

# Simple approximation algorithms for Polyamorous Scheduling

Yuriy Biktairov\*    Leszek Gąsieniec†    Wanchote Po Jiamjitrak‡  
Namrata†    Benjamin Smith†    Sebastian Wild§†

November 9, 2024

## Abstract

In Polyamorous Scheduling, we are given an edge-weighted graph and must find a periodic schedule of matchings in this graph which minimizes the maximal weighted waiting time between consecutive occurrences of the same edge. This NP-hard problem generalises Bamboo Garden Trimming and is motivated by the need to find schedules of pairwise meetings in a complex social group. We present two different analyses of an approximation algorithm based on the Reduce-Fastest heuristic, from which we obtain first a 6-approximation and then a 5.24-approximation for Polyamorous Scheduling. We also strengthen the extant proof that there is no polynomial-time  $(1 + \delta)$ -approximation algorithm for the Optimisation Polyamorous Scheduling problem for any  $\delta < \frac{1}{12}$  unless  $P = NP$  to the bipartite case. The decision version of Polyamorous Scheduling has a notion of density, similar to that of Pinwheel Scheduling, where problems with density below the threshold are guaranteed to admit a schedule (cf. the recently proven 5/6 conjecture, Kawamura, STOC 2024). We establish the existence of a similar threshold for Polyamorous Scheduling and give the first non-trivial bounds on the poly density threshold.

## 1 Introduction

We prove that simple heuristics from Bamboo Garden Trimming can be generalized to Polyamorous Scheduling and that they yield a constant-factor approximation; we complement this with hardness-of-approximation and structural results about the problem.

Polyamorous Scheduling is a recently introduced class of periodic scheduling problems. *Polyamorous Scheduling* problems each consider a “polycule”, i.e. a set of individuals connected by a set of pairwise relationships, each characterized by a value indicating its neediness, importance, or emotional weight. Their given objective is to devise a periodic schedule of pairwise meetings between couples that minimizes the maximum weighted waiting time between such meetings, subject to the constraint that each individual can meet with at most one person on any given day (see Section 2 for formal definitions).

---

\*University of Southern California, USA, [biktairo@usc.edu](mailto:biktairo@usc.edu)

†University of Liverpool, UK, [{lechu,n.namrata,b.m.smith,wild}@liverpool.ac.uk">{lechu,n.namrata,b.m.smith,wild}@liverpool.ac.uk}](mailto)

‡University of Helsinki, Finland, [wanchote.jiamjitrak@helsinki.fi](mailto:wanchote.jiamjitrak@helsinki.fi)

§University of Marburg, Germany, [wild@informatik.uni-marburg.de](mailto:wild@informatik.uni-marburg.de)

Wanchote Po Jiamjitrak is supported by the Research Council of Finland grant No. 346968.

Namrata and Sebastian Wild are supported by the Engineering and Physical Sciences Research Council of the UK grant EP/X039447/1.

Polyamorous Scheduling (Poly Scheduling in short) was first introduced by Gąsieniec, Smith and Wild in [GSW24]. We first consider the optimisation variant of Poly Scheduling namely, the *Optimisation Polyamorous Scheduling* (OPS) problem where the objective is to find a schedule that minimizes the “heat”, i.e. the worst pain of separation ever felt by any couple. Unlike simpler periodic scheduling problems, such as the Bamboo Garden Trimming (BGT) problem [GKL<sup>+</sup>17], Poly Scheduling is known to be hard to approximate: No poly-time  $(\frac{4}{3} - \varepsilon)$ -approximation exists unless  $P = NP$ .

Prior to this work, the asymptotically best approximation ratio is achieved by the “layering algorithm” from [GSW24], which gives a  $3\lg(\Delta + 1)$ -approximation where  $\Delta$  is the maximal number of partners of any person in the input. We will show in this paper that this layering algorithm is best possible up to constant factors in its class of algorithms, namely all methods that work by scheduling relationships as series of *disjoint* matchings, but that outside of this class, constant-factor approximations are possible. Indeed, we show that a simple generalization of the *Reduce-Fastest* strategy (with  $x = 2.89$ ) achieves a 5.24-approximation.

We also consider *Decision Polyamorous Scheduling* (DPS), where each couple has a maximum duration between meetings and a schedule must either be found or shown not to exist. Pinwheel Scheduling, the well-studied decision variant of Bamboo Garden Trimming, is a special case of DPS where the polycule is a star graph (similarly, BGT is the set of OPS polycules on star graphs). The recent proof of the long-standing  $\frac{5}{6}$ -density conjecture of Pinwheel Scheduling by [Kaw24] admits a strong sufficient criterion to certify feasibility for this class of polycules; prior to this work, this was lacking for Poly Scheduling. More precisely, [GSW24] presents the notion of “poly density” as a generalization of Pinwheel density, but the original definition based on an exponential-size linear program leaves open how to compute the poly density of an instance and hence doesn’t provide an efficiently testable sufficient criterion. In this paper, we show that poly density is always within a constant factor of the largest local Pinwheel density and can hence be well approximated.

A consequence of the above observation is that from the perspective of constant-factor approximations, any OPS instance’s cost is dominated by a single BGT instance embedded in it. This insight was the key ingredient to designing and analyzing a generalization of the Reduce-Fastest approximation [GKL<sup>+</sup>17] for Poly Scheduling.

A well-motivated special case of Poly Scheduling considers a *bipartite* graph of relationships (Bipartite Poly Scheduling). In the original motivation for scheduling meetings of polyamorous people, bipartite polycules correspond to a group of heterosexuals; it also models periodic scheduling scenarios with two explicit classes or hierarchies, such as the members of a sports team who each need to regularly train on a shared set of machines or with a shared set of trainers. Apart from applications, there are strong indications in [GSW24] that bipartite instances may be more well behaved: All examples of infeasible decision polycules where showing infeasibility was nontrivial involve odd cycles. Moreover, the hardness-of-approximation proof in [GSW24] reduces the Chromatic Index problem to unweighted DPS instances; the latter is trivial on bipartite graphs (since they can always be  $\Delta$ -edge coloured). We show that bipartiteness does indeed *not* make the problem easier: It remains NP-hard to find even approximate solutions to Bipartite Poly Scheduling. As a consequence, the hardness of Poly Scheduling is intrinsically linked to the frequencies resp. weights of edges, not simply the embedded edge colouring problem.

## 1.1 Related Work

Poly scheduling problems belong to the broader class of periodic scheduling problems, many of which have attracted considerable attention lately [AdBB<sup>+</sup>22, HvS23, Kaw24, BGL<sup>+</sup>22].

Arguably, the simplest periodic scheduling problem is *Pinwheel Scheduling*, which has been studied extensively since it was introduced in 1989 [HMR<sup>+</sup>89]. In Pinwheel Scheduling we are given  $n$  positive integer frequencies  $f_1 \leq f_2 \leq \dots \leq f_n$ . The goal is to find a Pinwheel schedule, i.e. an infinite sequence of tasks  $1, \dots, n$  such that any contiguous time window of length  $f_i$  contains at least one occurrence of  $i$  or to show that no such schedule exists. The *density* of a Pinwheel Scheduling instance is  $d := \sum_{i=1}^n \frac{1}{f_i}$ . It is easy to observe that  $d \leq 1$  is a necessary condition for an instance to be schedulable, but this is not sufficient; in fact, there exists a threshold  $d^*$  such that  $d \leq d^*$  implies schedulability. A long sequence of works [HMR<sup>+</sup>89, CC93, HRTV92, LL97, CC92, FL02, Din20, GSW22] successively improved bounds on  $d^*$ , culminating in the very recent proof of the best possible bound of  $\frac{5}{6}$  by Kawamura’s [Kaw24] – resolving the long-standing conjecture of Chan and Chin from 1992 [HRTV92]. Generalizations of Pinwheel Scheduling have also been studied, e.g., with jobs of different lengths [FC05, HL92].

We will extensively use the natural optimisation variant of Pinwheel Scheduling, known as the *Bamboo Garden Trimming* (BGT) problem, throughout the paper. In BGT, a garden  $G$  of  $n$  bamboos  $b_1, b_2, \dots, b_n$  with known respective daily growth rates  $g_1 \geq g_2 \geq \dots \geq g_n > 0$  is given. Initially, the height of each bamboo is set to zero. A gardener maintains the garden by trimming each bamboo to height zero according to some schedule. The objective of BGT is to design such a schedule, minimizing the maximum observed height of the tallest ever bamboo, subject to the constraint that the gardener can cut only one bamboo at the end of each day and is not allowed to attend the garden at any other time. This problem was introduced by Gaśieniec, Kalsing, Levcopoulos, Lingas, Min, and Radzik in 2017 [GKL<sup>+</sup>17] and has since attracted a lot of attention. This original paper proved that there exists a natural lower bound of  $H := g_1 + g_2 + \dots + g_n$  on the maximum height of a bamboo. One particularly popular line of research into BGT concerns constant-factor approximation ratios for offline scheduling. Gaśieniec et al. [GKL<sup>+</sup>17] introduced a 2 approximation, but this has since been improved to the current best of  $\frac{4}{3}$  by Kawamura [Kaw24]. For online scheduling, various greedy approaches were taken. Some of the most natural strategies are to trim the highest bamboo, aka, *Reduce-Max* [Als15], to trim the fastest growing bamboo whose height is above a certain threshold, aka, *Reduce-Fastest* [GKL<sup>+</sup>17], and to trim the bamboo above a certain threshold that will soonest achieve the desired height, aka, *Deadline-Driven* [Kus22].

The first approximation ratio to be proved for Reduce-Max was  $\mathcal{O}(\log H)$  [GKL<sup>+</sup>17], but after extensive computer experiments [DDSN19], it was conjectured that Reduce-Max keeps the maximum bamboo height within  $\mathcal{O}(H)$ . In fact, the authors conjectured a much stronger bound of  $2H$  for Reduce-Max. However, this cannot be arbitrarily low. The authors of [GJK<sup>+</sup>24] constructed an input BGT instance using which they proved that the approximation ratio of Reduce-Max cannot be less than  $9/8$ . On the other hand, Bilò et al. [BGL<sup>+</sup>22] proved a bound of  $9H$  on the maximum bamboo height under the Reduce-Max algorithm, which was later improved by Kuszmaul [Kus22] to  $4H$  – the best bound to be proved so far.

Another popular algorithm is Reduce-Fastest( $x$ ), in which the fastest-growing bamboo with a height of at least  $xH$  is trimmed each day. The bounds known for Reduce-Fastest are a bit better than that of Reduce-Max. The first approximation ratio of 4 was proved for Reduce-Fastest(2) in [GKL<sup>+</sup>17]. The same paper [DDSN19] that conjectured the  $2H$  bound of Reduce-Max, also conjectured that Reduce-Fastest(2) and Reduce-Fastest(1) achieve the maximum height of  $3H$  and  $2H$ , respectively. Subsequently, Bilò et al. [BGL<sup>+</sup>22] proved that for  $x = 1 + 1/\sqrt{5} \approx 1.45$ , the maximum bamboo height achieved is within the approximation ratio of  $\frac{(3+\sqrt{5})}{2} \approx 2.62$ . Kuszmaul [Kus22] showed recently that for any  $x \geq 2$ , Reduce-Fastest( $x$ ) keeps all bamboos below  $(x+1)H$ , thereby proving one of the conjectures from earlier [DDSN19]. Kuszmaul also proved that no matter the value of  $x$ , Reduce-Fastest( $x$ ) does not give a bound

better than  $(2.01)H$ , thus disproving the conjecture that Reduce-Fastest(1) keeps all bamboos below  $2H$ .

Kusmaul [Kus22] also proved an approximation ratio of 2 for the simple Deadline-Driven algorithm. The algorithm was first introduced by Liu and Layland [LL73] in the context of a related scheduling problem in the early 1970s.

The Bamboo Garden Trimming problem is closely related to *Cup Games*, which share the mechanic of growth and cutting and are used to model similar applications [BFCK19, Kus20]. Single processor cup games follow a similar daily routine in which an adversarial filler distributes 1 unit of water freely between a set of cups, then an emptier selects a cup and removes up to 1 unit of water from it. Multi-processor versions are also of interest, relating to the multi-gardener version of Bamboo Garden Trimming [Kus22]. As with Bamboo Garden Trimming, greedy algorithms perform well.

