution to the SetCover problem. Notice that, independently of which agents operate on graphs of type

2, they always pay cost 2. Every agent a_u, for $u \in U$, pays cost 2 if it operates on a type-2 graph, cost 13 if it operates on a type-1 graph and there exists $w_{iu} = 1$ where a_i operates on a type-2 graph, and 14 otherwise. Note that for each $w_{iu} = 1$, there exists a set C_i covering the element u in the SetCover instance Therefore, the only agents that can pay cost 13 are the element-agents. Since in any solution there are exactly k agents paying cost 2 (i.e., those operating on type-2 graphs) and at least $(\ell - k)$ agents paying cost 14 (i.e., the set-agents that operate on type-2 graphs), the minimum possible total cost is $k \cdot 2 + |U| \cdot 13 + (\ell - k)14$, which is achieved by assigning k set-agents to type-2 graphs, such that for each element-agent u it holds $w_{iu} = 1$ for at least one of the set-agents that were assigned to type-2 graphs. Such an assignment indicates that there exists a set of k set-agents $a_1, a_2, \ldots, a_k$ such that for each $u \in C$, there exists a $1 \le i \le k$ such that $w_{iu} = 1$. These k set-agents correspond to a solution of the SetCover instance. This concludes the proof. ◀