Further periodic scheduling problems with less direct connections to Poly Scheduling have also been studied. *Patrolling problems* typically involve periodic schedules: for example, [LL73] finds schedules for a fleet of  $n$  identical robots which patrol points in a metric space; the *Continuous BGT Problem* [GJK<sup>+</sup>24] sends a single robot to points with different frequencies requirements; the *Point Patrolling Problem* studied in [KS20] assigns one of  $n$  workers a single, daily recurring task subject to the constraint that each worker  $i$  requires a break of  $a_i$  days between work days; the *Replenishment Problems with Fixed Turnover Times*, studied in [BvEJ<sup>+</sup>22], considers vertices in a graph which must be visited with given frequencies, constrained by the goal of minimizing the length of the tour used to visit them.

## 1.2 Our Results

Despite being an interesting combinatorial optimisation problem, Poly Scheduling has only been formally studied in [GSW24]. This paper introduced several approximations, the notion of poly density, and the first inapproximability result for any periodic scheduling problem; we introduce better approximations, bound and exploit poly density, and extend this inapproximability result to the bipartite case.

An OPS instance  $\mathcal{P}_o = (P, R, g)$  consists of an undirected graph  $(P, R)$  of people  $P$  connected by relationships  $R$ , along with a desire growth rate  $g_e : R \mapsto \mathbb{R}_{>0}$  for each relationship  $e \in R$ . The *current heat* of an edge is the total desire grown since the last time it was scheduled; every day, the scheduler must choose a set of all-day meetings such that each person  $v \in P$  is not double-booked, with the goal of minimising the highest heat ever experienced.

In Section 3, we define a simple generalization of the Reduce-Fastest( $x$ ) algorithm introduced by [GKL<sup>+</sup>17]: Broadly speaking, Reduce-Fastest( $x$ ) is an online scheduling method which chooses a matching to schedule each day; the matching with the fastest-growing edges whose current heat is above some user-selected threshold  $x$ . In Section 3.2.2, we introduce a simple analysis that shows that Reduce-Fastest(4) gives a 6-approximation, while more in-depth analysis shows that Reduce-Fastest( $2 + \frac{2\sqrt{5}}{5} \approx 2.89$ ) gives a  $3 + \sqrt{5} \approx 5.24$  approximation. These both substantially improve on the  $O(\log \Delta)$  approximation ratio in [GSW24]. Formal definitions of OPS and Reduce-Fastest( $x$ ) for OPS are introduced in Section 2.

When looking for exact solutions for certain OPS instances, it is well established that schedules with exponentially long periods are inevitable, and whether any feasible instance of Pinwheel Scheduling admits a succinctly encodable schedule remains an open problem. We show that the same is not true for *approximate* solutions. In Section 3.3, we show that for any arbitrary OPS schedule, we can construct a new periodic schedule with polynomial length and a maximum heat of at most four times that of the original schedule; we thus never have to deal with superpolynomial schedules for constant-factor approximations.

Section 4.1 revisits the idea of poly density introduced by [GSW24] and shows that it is narrowly sandwiched by the maximum personal growth rate of any person in the polycule:  $G^* = \max_{v \in P} \sum_{e \in R: v \in e} g(e)$ .

**Theorem 1.1 (Poly Density Approximation):** *For an OPS instance with poly density  $\bar{h}^*$  and maximum personal growth rate  $G^*$ , we have  $G^* \leq \bar{h}^* < \frac{3}{2}G^*$ .*  $\triangleleft$

We then turn our attention to the poly density of Decision Polyamorous Scheduling. Each instance of DPS, or *DPS polycule*  $\mathcal{P}_d = (P, R, f)$  consists of a set  $P$  of people, a set  $R$  of relationships, and a set  $f$  of “meetup frequencies”  $f_e$ . The goal is to either find a schedule of pairwise meetings such that each couple  $e = \{i, j\}$  meets at least every  $f_e$  days or to report that no such schedule exists (subject to the constraint that each person can attend at most one meeting per day). A DPS polycule can naturally be represented as a graph of people with the edges representing their relationships; since each person can attend at most one meeting per day, the edges scheduled on any given day must form a matching in this graph.

Section 4.2 shows that the poly density of DPS generalizes the density of Pinwheel Scheduling and establishes a poly-density threshold.

**Theorem 1.2 (Density Threshold):** *Any DPS instance  $\mathcal{P}_d = (P, R, f)$  with poly density  $d \leq \frac{1}{4}$  is schedulable.*  $\triangleleft$

Finally, in Section 5 we strengthen the inapproximability result introduced by [GSW24, Theorem 1.3] by proving that it holds even in the bipartite case:

**Theorem 1.3 (SAT Hardness of approximation):** *Unless  $P = NP$ , there is no polynomial-time  $(1 + \delta)$ -approximation algorithm for the bipartite Optimisation Poly Scheduling problem for any  $\delta < \frac{1}{12}$ .*  $\triangleleft$

We prove this by constricting a bipartite DPS polycule  $\mathcal{P}_{d\varphi}$  from an arbitrary 3SAT instance  $\varphi$  such that  $\mathcal{P}_{d\varphi}$  is schedulable iff  $\varphi$  is satisfiable.

## 2 Preliminaries

This section introduces some general notations which will be used throughout the paper.

We write  $[m, n]$  for  $\{m, m+1, \dots, n\}$  and  $[n]$  for  $[1, n]$ . For a set  $\mathcal{A}$ , we denote its powerset by  $2^{\mathcal{A}}$ . All graphs  $(V, E)$  are simple and undirected, and the *maximum* degree in  $(V, E)$  is denoted by  $\Delta(V, E)$ . An edge colouring of a graph is an assignment of colours to its edges so that no two adjacent edges are assigned the same colour. The chromatic index of a graph  $(V, E)$ , denoted  $\chi_1(V, E)$ , is the smallest number of colours with which  $(V, E)$  can be edge coloured. Vizing [Viz64] has shown that for any graph  $(V, E)$ , the chromatic index  $\chi_1(V, E)$  is either its maximum degree  $\Delta$ , or else is  $\Delta + 1$ . We represent the set of vertices in the neighbourhood of  $v \in V$  by  $\mathcal{N}(v)$ .

We now give formal definitions for the optimisation and decision versions of the Polyamorous Scheduling problem, as introduced by [GSW24]. For a set of edges  $R$ , an (infinite) *schedule* is a function assigning days to subsets of edges:  $S : \mathbb{N}_0 \rightarrow 2^R$ . Given a schedule  $S$  and an edge  $e \in R$ , we define the *recurrence time*  $r(e) = r_S(e)$  of  $e$  in  $S$  as the maximal time between consecutive occurrences of  $e$  in  $S$ , formally

$$r_S(e) := \sup_{d \in \mathbb{N}} \begin{cases} d+1 & \exists t \in \mathbb{N}_0 : e \notin S(t) \cup S(t+1) \cup \dots \cup S(t+d-1); \\ 1 & \text{otherwise.} \end{cases} \quad (1)$$

(We can allow  $r_S(e) = \infty$  if  $S$  leaves growing gaps between consecutive occurrences of  $e$ , but all schedules in the following will have finite recurrence times for all edges.)

**Definition 2.1 (Optimisation Polyamorous Scheduling):** An OPS instance  $\mathcal{P}_o = (P, R, g)$  consists of an undirected graph  $(P, R)$ , where the vertices  $P := \{p_1, \dots, p_n\}$  are  $n$  persons and the edges  $R$  are pairwise relationships, along with a desire growth rate  $g : R \rightarrow \mathbb{R}_{>0}$  for each relationship in  $R$ . A schedule  $S : \mathbb{N}_0 \rightarrow 2^R$  is valid if, for all days,  $t \in \mathbb{N}_0$ ,  $S(t)$  is a matching in  $\mathcal{P}_o$ . The goal is to find a valid schedule that minimizes the heat  $h := h(S) := \max_{e \in R} g(e) \cdot r_S(e)$ .  $\triangleleft$

**Definition 2.2 (Decision Polyamorous Scheduling):** A DPS instance  $\mathcal{P}_d = (P, R, f)$  consists of an undirected graph  $(P, R)$  with integer frequencies  $f : R \rightarrow \mathbb{N}$  for each relationship.

The goal is to find a schedule  $S : \mathbb{N}_0 \rightarrow 2^R$ , such that

- (a) (no conflicts) for all days  $t \in \mathbb{N}_0$ ,  $S(t)$  is a matching in  $\mathcal{P}_d$ , and
- (b) (frequencies) for all  $e \in R$ , we have  $r_S(e) \leq f(e)$ ,

or to report that no such schedule exists. In the latter case,  $\mathcal{P}_d$  is called infeasible.  $\triangleleft$

We write  $f_e$  as a shorthand for  $f(\{p_i, p_j\})$  resp.  $f(e)$ . It is easy to show that an infinite schedule exists if and only if a periodic schedule exists, i.e. a schedule where there is a  $T \in \mathbb{N}$  such that for all  $t$ , we have  $S(t) = S(t + T)$ .

It will be convenient for much of the following to use *Normalised* OPS instances, which we define in the following way: Considering an OPS polycule  $\mathcal{P}_o = (P, R, g)$ , define the personal growth rate  $G_v$  for person  $v \in P$  as  $G_v := \sum_{e \in R: v \in e} g(e)$  and the maximum personal growth rate across all  $v \in P$  as  $G^* := \max_{v \in P} G_v$ .  $\mathcal{P}_o$  is normalised, or in normal form, if  $G^* = 1$ . Any OPS polycule  $\mathcal{P}_o$  can be trivially normalised by dividing each growth rate  $g_e \in g$  by  $G^*$ , and restored by inverting this.

The following pair of lemmas are taken from [GSW24]; the first shows how to reduce OPS to DPS, while the second shows how DPS solves OPS.

**Lemma 2.3 (OPS to DPS, [GSW24]):** For every combination of OPS instance  $\mathcal{P}_o = (P, R, g)$  and heat value  $h$ , there exists a DPS instance  $\mathcal{P}_d = (P, R, f)$  such that

- (a) any feasible schedule  $S : \mathbb{N}_0 \rightarrow 2^R$  for  $\mathcal{P}_d$  is a schedule for  $\mathcal{P}_o$  with heat  $\leq h$ , and
- (b) any schedule  $S'$  for  $\mathcal{P}_o$  with heat  $h' > h$  is not feasible for  $\mathcal{P}_d$ .  $\triangleleft$

**Lemma 2.4 (DPS to OPS, [GSW24]):** Let  $\mathcal{P}_d = (P, R, f)$  be a DPS instance. Set  $F := \max_{e \in R} f(e)$ . There is an OPS instance  $\mathcal{P}_o = (P, R, g)$  such that the following holds.

- (a) If  $\mathcal{P}_d$  is feasible, then  $\mathcal{P}_o$  admits a schedule of heat  $h \leq 1$ .
- (b) If  $\mathcal{P}_d$  is infeasible, then the optimal heat  $h^*$  of  $\mathcal{P}_o$  satisfies  $h^* \geq (F + 1)/F$ .  $\triangleleft$

### 3 Approximating Optimisation Polyamorous Scheduling

This section will examine approximation algorithms for OPS in 3 ways: Section 3.1 considers several lower bounds on OPS, Section 3.2 evaluates reduce-fastest as an approximation algorithm for OPS, and Section 3.3 introduces a method for constructing polynomial schedules for OPS and BGT.

As OPS generalises BGT, with instances on star-graphs being fully equivalent, these problems share both features and algorithms. Particularly of interest here is the *Reduce-Fastest* heuristic, which we generalise here and will use in multiple parts of this section:

**Definition 3.1 (Reduce-Fastest( $x$ ) for OPS):** Begin with an OPS instance  $\mathcal{P}_o = (P, R, g)$  and relabel to sort edges  $e \in R$  by decreasing growth rates  $g(e)$ , breaking ties arbitrarily.

Construct an online schedule using the fixed parameter  $x$  by scheduling a greedy matching  $M$  each day, built the following way:

For all edges  $e = \{u, v\} \in R$  in order of decreasing growth rate, add  $e$  to  $M$  if and only if

- the current heat of  $e$  is at least  $x \cdot G^*$ , and
- no edge incident at  $u$  or  $v$  is already in  $M$ .

◁

Note that on a star graph, each matching  $M$  will contain a single edge, so our OPS Reduce-Fastest behaves like the BGT Reduce-Fastest heuristic introduced by [GKL<sup>+</sup>17] on such instances.

### 3.1 Lower Bounds

We will begin by defining Disjoint Matching algorithms and demonstrating them to be a dead end in Section 3.1.1, then show a simple but general lower bound from the local BGT instance in Section 3.1.2, then a lower bound for the Reduce-Fastest( $x$ ) algorithm in Section 3.1.3.

#### 3.1.1 Disjoint Matching Algorithms

In addition to introducing OPS, [GJK<sup>+</sup>24] introduced multiple proofs that OPS is NP-hard, several lower bounds for it, and the *Layering Algorithm* – an  $\mathcal{O}(\log \Delta)$ -approximation for OPS.

This algorithm partitions edges into layers based on their growth rates, uses a fixed round-robin schedule for each layer, and then interleaves these schedules into another fixed round-robin schedule. Further details are not needed here, only that each edge only appears in a single matching, and thus, all matchings in schedules produced by the Layering Algorithm are *disjoint*.

We define Disjoint Matching algorithms as algorithms whose schedules contain only disjoint matchings, and will show that such algorithms obey the following:


**Theorem 3.2:** The approximation ratio of any Disjoint Matching algorithm for OPS is  $\Omega(\log n)$ .

◁

**Proof:** Consider an OPS instance  $\mathcal{P}_o^* = (P, R, g)$  consisting of  $d$  disjoint stars  $*_1, \dots, *_i, \dots, *_d$ , such that star  $i$  has  $i$  edges and  $g = \frac{1}{i}$  for all edges in star  $i$ . The first four such stars are shown in Figure 1. The optimal heat for  $\mathcal{P}_o^*$  is 1, and can be obtained by scheduling each star independently in a round-robin fashion.

Any set of disjoint matchings  $\mathcal{M}$  which covers all edges  $e \in R$  will consist of at least  $d$  matchings due to the largest star  $*_d$ . Let  $g_M^*$  be the maximum growth rate of any edge in matching  $M_i \in \mathcal{M}$ , i.e.  $g_i^* := \max_{e \in M_i} g(e)$  and  $g^* = \max_{i \in \mathcal{M}} g_i^*$ . In the given disjoint stars, there is at least one matching with  $g_i^* = 1$ , at least two matchings with  $g_i^* \geq 1/2$ , at least three with  $g_i^* \geq 1/3$ , and so on. Once we commit to scheduling the matchings as atomic objects, the problem has become equivalent to a BGT instance with  $d$  bamboos of growth rates  $g_1^*, \dots, g_d^*$ . The optimal height for any schedule of this instance is at least  $\sum_{i=1}^d g_i^* \geq \sum_{i=1}^{\Delta} 1/i \sim \ln \Delta$ . Since  $n = \sum_{i=1}^{\Delta} 2i - 1 = \Delta^2$ , any such schedule for  $\mathcal{P}_o^*$  thus has heat  $\Omega(\log \Delta) = \Omega(\log n)$ , compared to the optimal heat of 1.  $\square$

This means that, in its “league” of algorithms that first divide the graph into a set of disjoint matchings and then schedule these matchings as atomic objects, the Layering Algorithm’s approximation ratio is best possible up to constant factors. To do much better than this, different kinds of approach will be needed



**Figure 1:** An instance of disjoint set of four stars that give the worst-case lower bound of  $\Omega(\log n)$  for the maximum heat for an OPS instance when disjoint matching is scheduled.

### 3.1.2 Lower bound due to the local BGT instance

Recall the BGT problem introduced in Section 1.1, and that BGT is the special case of OPS on star graphs. Thus, the largest lower bound obtained from the embedded BGT instances in the OPS instance will be a natural lower bound for OPS. More formally, consider a vertex  $v \in P$ , recalling that the personal growth rate  $G_v$  for person  $v \in P$  as  $G_v := \sum_{e \in R: v \in e} g(e)$  and the maximum personal growth rate across all  $v \in P$  as  $G^* := \max_{v \in P} G_v$ . We claim that  $G^*$  is a natural lower bound on the maximum heat of any schedule of an OPS instance.

**Theorem 3.3:** *For any OPS instance  $\mathcal{P}_o = (P, R, g)$ ,  $G^*$  is a lower bound on the maximum heat of any schedule for  $\mathcal{P}_o$ .*  $\triangleleft$

**Proof:** This can be easily proved by contradiction. Let  $v^* = \arg \max_{v \in P} \sum_{e: v \in e} g(e)$  be the vertex that corresponds to  $G^*$ . Assume that there exists a periodic schedule that keeps the heat of all edges incident on  $v^*$  to some heat  $h < G^*$ . In the neighbourhood of  $v^*$ ,  $\mathcal{N}(v^*)$ , we can only schedule one edge on any given day. After  $\lfloor \frac{|N(v^*)| \cdot h}{G^* - h} \rfloor + 1$  days, the heat of at least one edge must be greater than  $h$ , a contradiction.  $\square$

### 3.1.3 Lower bound for Reduce-Fastest

The next section will analyse the approximation ratio of the Reduce-Fastest heuristic, but first, we will demonstrate a lower bound for this heuristic. The proof concerns a fully connected OPS polycule which is nominally fair (all growth rates are the same), but which is labelled adversarially with respect to a single unfortunate individual.

**Theorem 3.4:** *When applied to normalised OPS instances  $\mathcal{P}_o = (P, R, g)$ , Reduce-Fastest( $x$ ) with  $x \geq 2$  can produce schedules with heat  $h \geq x + 2 - \frac{1}{|P|-1}$ .*  $\triangleleft$

**Proof:** For some odd  $n$ , consider a normalised OPS instance  $\mathcal{P}_o = (P, R, g)$  where  $(P, R)$  is the complete graph  $K_n$  on  $n$  vertices, and has uniform growth rates of  $\frac{1}{n-1}$ . Reduce-Fastest( $x$ ) will not schedule edges until they reach the threshold heat  $x$ , and will then schedule edges according to growth rate (all tied), then to an arbitrary choice of their ordering to break ties. To construct an adversarial ordering of edges, we will consider two subgraphs: the arbitrary induced subgraph  $K_{n-1}$  of  $P$ , and the subgraph of edges in  $K_n$ , but not  $K_{n-1}$ , called  $K_n - R(K_{n-1})$ .

Since  $n$  is odd,  $n-1$  is even, and  $K_{n-1}$  is 1-factorizable, i.e. edges of  $K_{n-1}$  can be decomposed into  $n-2$  disjoint matchings, each covering all vertices in  $K_{n-1}$ . As only one vertex appears in  $K_n$  but not  $K_{n-1}$ ,  $K_n - R(K_{n-1})$  requires a further  $n-1$  matchings to cover, as each matching can only contain a single edge from  $K_n - R(K_{n-1})$ .

So, for some initial ordering of edges  $e \in R$ , Reduce-Fastest( $x$ ) will first wait for all edges to grow to heat  $x$ , then schedule matchings in  $K_{n-1}$ , then matchings covering  $K_n - R(K_{n-1})$ . This means that the last edge is scheduled  $(n-2 + n-1) = 2n-3$  days later. This last edge thus starts with heat at least  $x$  and grows a further heat of  $(2n-3) \cdot \frac{1}{n-1} = 2 - \frac{1}{n-1}$  before being scheduled.



Note that in this  $2n - 3$  day period, the heat of each edge grows by  $\frac{1}{n-1}$  per day, so no previously scheduled edge will reach the threshold of  $x \geq 2$ , and thus no edge will interfere with the scheduling sequence.  $\square$

**Remark 3.5:** Note that if Reduce-Fastest is applied to the above graph with  $x < 2$ , the maximum heat will tend to infinity for an appropriately large  $n$ . This is because the edges which are scheduled first will cross the threshold  $x$  for a second time before the last edge is scheduled, causing this last edge to grow indefinitely.  $\triangleleft$

## 3.2 Reduce-Fastest

This section gives two analyses for the approximation ratio of the OPS version of Reduce-Fastest ( $x$ ), defined by Definition 3.1.

Both analyses use a notion of *blocking*: Each day, Reduce-Fastest constructs a matching  $M$  by greedily selecting edges with heat at least  $G^*x$  according to their growth rates (with ties broken by some initial ordering). This means that for some edge  $e$  with heat  $h \geq G^*x$  not to be added to  $M$ , a higher priority edge containing one of the vertices in  $e$  must block  $e$  by being added to  $M$  before  $e$  was considered.

### 3.2.1 6-Approximation Algorithm

In this section, we prove that Reduce-Fastest(4) is a 6-approximation for OPS by proving Theorem 3.6. This analysis is inspired by the similar proof of an  $(x+1)$ -approximation algorithm for the BGT problem in [Kus22].

**Theorem 3.6:** *Given a normalised OPS instance  $\mathcal{P}_o = (P, R, g)$ ; for all  $x \geq 4$ , the maximum heat achieved by an edge in Reduce-Fastest( $x$ ) is strictly less than  $x + 2$ .*  $\triangleleft$

**Proof:** Suppose for the sake of contradiction that some edge  $e_i \in R$  achieves heat of at least  $x + 2$  at time  $t_2$  after most recently reaching heat  $x$  at time  $t_1$ . This implies that on every day in the interval  $[t_1, t_2)$ ,  $e_i$  was blocked by a higher-priority adjacent edge. Consider the sequence of these  $t_2 - t_1$  blocking edges and denote the set of distinct edges in this sequence by  $S$ .  $e_i$  could be blocked by two separate higher-priority edges at the same time – in this case, arbitrarily select one to include in the sequence. For all  $e_j \in S$ , we denote by  $m_j$  the number of times that  $e_j$  appears in the sequence. For an edge to appear  $m_j$  times in our sequence, its growth rate must follow:

**Claim 3.7:** For all  $e_j \in S$ , we have  $g(e_j) > m_j \cdot g(e_i)$ .  $\triangleleft$

**Proof:** Begin by considering the case of  $m_j = 1$ , which simply implies that  $g(e_j) \geq g(e_i)$ . This follows from the definition of Reduce-Fastest( $x$ ) (Definition 3.1) – the edge with faster heat growth rate is scheduled earlier and hence blocks the edge  $e_i$ .

Given that edge  $e_j$  can have any initial heat, it could be scheduled as early as  $t_1$ . After this, it must grow to heat  $x$  before each time it is scheduled again,  $m_j - 1$  times total. Thus, to achieve  $m_j \geq 2$ , edge  $e_j$  has to grow a total heat of at least  $x(m_j - 1)$  in the interval  $(t_1, t_2)$ . In the same time interval  $\mathcal{T}$ ,  $e_i$  must grow less than 2 units of heat, as it starts with heat  $x$  and doesn't reach a heat of  $x + 2$  until  $t_2$ . Thus

$$\frac{2}{g(e_i)} > \mathcal{T} \geq \frac{x(m_j - 1)}{g(e_j)}$$

and

$$g(e_j) > x(m_j - 1) \frac{g(e_i)}{2} > 2(m_j - 1)g(e_i) > m_j g(e_i),$$

for  $x \geq 4$  and  $m_j \geq 2$ . This completes the proof of the claim.  $\square$

Using Claim 3.7, we can argue from the number of blockers:

$$t_2 - t_1 \leq \sum_{e_j \in S} m_j < \sum_{e_j \in S} \frac{g(e_j)}{g(e_i)} = \frac{1}{g(e_i)} \sum_{e_j \in S} g(e_j) \leq \frac{2}{g(e_i)}(1 - g(e_i)),$$

where the last inequality is a product of normalisation – since  $G^* = 1$ , for each edge  $e_i = \{u, v\}$ , we have  $G_u \leq 1$  and  $G_v \leq 1$ .

Since the heat of  $e_i$  is less than  $x + g(e_i)$  at  $t_1$ , the heat level reached by  $t_2$  is strictly less than

$$x + g(e_i) + (t_2 - t_1)g(e_i) \leq x + g(e_i) \left(1 + \frac{2}{g(e_i)}(1 - g(e_i))\right) = x + 2 - g(e_i) < x + 2,$$

which is a contradiction.  $\square$

**Corollary 3.8:** *Reduce-Fastest(4) is a 6-approximation of OPS.*  $\triangleleft$

From Theorems 3.4 and 3.6, the following simple corollary follows.

**Corollary 3.9:** *Given an OPS instance  $\mathcal{P}_o = (P, R, g)$ , the  $x + 2$  bound on the maximum heat achieved by Reduce-Fastest( $x$ ) is tight for  $x \geq 4$  and large  $n$ .*  $\triangleleft$

### 3.2.2 5.24-Approximation Algorithm

In this section, we prove our best bound on Reduce-Fastest: Reduce-Fastest(2.89) is a 5.24 approximation for OPS. The analysis follows the related proof of approximation algorithm for the BGT problem in [BGL<sup>+</sup>22].

**Theorem 3.10:** *Given a normalised OPS instance  $\mathcal{P}_o = (P, R, g)$ ; for all  $x > 2$ , the maximum heat achieved by an edge in Reduce-Fastest( $x$ ) is bounded by*

$$\max \left\{ x + \frac{x^2}{4(x-2)}, x + \frac{1}{2} + \frac{x^2}{4(x-1)} \right\}. \quad \triangleleft$$

**Proof:** Let  $S_{\text{rf}}$  be the schedule produced when Reduce-Fastest( $x$ ) is applied to  $\mathcal{P}_o$ ,  $h = h(S_{\text{rf}})$  be the heat of this schedule, and  $e_i = \{u, v\}$  be an edge which reaches heat  $h$  on day  $T$ . Before reaching day  $T$ , the heat of  $e_j$  has grown for some period without being scheduled, reaching threshold  $x$  on day 0 in this period. As  $e_i$  was not scheduled in the interval  $[0, T-1]$ , for each day in this interval, there must have been some higher-priority edge adjacent to  $e_i$  which was scheduled instead; construct the *blocking sequence*  $S$  from these  $T$  blocking edges. On any day which has two separate blocking edges, arbitrarily choose one to add to  $S$  and add the other to the *suppressed sequence*  $Q$ . Let  $N$  be the number of distinct edges in  $S$ .

Let the *volume*  $V$  be the overall growth of  $e_i$  and its adjacent edges in the days of the interval  $[0, T-1]$ . Because of normalisation,  $G_u \leq 1$  and  $G_v \leq 1$ , so the heat of this set of edges grows by at most  $2 - g(e_i)$  per day, and  $V \leq T(2 - g(e_i))$ .

For each edge  $e_j \in S$ , call the first time it was scheduled in  $S$  the *first-time scheduling* and further times *repeated schedulings*. The volume of heat associated with edge  $e_j$  is either removed by the first-time scheduling, removed by repeated scheduling, or is *leftover volume* – heat that remains after day  $T-1$ .

In order to bound the amount of volume  $V_r$  removed by repeated schedulings, we first calculate the volume  $V_f$  that is accumulated before the first time each edge is scheduled in  $S$ , then the volume  $V_l$  which is leftover after the last time each edge is scheduled in  $S$ . Notice that for each edge  $e_j \in S$ , it holds that  $g(e_j) \geq g(e_i)$ . If  $e_j$  appears in  $S$  for the first time on day  $d$ , then the volume removed on day  $d$  in addition to the total volume removed by the preceding schedulings of  $e_j$  in  $Q$  will be at least  $(d+1)g(e_j) \geq (d+1)g(e_i)$ . This is because  $(d+1)g(e_j)$  is the total growth of  $e_j$  up to and including the day  $d$ , all of which will be removed when  $e_j$  is scheduled. This implies that the amount of volume removed by all first-time schedulings in addition to the preceding schedulings of the corresponding edges in the suppressed sequence  $Q$  will be given by  $V_f \geq \sum_{j=1}^N jg(e_i) = \frac{N(N+1)}{2} \cdot g(e_i)$ , since exactly one edge is added to the sequence  $S$  per day.

Moreover, if  $e_j$  appears in the sequence  $S$  for its last time at day  $T-1-d$ , then the leftover volume of  $e_j$  in addition to the volume removed by the following schedulings of  $e_j$  in  $Q$  is at least  $d' \cdot g(e_i)$ . Therefore, for all  $N$  edges, the total volume left after the last scheduling of  $e_j$  in  $S$  will be  $V_l = \sum_{j=1}^N (j-1)g(e_i) = \frac{N(N-1)}{2} \cdot g(e_i)$ . Finally, the edge  $e_i$  is not scheduled in the interval  $[0, T-1]$  but does grow by exactly  $Tg(e_i)$  during this time.

We can then bound the total volume removed by repeated schedulings of edges in  $S$  as

$$\begin{aligned} V_r &\leq V - \left( \frac{N(N+1)}{2} + \frac{N(N-1)}{2} \right) g(e_i) - Tg(e_i) \\ &= V - N^2 \cdot g(e_i) - Tg(e_i) \\ &\leq T(2 - g(e_i)) - N^2 g(e_i) - Tg(e_i) \\ &= 2T(1 - g(e_i)) - N^2 g(e_i). \end{aligned}$$

Each of these removes heat at least  $x$ , so the number  $N_r$  of repeated schedulings is upper bounded by  $\frac{V_r}{x} = \frac{2T(1-g(e_i)) - N^2 g(e_i)}{x}$ , which gives us the following bound on the total number of elements in the sequence

$$|S| = T = N + N_r \leq N + \frac{2T(1 - g(e_i)) - N^2 g(e_i)}{x},$$

which, with  $x \geq 2$ , implies

$$T \leq \frac{Nx - N^2 g(e_i)}{2g(e_i) + x - 2}$$

Consider the minimum value of the right-hand side; as  $g(e_i)$  and  $x$  are fixed, we can view it as a function of  $N$ :  $R(N) = (Nx - N^2 g(e_i)) / (2g(e_i) + x - 2)$ . The function  $R(N)$  is a concave downward parabola that attains its maximum at  $N = \frac{x}{2g(e_i)}$ , giving us

$$T \leq R\left(\frac{x}{2g(e_i)}\right) = \frac{x^2}{4g(e_i)(2g(e_i) + x - 2)}$$

Using this upper bound to  $T$ , we now bound the overall growth of the edge  $e_i$ , i.e. the maximum heat  $h(S_{\text{rf}})$ . At day  $d = 0$ ,  $e_i$  has heat at most  $x + g(e_i)$  by our choice of the time interval, and in the next  $T$  days it grows by  $Tg(e_i)$ . Hence:

$$h(S_{\text{rf}}) \leq x + g(e_i) + Tg(e_i) \leq x + g(e_i) + \frac{x^2}{4(2g(e_i) + x - 2)}.$$

We now show that the claimed bound follows. We distinguish two cases; first assume  $g(e_i) \leq \frac{1}{2}$ . Viewed as a function of  $g(e_i) \in [0, \frac{1}{2}]$ , the upper bound on  $h(S_{\text{rf}})$  is a convex function since its

second derivative w.r.t.  $g(e_i)$  is

$$\frac{2x^2}{(2g(e_i) + x - 2)^3} > 0$$

(for  $x > 2$ ). The maximum (worst upper bound) is thus attained at the extreme points, either  $g(e_i) = 0$  or  $g(e_i) = \frac{1}{2}$ :

$$h(S_{\text{rf}}) \leq \max \left\{ x + \frac{x^2}{4(x-2)}, x + \frac{1}{2} + \frac{x^2}{4(x-1)} \right\}. \quad (2)$$

The second case,  $g(e_i) > \frac{1}{2}$ , is trivial because then there are no possible blocking edges, otherwise one of  $G_u$  and  $G_v$  would be greater than 1. The edge  $e_i$  is then always scheduled immediately upon reaching  $x$ , and, thus,  $h(S_{\text{rf}}) \leq x + 1 \leq x + \frac{1}{2} + \frac{x^2}{4(x-1)}$  for  $x \geq 2$ .  $\square$

This bound on  $h(S_{\text{rf}})$  gives us the following corollaries.

**Corollary 3.11:** *For Reduce-Fastest(4), we have  $h(S_{\text{rf}}) \leq 6$ .*  $\triangleleft$

The above corollary gives the upper bound on the maximum heat achieved by Reduce-Fastest(4), which also meets the bound from Theorem 3.6. Optimizing the bound in Equation (2) allows to slightly improve upon this by changing to  $x = 2 + \frac{2\sqrt{5}}{5}$ .

**Corollary 3.12:** *Reduce-Fastest( $2 + \frac{2\sqrt{5}}{5} \approx 2.89$ ) gives  $h(S_{\text{rf}}) \leq 3 + \sqrt{5} < 5.24$ .*  $\triangleleft$

### 3.3 Length of schedules

In this section, we will show how an arbitrary schedule for an OPS problem can be modified to produce a periodic schedule with a polynomial-sized period and a maximum heat of at most 4 times the maximum heat of the original schedule.

This proof uses an overloading of the definition of heat  $h$  such that  $h(S, R')$  is the heat achieved by applying some cyclic schedule  $S$  to an OPS instance  $\mathcal{P}'_o = (P, R', g)$  where  $R' \subseteq R$ , and  $h(S, e, d)$  is the heat of a single edge  $e \in R$  that results from applying the same schedule  $S$  such that  $d$  is the duration since  $e$  was last scheduled.

**Lemma 3.13:** *Consider an OPS instance  $\mathcal{P}_o = (P, R, g)$ . Given an **Arbitrary** schedule  $S_A$  for  $\mathcal{P}_o$ , we can construct a **Polynomial** schedule  $S_P$  with period  $2\chi_1(P, R)$  and heat  $h(S_P) \leq 4h(S_A)$ .*  $\triangleleft$

**Proof:** Recall that  $r_S(e)$  is the recurrence time of an edge  $e$  in any schedule  $S$ . First, consider a **Truncated** schedule  $S_T$  which consists of the first  $\chi_1$  characters of  $S_A$ . Let the subset of edges  $R_T \subseteq R$  consist of the edges  $e \in R_T$  that appear in  $S_T$ . Using these, prove the following claim:

**Claim 3.14:** *Given the schedules  $S_A$  and  $S_T$ , defined above, it holds that  $h(S_T, R_T) \leq 2h(S_A)$ .*  $\triangleleft$

**Proof:** For the first  $\chi_1$  days,  $S_A$  and  $S_T$  are identical, so both have heat at most  $h(S_A)$  within that period. Any edge  $e \in R_T$  will have a heat  $h(S_T, e, d) = g(e) \cdot d$ , where  $d$  is again the duration since  $e$  was last scheduled. The set of such durations  $D$  includes the duration  $d_b$  before the first meeting, the duration  $d_a$  after the last meeting, and the durations  $D_i = d_1, d_2 \dots d_i$  between meetings, so  $D = \{d_b\} \cup \{d_a\} \cup D_i$ . Note that the heats of edges grow on days when they are scheduled, so the sum of all such durations  $\sum_{d \in D} d$  is the period  $\chi_1$  of  $S_T$ .

For the first  $\chi_1$  days,  $S_T$  does not repeat, so it holds that:

$$h(S_A) \geq \max_{d \in D}(h(S_T, e, d)) = g(e) \cdot \arg \max(\{d_b\} \cup \{d_a\} \cup D_i)$$

Thus:

$$g(e) \leq \frac{h(S_A)}{\arg \max(\{d_b\} \cup \{d_a\} \cup D_i)} \quad (3)$$

After  $\chi_1$  days,  $S_A$  and  $S_T$  diverge, with  $S_T$  repeating itself. The maximum heat of  $e$  in this period  $h(S_T, e)$  is therefore given by:

$$h(S_T, e) = \max_{d \in D}(h(S_T, e, d)) = g(e) \cdot \arg \max(\{(d_b + d_a)\} \cup D_i)$$

If  $h(S_T, e) > h(S_A)$ , it follows that  $h(S_T, e) = g(e) \cdot (d_b + d_a)$ . Which, along with Equation (3) shows that for all  $e \in R_T$ :

$$h(S_T, e) \leq h(S_A) \frac{d_b + d_a}{\arg \max(\{d_b\} \cup \{d_a\} \cup D_i)} \leq 2h(S_A)$$

Hence,  $h(S_T, R_T) \leq 2h(S_A)$ .  $\square$

Let  $C := \chi_1(P, R)$  be the number of colours required for edge colouring the graph  $(P, R)$ . Observe that for an OPS instance  $\mathcal{P}_o$ , one can always obtain a feasible schedule from a proper edge colouring  $c : e \rightarrow [C]$  of the graph  $(P, R)$ , where  $e \in R$ , by taking any round-robin schedule of the  $C$  colours. More formally, we can define a **coloured** schedule  $S_C(t) := \{e \in R : c(e) \equiv t \pmod{C}\}$ . As shown by [GSW24][Proposition 5.1], unweighted OPS is the same problem as edge colouring; thus, the round-robin schedule  $S_C$  of the coloured edges will cover all of the edges.

Finally, we construct the Polynomial schedule  $S_P$  by interleaving the Truncated schedule  $S_T$  with the coloured schedule  $S_C$ . This means that for any edge  $e \in R_T$ , the recurrence times obey  $2r_{S_T}(e) \leq r_{S_P}(e)$ . Now consider the complement to  $R_T$ ,  $R_T^C := R \setminus R_T$ . As edges  $e' \in R_T^C$  do not appear in the first  $\chi_1$  days of  $S_A$ , their growth rate is bounded by  $g(e') < \frac{h(S_A)}{\chi_1}$ , and they grow to heat  $h(S_C, R_T^C) < h(S_A)$  over the course of  $S_C$ . This completes the claim: for any edge  $e' \in R_T^C$ , we have  $2r_{S_C}(e') = r_{S_P}(e')$ . In other words, the recurrence time of any edge in  $S_P$  is at most twice its recurrence time in either  $S_T$  or  $S_C$ .

This finally demonstrates that  $h(S_P) \leq 2h(S_C) \leq 2h(S_T, R_T)$  and therefore, from Claim 3.14, we have  $h(S_P) \leq 4h(S_A)$ .  $\square$

**Corollary 3.15:** As BGT is a special case of OPS on star graphs, Lemma 3.13 also applies to BGT.  $\triangleleft$

**Remark 3.16:** Only the first  $\chi_1$  days of  $S_A$  appear in  $S_P$ , and only edges  $e' \in R_T^C$  need to appear in  $S_C$ . Further works may apply this process to unsustainable schedules, taking advantage of these facts to reduce overall heat.  $\triangleleft$

## 4 Poly Density

This section introduces two key results: Section 4.1 provides improved upper and lower bounds by proving Theorem 1.1, while Section 4.2 shows that poly density for DPS generalizes density for PWS and proves Theorem 1.2, showing that OPS instances with  $d \leq \frac{1}{4}$  are schedulable.

#### 4.1 Bounding Poly Density for OPS polycules

Recall the notion of poly density introduced by [GSW24], which follows the same intuition as Pinwheel Scheduling: relax the requirement that one complete task must be done each day, instead allowing any fractions of a day to be spent on each task. Use this relaxation to identify classes of schedulable or unschedulable instances. Pinwheel Scheduling uses an implicit schedule of unit length in which each task  $i$  with frequency  $f_i$  has contribution  $y_i \geq \frac{1}{f_i}$ . Task  $i$  is then considered to be performed after  $\frac{1}{y_i}$  days, with the density of the PWS instance given by the sum of these contributions  $d = \sum_{i=1}^k \frac{1}{f_i}$ . If  $d > 1$  then these contributions cannot fit into a single day, and the instance clearly cannot be scheduled.

[GSW24] applied similar reasoning to determine the poly density of Optimisation Polyamorous Scheduling problems, obtaining a lower bound on the optimal heat  $h^*$ . First, compose an optimisation problem using the same logic as Pinwheel Scheduling: construct a fractional schedule of unit length by assigning contribution  $y_M \in [0, 1]$  to each matching  $M$ . The heat  $\bar{h}$  of the best such schedule provides a lower bound on  $h^*$ .

$$\min \bar{h} \tag{4}$$

$$\text{s.t.} \quad \sum_{M \in \mathcal{M}} y_M \leq 1 \tag{5}$$

$$\frac{g(e)}{\sum_{M \in \mathcal{M}: e \in M} y_M} \leq \bar{h} \quad \forall e \in R \tag{6}$$

$$y_M \in [0, 1] \tag{7}$$

[GSW24] defines the poly density  $\bar{h}^*$  of an OPS polycule  $\mathcal{P}_o = (P, R, g)$  as the optimal objective value of the fractional-schedule LP.<sup>1</sup> From the dual of this LP, [GSW24] obtain a more explicit term for  $\bar{h}^*$ . A dual solution consists of weights  $z_e$  for edges  $e \in R$  with  $\sum_{e \in R} z_e = 1$ . Intuitively these account for their growth rates  $g_e$  and the (inclusion-maximal) matchings  $M \in \mathcal{M}$  in which they appear. Any such weighting implies a lower bound  $\bar{h}(z)$  on  $\bar{h}^*$  and hence on  $h^*$ , with the best choice of  $z$  yielding the poly density

$$\bar{h}^* = \max_{z \in \mathcal{Z}} \bar{h}(z), \quad \text{where } \bar{h}(z) = \min_{M \in \mathcal{M}} \frac{1}{\sum_{e \in M} \frac{z_e}{g(e)}}, \tag{8}$$

where  $\mathcal{Z}$  is the set of all weightings.

Recall from Section 2 that the maximum personal growth over some polycule  $\mathcal{P}_o = (P, R, g)$  is given by  $G^* = \max_{v \in P} \sum_{e \in R: v \in e} g(e)$ . We now show that the poly density is within a constant factor of  $G^*$ , proving Theorem 1.1 in two parts: The lower bound by Lemma 4.1 and the upper bound by Lemma 4.3.

**Lemma 4.1:** *For any OPS instance  $\mathcal{P}_o = (P, R, g)$ , the poly density  $\bar{h}^*$  is bounded by  $\bar{h}^* \geq G^*$ .<sup>◁</sup>*

**Proof:** Let  $v^*$  be the vertex with the maximum local heat  $G^*$ , i.e.

$$v^* = \arg \max_{v \in P} \sum_{e \in R: v \in e} g(e).$$

Consider the weighting  $z^*$  with  $z_e = \frac{g(e)}{G^*}$  when  $v^* \in e$  and  $z_e = 0$  otherwise. For  $z^*$ , it holds that

$$\max_{M \in \mathcal{M}} \sum_{e \in M} \frac{z_e}{g(e)} = \frac{1}{G^*}$$

<sup>1</sup>A change of variables  $\ell = 1/\bar{h}$  and taking reciprocals in the second constraint yields this LP.

and  $\bar{h}(z) = G^*$ . Since  $\bar{h}^* \geq \bar{h}(z)$  for all weights  $z$ , the inequality holds.  $\square$

**Lemma 4.2:** *For any OPS instance  $\mathcal{P}_o = (P, R, g)$  with integer growth rates,  $\bar{h}^* \leq \frac{3}{2}G^*$ .*  $\triangleleft$

**Proof:** Begin by constructing an unweighted multigraph  $\mathcal{P}'_o = (P, E)$  such that each edge  $e \in R$  is replaced by a bundle of  $g(e)$  parallel edges in  $E$ , noting that the maximum degree  $\Delta$  of  $\mathcal{P}'_o$  is exactly  $G^*$ . By [Sha49], we can colour edges in  $\mathcal{P}'_o$  using  $C \leq \frac{3}{2}\Delta$  colours. Given that each colour induces a matching in  $\mathcal{P}_o$  such that each edge  $e \in R$  appears in exactly  $g(e)$  colours, construct a fractional solution by assigning a contribution  $y_M = \frac{1}{C}$  to each colour class  $M$ . Checking Equations (5)–(7) shows that this is a valid fractional schedule with heat

$$\bar{h} \geq g(e) \Big/ \sum_{M \in \mathcal{M}: e \in M} y_M = C \quad \square$$

for all edges. Hence,  $\bar{h} = C \leq \frac{3}{2}G^*$  for this assignment, and  $\bar{h}^* \leq \frac{3}{2}G^*$  globally.

**Lemma 4.3:** *The poly density  $\bar{h}^*$  of any OPS instance  $\mathcal{P}_o = (P, R, g)$  is bounded by  $\bar{h}^* \leq \frac{3}{2}G^* + \varepsilon$  for any  $\varepsilon > 0$ .*  $\triangleleft$

**Proof:** When all heat growth rates are rational numbers, we can scale all of them by their common denominator to reduce them to the integral case. If any growth rate is irrational, we can upper-bound it with a rational number that is arbitrarily close.  $\square$

## 4.2 Poly Density of DPS polycules

In this section, we turn our attention to the poly density of Decision Polyamorous Scheduling: showing that it generalizes the density of Pinwheel Scheduling with Theorem 4.5, introducing the PolyGreedy algorithm and analysing it to prove Theorem 4.6 and that  $d \leq \frac{1}{4}$  implies schedulability with a proof for Theorem 1.2. When introduced by [GSW24], poly density was chiefly discussed in the context of OPS, but the poly density of a DPS instance  $\mathcal{P}_d = (P, R, f)$  was defined as the poly density of a paired OPS instance  $\mathcal{P}_o = (P, R, \frac{1}{f})$ :

**Definition 4.4 (Poly Density of DPS instances):** *The poly density  $d$  of a DPS polycule  $\mathcal{P}_d = (P, R, f)$  is given by*

$$d = \max_{z \in \mathcal{Z}} \frac{1}{\max_{M \in \mathcal{M}} \sum_{e \in M} z_e \cdot f_e}$$

where  $\mathcal{Z}$  is the set of all weightings applying  $z_e \in [0, 1]$  to each edge  $e \in R$  such that  $\sum_{e \in R} z_e = 1$  and  $\mathcal{M}$  is the set of inclusion maximal matchings  $M$  in the graph  $(P, R)$ .  $\triangleleft$

Given the equivalence between DPS polycules on star graphs  $\mathcal{P}_d^* = (P, R, f)$  and Pinwheel Scheduling instances with the same frequency set  $f$ , this is a generalisation of the density of Pinwheel Scheduling:

**Theorem 4.5:** *The poly density  $d$  of any DPS polycule  $\mathcal{P}_d^* = (P, R, f)$  whose graph  $(P, R)$  is a star is given by  $d = \sum_{e \in R} \frac{1}{f_e}$*   $\triangleleft$

**Proof:** Begin with the poly density defined by Definition 4.4. As  $\mathcal{P}_d^*$  is a star graph, each matching  $m \in \mathcal{M}$  contains a single edge  $e \in R$ , so  $d = \max_{z \in \mathcal{Z}} 1/(\max_{e \in R} z_e \cdot f_e)$ . This will be solved by the value of  $z$  which minimizes  $\max_{e \in R} z_e \cdot f_e$ . As  $\sum_{e \in R} z_e = 1$ , any assignment  $\mathcal{Z}$

which does not set  $z_e \cdot f_e = z_{e'} \cdot f_{e'}$  for all  $e, e' \in R$  will have a higher maximum than one that does. Try

$$z_e = \frac{1}{f_e \sum_{e' \in R} \frac{1}{f_{e'}}},$$

observing that

$$\sum_{e \in R} z_e = \frac{1}{\sum_{e' \in R} \frac{1}{f_{e'}}} \cdot \sum_{e \in R} \frac{1}{f_e} = 1,$$

and that  $d = \sum_{e \in R} \frac{1}{f_e}$ . □

We continue by identifying two classes of schedulable DPS instances which instead restrict the *local density* of vertices – the density of the portion of the instance directly connected to any single vertex. This analysis is motivated by [HMR<sup>+</sup>89] and uses the following PolyGreedy algorithm for DPS:

---

**Algorithm 1** PolyGreedy( $P, R = \{e_1, \dots, e_{|R|}\}, f$ )

---

```

1: /* Assumes  $f(e_i) = 2^{k_i}$  for some  $k_i \in \mathbb{N}$  for all  $e_i \in R$  */
2: /* W.l.o.g. let  $f(e_1) \leq f(e_2) \leq \dots \leq f(e_{|R|})$  */
3:  $f_{\max} = f(e_{|R|})$ 
4: for  $t = 0, \dots, f_{\max}$  do
5:    $S[t] \leftarrow \emptyset$ 
6: end for
7: for  $i = 1, \dots, |R|$  do
8:    $s_i \leftarrow \min\{t \in [0..f_{\max}) : S[t] \cup \{e_i\} \text{ is a matching}\}$ 
9:   for  $k = 1, \dots, f_{\max}/f_i$  do
10:     $t \leftarrow s_i + (k-1) \cdot f_i$ 
11:     $S[t] \leftarrow S[t] \cup \{e_i\}$ 
12:   end for
13: end for
14: return  $S[0..f_{\max})$ 
```

---

**Theorem 4.6:** Let  $\mathcal{P}_d = (P, R, f)$  be a DPS instance such that

1.  $\forall e \in R, f_e$  is a power of two

2.  $\forall v \in V, \sum_{e \in R: v \in e} \frac{1}{f_e} \leq \frac{1}{2}$ .

Then, there exists a valid schedule for  $\mathcal{P}_d$ . ◁

**Proof:** Let  $f_{\max}$  be a maximum frequency among all edges and construct a periodic schedule of size  $f_{\max}$  using Algorithm 1. By Definition 2.2, valid schedules lack conflicts and respect frequencies; we proceed by showing inductively that Algorithm 1 respects both properties.

In the initial state,  $S$  is empty and thus contains no conflicts. Assuming that there are no conflicts after  $i-1$  edges have been added, observe that adding edge  $e_i$  initially places it in a conflict-free slot  $s_i$ , then in slots  $s_i + f_i, s_i + 2f_i, s_i + 3f_i, \dots$ . As all frequencies  $f_e$  are powers of two and all edges placed thus far obey  $f_j \leq f_i$ , slots  $s_i + f_i, s_i + 2f_i, s_i + 3f_i, \dots$  must also be conflict-free for  $e_i$ .



If we again consider an arbitrary edge  $e_i$  at any point in  $S$  after its initial inclusion  $s_i$ , it is clear that the schedule is frequency respecting. Initially,  $e_i$  is placed in the first slot with no edges that conflict with  $e_i$ . If  $e_i = \{u, v\}$ , the number of slots with conflicting edges is at most

$$\begin{aligned} \sum_{\substack{e_j: u \in e_j \\ f_j < f_i}} \frac{f_i}{f_j} + \sum_{\substack{e_j: v \in e_j \\ f_j < f_i}} \frac{f_i}{f_j} &= f_i \left( \sum_{\substack{e_j: u \in e_j \\ f_j < f_i}} \frac{1}{f_j} + \sum_{\substack{e_j: v \in e_j \\ f_j < f_i}} \frac{1}{f_j} \right) \\ &< f_i \left( \frac{1}{2} + \frac{1}{2} \right) \\ &= f_i \end{aligned}$$

Hence, there must be at least one slot  $s_i$  among the first  $f_i$  slots with no conflicting edges. The last occurrence of  $e_i$  in  $S$  is at  $s_i + f_{\max} - f_i$ , so the recurrence time  $r$  between the last and first instances of  $e_i$  in the periodic solution is given by  $f_i$   $\square$

**Proof of Theorem 1.2:** Begin by constructing an OPS instance  $\mathcal{P}'_o = (P, R, g)$  with  $g_e = \frac{1}{f_e}$  for all edges  $e \in R$ ; by Definition 4.4, the poly density  $\bar{h}^*$  of  $\mathcal{P}'_o$  equals the density  $d$  of the corresponding  $\mathcal{P}_d = (P, R, f)$ . By Lemma 4.1, the maximum personal growth rate  $G^*$  of  $\mathcal{P}'_o$  is at most  $\frac{1}{4}$ , so

$$\frac{1}{4} \geq G^* = \max_{v \in P} \sum_{e \in R: v \in e} g(e) = \max_{v \in P} \sum_{e \in R: v \in e} \frac{1}{f_e},$$

and  $\forall v \in V, \sum_{e \in R: v \in e} \frac{1}{f_e} \leq \frac{1}{4}$ .

Now construct  $\mathcal{P}'_d = (P, R, f')$  such that  $\forall e \in R, f'(e) = 2^{\lfloor \log_2 f_e \rfloor}$ , i.e. round all the frequencies down to the nearest power of two such that  $\sum_{e: v \in e} \frac{1}{f'_e} \leq \sum_{e: v \in e} \frac{2}{f_e} \leq 2 \cdot \frac{1}{4} = \frac{1}{2}$ . Applying Theorem 4.6 then gives a feasible schedule for  $\mathcal{P}'_d$ , which must also solve  $\mathcal{P}_d$ .  $\square$

## 5 Inapproximability

In this section, we prove Theorem 1.3; showing, via a reduction from 3SAT, that bipartite OPS does not allow efficient  $(1 + \delta)$ -approximation algorithms for  $\delta < \frac{1}{12}$  unless  $P = NP$ .

### 5.1 Overview of Proof

The proof of Theorem 1.3 has two steps: a reduction from 3SAT to bipartite DPS, and a conversion to bipartite OPS such that a  $\frac{13}{12}$  approximation is shown to be NP-hard.

3SAT is the following problem: given a 3-CNF formula  $\varphi = c_1 \wedge c_2 \wedge \dots \wedge c_m$  with clauses  $C = \{c_1, c_2, \dots, c_m\}$ , each consisting of exactly 3 literals over variables  $X = \{x_1, x_2, \dots, x_n\}$ , decide whether there is an assignment of Boolean values to the variables in  $X$  such that all  $m$  clauses in  $C$  are satisfied.

**Lemma 5.1 (3SAT  $\leq_p$  bipartite DPS):** *For any 3-CNF formula  $\varphi$  with  $m$  clauses, we can construct in polynomial time a bipartite decision polycule  $\mathcal{P}_{d\varphi}$  which has a valid schedule if and only if all clauses of  $\varphi$  can be simultaneously satisfied.*  $\triangleleft$

The bulk of Section 5 will introduce our construction of  $\mathcal{P}_{d\varphi}$ , which we will use to prove this Lemma. This construction is designed to be modular and communicable, rather than efficient with constant factors.

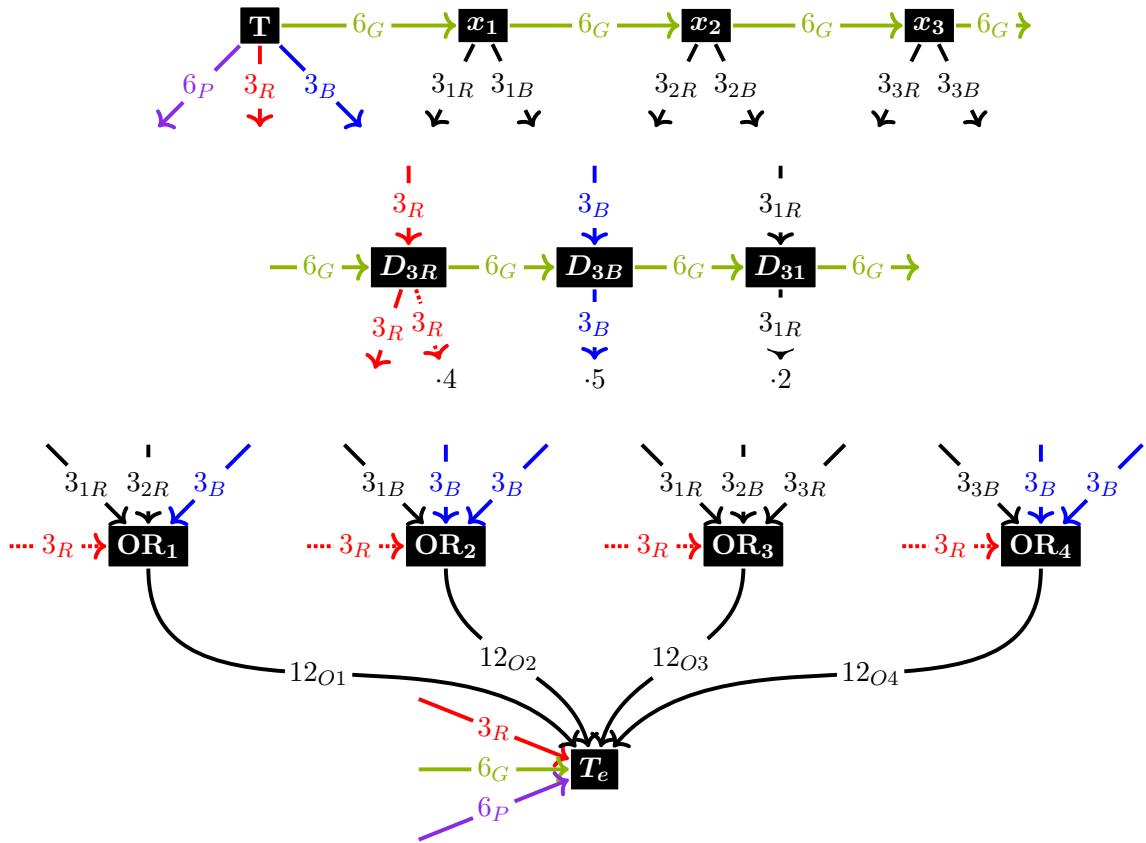
The second step in the proof of Theorem 1.3 is to convert the decision polycule  $\mathcal{P}_{d\varphi}$  from Lemma 5.1 to an optimisation polycule  $\mathcal{P}_{o\varphi}$  using Lemma 2.4. It will be immediately obvious

from the construction that the largest frequency in  $\mathcal{P}_{d\varphi}$  is  $F = 12$ . So by Lemma 2.4,  $\mathcal{P}_{o\varphi}$  either has  $h^* = 1$  (if  $\mathcal{P}_{d\varphi}$  is feasible) or else has  $h^* \geq \frac{13}{12}$ .

**Proof of Theorem 1.3:** Assume that there is a polynomial-time approximation algorithm  $A$ , which solves bipartite OPS with approximation ratio  $\alpha < \frac{13}{12}$ . If  $\mathcal{P}_{o\varphi}$  has optimal heat  $h^* = 1$ ,  $A$  produces a schedule with heat  $1 \leq h \leq \alpha < \frac{13}{12}$ ; whereas if  $h^* \geq \frac{13}{12}$ ,  $A$  must produce a schedule of heat  $\frac{13}{12} \leq h \leq \alpha \cdot \frac{13}{12}$ . By running  $A$ , we are therefore able to distinguish between  $h^* \leq 1$  and  $h^* \geq \frac{13}{12}$  for  $\mathcal{P}_{o\varphi}$ , and hence between the feasibility or infeasibility of  $\mathcal{P}_{d\varphi}$ . This then discriminates between Yes and No instances of 3SAT, via the polynomial-time reduction from Lemma 5.1. 3SAT is NP-hard, hence  $P = NP$  follows.  $\square$

Let us denote by  $\alpha^*$  the approximability threshold for bipartite OPS, that is: efficient polynomial-time approximation algorithms with approximation ratio  $\alpha$  exist if and only if  $\alpha \geq \alpha^*$  (assuming  $P \neq NP$ ). Theorem 1.3 shows that  $\alpha^* \geq \frac{13}{12}$ , while Theorem 3.10 shows that  $\alpha^* \leq 5.24$ , leaving a substantial gap. We conjecture that the constant  $\frac{13}{12}$  can be improved by careful analysis of our construction, but we leave this to future work.

**Conjecture 5.2:** The approximability threshold for bipartite OPS is  $\alpha^* \geq \frac{4}{3}$ .  $\triangleleft$



**Figure 2:** A DPS polycule which is schedulable iff there is some assignment that satisfies  $(x_1 \vee x_2), (\overline{x_1}), (x_1 \vee \overline{x_2} \vee x_3)$ , and  $(\overline{x_3})$  (which is not possible). Connections between layers are omitted for clarity but flow from top to bottom, starting with the variable layer, then a layer duplicating nodes with frequency 3, the  $OR$  layer, and finally the tensioning layer. Note that larger problems will also require a  $D_6$  layer to support additional tension gadgets.

## 5.2 Overview of Reduction

We now give the proof of Lemma 5.1. For the remainder of this section, we will assume that a 3-CNF formula  $\varphi = c_1 \wedge \dots \wedge c_m$  over variables  $X = \{x_1, \dots, x_n\}$  is given. Using this, we will show how to construct DPS instances  $\mathcal{P}_{d\varphi}$  that admits a schedule iff there is a variable assignment  $v : X \rightarrow \{\text{True}, \text{False}\}$  that satisfies all clauses in  $C = \{c_1, \dots, c_m\}$ . This construction directly represents components of Boolean formulas as bipartite DPS “gadgets”:

- *variables* (Section 5.4),
- *clauses (OR gadgets)* (Section 5.7), and
- a check that all clauses are True, the *tension gadget* (Section 5.8).

To make these gadgets work, we require further auxiliary gadgets:

- a “*True Clock*” to break ties between symmetric choices for schedules (Section 5.3),
- slot duplication gadgets:  $D_3$  *duplicators* (Section 5.6.1), and  $D_6$  *duplicators* (Section 5.6.2).

The overall conversion algorithm is given by Definition 5.3 and a worked example is shown in Figure 2.

**Definition 5.3 ( $\mathcal{P}_{d\varphi}$  polycules):** The decision polycule  $\mathcal{P}_{d\varphi} = (P, R, f)$  is constructed in layers as follows:

(a) **Variable layer:**

The variable layer consists of a variable gadget with outputs  $3_{iR}$  and  $3_{iB}$  for each variable  $x_i \in X$ , as well as a single special variable: the True Clock.

(b) **Duplication layer:**

The duplication layer duplicates outputs of the variable layer:  $D_3$  duplicators create one  $3_{iR}$  edge for each  $x_i$  in clauses  $c \in C$ , and one  $3_{iB}$  edge for each occurrence of  $\bar{x}_i$  in clauses  $c \in C$ . They also create as many  $3_R$  and  $3_B$  edges as are needed, while  $D_6$  duplicators do the same for  $6_G$  and  $6_P$  edges. Both male and female constant edges are created, and unused edges are connected to pendent nodes.

(c) **Clause layer:**

The clause layer consists of one OR gadget for each clause  $c_j \in C$ . OR gadgets have three inputs, each corresponding to a literal in  $c_j$ :  $3_{iR}$  for  $x_i$ ,  $3_{iB}$  for  $\bar{x}_i$ . For clauses with less than 3 literals,  $3_B$  edges fill the OR gadget’s unused inputs.

(d) **Tension layer:**

The tensioning layer attaches tension gadgets to all outputs of the sorting layer; any spare inputs to the tensioning layer are connected to pendant nodes. ◁


## 5.3 The True Clock & Colour Slots

Figure 3 introduces gadgets representing sample variables  $x_1$  and  $x_2$ , as well as a special variable: the True Clock  $T$ , which acts as a drumbeat for the polycule as a whole. Variables, including  $T$ , have four relationships:  $[3, 3, 6, 6]$ , so their local schedules must all have the following form:  $[3_a, 3_b, 6_a, 3_a, 3_b, 6_b]$ . As schedules are cyclic, we can choose to start this schedule with the  $3_a$  edge of  $T$  without loss of generality. We then assign names to these edges as dictated by the slots they are in:

**Definition 5.4 (Slots):** We call days  $t \in \mathbb{N}_0$  with  $t \equiv 0 \pmod{3}$  the *red slots*, days  $t \equiv 1 \pmod{3}$  *blue slots*, days  $t \equiv 2 \pmod{6}$  *green slots*, and days  $t \equiv 5 \pmod{6}$  *purple slots*.  $\triangleleft$

The local schedule of the True Clock is therefore given by  $[3_R, 3_B, 6_G, 3_R, 3_B, 6_P]$ . As  $3_R$  is scheduled on day 0, it will always be assigned in red slots, with  $3_B$ ,  $6_G$ , and  $6_P$  edges also restricted to slots of their respective colours. We will sometimes represent this by underlining elements or gaps in a schedule, e.g.,  $[\underline{\phantom{0}}, \underline{\phantom{0}}, \underline{\phantom{0}}, \underline{\phantom{0}}, \underline{\phantom{0}}, \underline{\phantom{0}}]$ , or by referring to edges as being red, blue, green, or purple.

All gadgets introduced below are constructed such that the lengths of their schedules are integer multiples of 6; in the final polycule  $\mathcal{P}_{d\varphi}$  they will be connected (usually through intermediaries) to the True Clock, such that their edges must stick to certain slots. To keep correctness proofs of individual gadgets readable, we call a schedule  $S$  that schedules all coloured edges in slots of the given colour *slot-respecting*.



**Figure 3:** Gadgets for the True Clock and for sample variables  $x_1$  and  $x_2$  (top, left to right), with shorthand versions shown below. Gadgets are shown connected as they would be in a sample variable layer, and their colours and schedules are discussed in Section 5.3. Further variables can be added to the right.

**Drawing Conventions.** Gadgets are connected by input and output edges, represented by incoming and outgoing arrows respectively – each one being half of a relationship between two people from different gadgets. In addition to their frequencies, input and output edges share restrictions on their permissible schedules, carrying them from one gadget to another as discussed in the proofs associated with each gadget.

In shorthand gadgets, vertical incoming edges are the primary *input* to a gadget, encoding the value of some variable or logical function; horizontal inputs contain edges of fixed colour, which we will refer to as *constants*. Similarly, vertical outgoing edges represent primary *outputs* that encode the result of the gadget, while horizontal outgoing edges represent incidentally created constants which may either be used by other gadgets or connected to pendent vertices.

Some incoming and outgoing edges will have end labels of the form “ $\cdot i$ ”, indicating  $i$  connections of the given type, each with a different person.

To show that  $\mathcal{P}_{d\varphi}$  is bipartite, we assign a sex to each node: male nodes with solid edges, and female nodes with dotted edges. We will also use a directed graph to show the sex of the originating node: edges coming from males have solid edges, while edges coming from female nodes will have dotted edges. Note that this convention is only used to show that the complete graph is bipartite while reasoning about parts of it; in all other respects, the graph is undirected (as we would expect for poly scheduling).

## 5.4 Variables

Figure 3 also introduces the gadget for a sample variable,  $x_1$ , which again has four relationships:  $[3_{1R}, 3_{1B}, 6_G, 6_P]$ . The key property of variable gadgets is summarized in the following Lemma.

**Lemma 5.5 (Variable gadget schedules):** *The  $x_i$  node in each variable gadget  $x_i$  has a valid local schedule which must be of the form  $[3_a, 3_b, 6_G, 3_a, 3_b, 6_P]$  in any slot-respecting global schedule.*  $\triangleleft$

**Proof:**  $x_i$  nodes have tasks  $[3_{1R}, 3_{1B}, 6_G, 6_P]$  and local density  $D = 1$ , so  $3_{iR}$  and  $3_{iB}$  must be scheduled exactly once in each 3-day period, forcing every 3 days to be of the form  $[3_a, 3_b, \_]$ , and every 6-day schedule to be of the form  $[3_a, 3_b, \_, 3_a, 3_b, \_]$ ; this leaves two remaining slots, which must contain  $6_G$  and  $6_P$ .

The  $6_G$  edge of the first variable node,  $x_1$ , is shared with  $T$  such that it must be green, so valid schedules for  $x_1$  are of the form  $[3_a, 3_b, 6_G, 3_a, 3_b, \_]$ , which must be completed as  $[3_a, 3_b, 6_G, 3_a, 3_b, 6_P]$  – a valid schedule for  $x_1$ .

The  $6_P$  edge leaving  $x_1$  is then connected to an  $F_6$  flipper gadget, which returns a  $6_G$  edge by Lemma 5.6. Further variables are each connected to the  $6_G$  edge returned from the  $F_6$  flipper gadget of the previous variable gadget, so their schedules must be of the same form. It therefore follows inductively that each variable gadget  $x_i$  has a valid schedule, and that it must be of the form  $[3_a, 3_b, 6_G, 3_a, 3_b, 6_P]$ .  $\square$

According to Lemma 5.5, the incident  $6_G$  edge is used to restrict the valid schedules for  $x_1$ , leaving two possibilities:  $[3_{1R}, 3_{1B}, 6_G, 3_{1R}, 3_{1B}, 6_P]$  and  $[3_{1B}, 3_{1R}, 6_G, 3_{1B}, 3_{1R}, 6_P]$ . The former schedule, where  $3_{1R}$  is scheduled in red slots, corresponds to a variable assignment where  $x_1$  is True, whereas the the second schedule corresponds to  $x_1$  being assigned False.

This technique of connecting  $3_R$ ,  $3_B$ ,  $6_G$ , or  $6_P$  edges to nodes in a gadget in order to limit their valid local schedules and force relationships between edges, slots, and particular meanings will be used extensively in what follows.

Note that  $3_{iB}$  has the opposite value to  $3_{iR}$ , so using  $3_{iB}$  edges in the polycule corresponds to the negated literal  $\bar{x}_i$ , just as  $3_{iR}$  edges correspond to the literal  $x_i$ .

## 5.5 Flippers


Flippers are introduced by Figure 4, and are used to ensure that  $\mathcal{P}_{d\varphi} = (P, R, f)$  is bipartite by converting between male and female edges. They allow variable and duplicator gadgets to consume a constant, rather than linear, amount of  $3_b$ ,  $6_G$ , and  $6_P$  edges.

**Lemma 5.6 (F-6 Flipper gadget schedules):** *The  $F_6$  node in each  $F_6$  flipper gadget has a valid local schedule which must be of the form  $[3_a, 3_b, 6_G, 3_a, 3_b, 6_P]$  in any slot-respecting global schedule.*  $\triangleleft$

**Proof:** Note that  $F_6$  flippers and variables share the same tasks:  $[3, 3, 6, 6]$ . By the proof of Lemma 5.5, schedules for  $F_6$  flippers must be of the form  $[3_a, 3_b, \_, 3_a, 3_b, \_]$ , where the remaining slots must contain  $6_G$  and  $6_P$  edges. The incoming edge, either a  $6_G$  or  $6_P$ , must be in a corresponding slot forcing schedules to be of the form  $[3_a, 3_b, 6_G, 3_a, 3_b, \_]$  or  $[3_a, 3_b, \_, 3_a, 3_b, 6_P]$  respectively. Both must yield the complete and valid local schedule  $[3_a, 3_b, 6_G, 3_a, 3_b, 6_P]$ .  $\square$

**Lemma 5.7 (F-3 Flipper gadget schedules):** *The  $F_3$  node in each  $F_3$  flipper gadget has a valid local schedule which must be of the form  $[3_R, 3_B, 6_G, 3_R, 3_B, 6_P]$  in any slot-respecting global schedule.*  $\triangleleft$

**Proof:**  $F_3$  flipper nodes have two inputs,  $[3_i, 6_i]$  and two outputs,  $[3_o, 6_o]$ .  $3_i$  and  $6_i$  are indirectly connected to the True Clock such that  $3_i$  must be red or blue, and  $6_i$  must be either green or purple. This forces partial schedules to be of the form  $[3_i, \_, 6_i, 3_i, \_, \_]$  or  $[3_i, 6_i, \_, 3_i, \_, \_]$ . As  $F_3$  nodes have density  $D = 1$ , the  $6_o$  edges must be placed 3 days after the  $6_i$  edge to avoid violating the condition of the  $3_o$  edge, for schedules of the form  $[3_i, 3_o, 6_i, 3_i, 3_o, 6_o]$  and  $[3_i, 6_i, 3_o, 3_i, 6_o, 3_o]$  respectively. If  $3_i$  is red,  $3_o$  will be blue, and vice versa; similarly, one 6 edge will be purple and the other green. In either case, the resultant schedule will be valid and of the form  $[3_R, 3_B, 6_G, 3_R, 3_B, 6_P]$ .  $\square$



**Figure 4:** Four sample flipper gadgets, each with shorthand versions below them:  $6_G$  male to  $6_P$  female (top left),  $6_P$  female to  $6_G$  male (top right),  $3_B$  male to  $3_R$  female (bottom left), and  $3_R$  female to  $3_B$  male (bottom right). Note that  $F_3$  flippers act incidentally as  $F_6$  flippers, but are used separately in our construction.

## 5.6 Duplication of Variables and Constants

Variables may appear in multiple clauses, while the constant  $3_R$ ,  $3_B$ ,  $6_G$ , and  $6_P$  edges are used in multiple gadgets, engendering a need for the duplication of variables and constants.

### 5.6.1 3-Duplicators


A gadget for duplicating edges with period 3 is shown in Figure 5 and proven to accurately reproduce its input by Lemma 5.8.

**Lemma 5.8 (3-Duplicator gadget schedules):** *The non-flipper nodes in each 3-duplicator gadget  $D_3$  have valid local schedules which must be of the form of either*

$$[3_a, 9_b, 6_G, 3_a, 9_{b'}, 6_P, 3_a, 9_{b''}, 6_G, 3_a, 9_b, 6_P, 3_a, 9_{b'}, 6_G, 3_a, 9_{b''}, 6_P] \text{ or} \\ [9_b, 3_a, 6_G, 9_{b'}, 3_a, 6_P, 9_{b''}, 3_a, 6_G, 9_b, 3_a, 6_P, 9_{b'}, 3_a, 6_G, 9_{b''}, 3_a, 6_P]$$

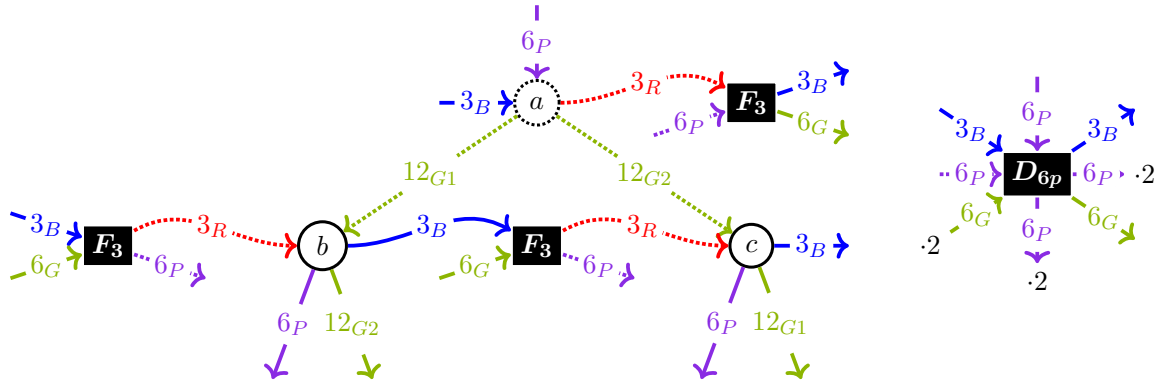
in any slot-respecting global schedule.

◁









**Figure 6:** A gadget for duplicating  $6_P$  input edges, with a shorthand version on the right. The gadget can be extended by adding further layers below to create as many  $6_P$  edges as are needed. Constant edges produced by nodes in high layers should be consumed by gadgets in lower layers, but these connections are not shown. Odd layers will produce female  $6_P$  edges and should have  $F_3$  flippers to the right of their nodes as shown in upper layer; even layers will produce male  $6_P$  edges, and should have  $F_3$  flippers to the left of their nodes as shown in the lower layer. The figure shows one female node ( $a$ ) and two female nodes ( $b, c$ ) – this creates an imbalance between the male  $6_G$  and female  $6_P$  nodes consumed and produced and should be avoided by adding nodes until there are as many male nodes as female nodes.  $6_G$  edges can be duplicated with a very similar  $D_{6_G}$  gadget simply by replacing the topmost  $6_P$  edge with a  $6_G$  edge.

As with  $a$ , exactly four of these slots must schedule  $3_B$  edges, and these must be the blue spaces for the  $3_B$  constraint not to be violated, forcing schedules of the form

$$[3_R, 3_B, 12_{Ga}, 3_R, 3_B, \_, 3_R, 3_B, \_, 3_R, 3_B, \_].$$

One of these slots must contain the  $12_{Gb}$  edge, with the other two scheduling the  $6_P$  edge. If the  $12_{Gb}$  edge is not scheduled in the remaining green slot, the constraint on the  $6_P$  edge will be violated, so the schedule must be as shown in the Lemma.  $\square$


**Corollary 5.10:** *Note that by replacing the  $6_P$  input edge with a  $6_G$  input edge, a  $D_{6G}$  replicator will be made instead that produces  $6_G$  edges according to the same logic.*  $\triangleleft$

## 5.7 Clauses

Clauses in  $C$  are disjunctions of at most 3 literals, i.e. the logical OR of at most 3 variables  $x_i$ , or their negations  $\bar{x}_i$ . A gadget which determines the truth value of a clause given the values of its variables is shown in Figure 7. Lemma 5.11 shows that the  $12_O$  output edges of this gadget can be blue iff the corresponding clause is evaluated to be True; Definition 5.3 and Lemma 5.12 ensure that all  $12_O$  edges must be blue in any valid global schedule. Thus, for  $\mathcal{P}_{d\varphi}$  to have a valid schedule,  $\varphi$  must have a satisfying assignment.

**Lemma 5.11 (OR gadget schedules):** *Consider an OR gadget whose input edges  $3_{1R}$ ,  $3_{2R}$ , and  $3_{3R}$  are either red or blue. There exists a slot-respecting schedule in which the  $12_O$  output edge of this OR gadget is blue iff at least one of its input edges is red.*  $\triangleleft$

While this Lemma only considers the clause  $x_1 \vee x_2 \vee x_3$ , all other clauses can be handled similarly: For a literal  $x_j$ , we use  $3_{jR}$  as input and for negated literals  $\bar{x}_j$ , we instead use  $3_{jB}$ .



**Figure 7:** A gadget which computes  $x_1 \vee x_2 \vee x_3$  (left), along with a shorthand version (right). To compute  $x_1 \vee x_2 \vee \bar{x}_3$ , replace the incoming  $3_{3R}$  edge with a  $3_{3B}$  edge. To instead compute  $x_1 \vee x_2$ , replace the incoming  $3_{3R}$  edge with a male  $3_B$  edge. Lemma 5.11 shows that the  $12_O$  edge can be blue iff at least one input is assigned True, while Definition 5.3 and Lemma 5.12 show that  $12_O$  edges must be blue in any valid global schedule.

Clauses with fewer literals also share the same gadget and proof: here, we replace each unused  $3_{jR}$  edge with a male  $3_B$  edge, effectively replacing clauses such as  $x_1 \vee x_2$  with analogous clauses in the form  $x_1 \vee x_2 \vee \perp$ .

**Proof:** The node labelled  $OR$  has tasks  $[3_R, 6_1, 6_2, 12_1, 12_2, 12_3, 12_O]$  and local density  $D = 1$ , so each task with frequency  $f$  must appear exactly once in every  $f$  day period. Any slot-respecting schedule must assign the  $3_R$  edge to red slots, so schedules for the  $OR$  node must be of the form  $[3_R, \_, \_, 3_R, \_, \_]$ .

Consider an inverter node  $I_{i \in \{1,2,3\}}$ , which has tasks  $[3_{iR}, 12_i]$  where  $3_{iR}$  is either red or blue by assumption. If the input edge of this node,  $3_{iR}$ , is red, then partial schedules for  $I_i$  must be of the form  $[3_{iR}, \_, \_, 3_{iR}, \_, \_]$  and the  $12_i$  edge must be either green, purple, or blue. Similarly, if  $3_{iR}$  is scheduled in blue slots, then partial schedules for  $I_i$  must instead be of the form  $[\_, 3_{iR}, \_, \_, 3_{iR}, \_]$  and the  $12_i$  edge must be either green, purple, or red; however, this  $12_i$  edge is also connected to the  $OR$  node, which has no empty red slots, further restricting it to be either green or purple.

Suppose that one input edge,  $3_{tR}$ , is scheduled in red slots, while the others,  $3_{iR}$  and  $3_{i'R}$ , may be scheduled in either red or blue slots. As  $3_{tR}$  is red, the output of the associated inverter,  $12_t$ , may be blue, green, or purple. Consider the schedule

$$[3_R, 12_t, 6_a, 3_R, 6_b, 12_i, 3_R, 12_O, 6_a, 3_R, 6_b, 12_{i'}],$$

observing that  $12_O$  is scheduled in blue slots and that no constraints are violated – a valid schedule.

Now suppose that all three inputs,  $3_{1R}$ ,  $3_{2R}$ , and  $3_{3R}$ , are scheduled in blue slots. By the reasoning above,  $12_1$ ,  $12_2$ , and  $12_3$  must now be scheduled in green or purple slots. This will force schedules for the  $OR$  node to be of the form

$$[3_R, \_, 12_a, 3_R, \_, 12_b, 3_R, \_, 12_c, 3_R, \_, \_] \text{ or } [3_R, \_, 12_a, 3_R, \_, 12_b, 3_R, \_, \_, 3_R, \_, 12_c],$$


for some mapping of  $12_1$ ,  $12_2$ , and  $12_3$  onto  $12_a$ ,  $12_b$ , and  $12_c$ . Assume towards a contradiction that the  $12_O$  edge is blue, observing an immediate constraint violation when scheduling  $6_1$  and  $6_2$ . This demonstrates that the  $12_O$  edge cannot be blue and the schedule for the  $OR$  node must be of the form

$$[3_R, 6_a, 12_a, 3_R, 6_b, 12_b, 3_R, 6_a, 12_c, 3_R, 6_b, 12_d]$$

(likewise, for some mapping of  $12_1$ ,  $12_2$ ,  $12_3$ , and  $12_O$  onto  $12_a$ ,  $12_b$ ,  $12_c$ , and  $12_d$ ). Thus, if all input edges are blue, then the  $12_O$  edge must be either green or purple.  $\square$

## 5.8 Tension

$P_{d\varphi}$  is intended to have a valid schedule iff the underlying  $\varphi$  is satisfied by some assignment of values to its variables. Lemma 5.11 shows that the  $12_O$  output edge of an  $OR$  gadget can be blue iff the associated clause is satisfied by this assignment. This section introduces Tension gadgets (shown in Figure 8), which ensure that in any global schedule for  $P_{d\varphi}$  all  $12_O$  edges must be blue, and hence that all clauses must be True.



**Figure 8:** A gadget (left) which applies tension to four inputs, ensuring that either  $12_{O1}$ ,  $12_{O2}$ ,  $12_{O3}$ , and  $12_{O4}$  are scheduled in blue slots or no schedule can be found for  $T_e$ . To apply tension to less than 4 inputs, connect any unneeded inputs to its own male pendant node.

**Lemma 5.12 (Tension gadget schedules):** *Each Tension gadget  $T_e$  has a valid local schedule iff all  $12_O$  input edges are scheduled in blue slots.*  $\triangleleft$

**Proof:** The single node of each tension gadget  $T_e$ , as shown in Figure 8, has tasks  $[3_R, 6_G, 6_P, 12_{O1}, 12_{O2}, 12_{O3}, 12_{O4}]$  and density  $D = 1$ , so each task with frequency  $f$  must appear exactly once in every  $f$  day period. Further, the  $3_R$ ,  $6_G$ , and  $6_P$  edges are indirectly connected to the True Clock such that they must be scheduled in slots of their respective colours. Partial schedules for  $T_e$  nodes must therefore be of the form  $[3_R, \_, 6_G, 3_R, \_, 6_P]$ . All empty slots in schedules of this form are blue, so the remaining  $12_{O1}, 12_{O2}, 12_{O3}, 12_{O4}$  edges must be scheduled in blue slots, creating valid schedules of the form

$$[3_R, 12_{Oa}, 6_G, 3_R, 12_{Ob}, 6_P, 3_R, 12_{Oc}, 6_G, 3_R, 12_{Od}, 6_P],$$

for any mapping of  $12_{O1}$ ,  $12_{O2}$ ,  $12_{O3}$ , and  $12_{O4}$  onto  $12_{Oa}$ ,  $12_{Ob}$ ,  $12_{Oc}$ , and  $12_{Od}$ .  $\square$

## 5.9 Proof of Reduction

With these preparations we can, at long last, prove Lemma 5.1, and hence complete the proof of Theorem 1.3.

**Proof of Lemma 5.1 on page 17:** Consider a polycule  $\mathcal{P}_{d\varphi}$  built from some 3-CNF formula  $\varphi = c_1 \wedge \dots \wedge c_m$  over variables  $X = \{x_1, \dots, x_n\}$  using the algorithm defined by Definition 5.3.

First, suppose that there is a variable assignment  $v : X \rightarrow \{\text{True}, \text{False}\}$  which satisfies all clauses  $c_i \in C$ . We construct a slot-respecting schedule  $S$  for  $\mathcal{P}_{d\varphi}$  as follows:

- (a) At the variable gadget for each  $x_i \in X$ , schedule  $3_{iR}$  in red slots if  $v(x_i) = \text{True}$  and in the blue slots otherwise. This fixes a schedule for all edges in the variable layer.
- (b) Lemma 5.8 and Lemma 5.9 give schedules for duplication gadgets, whose number should be sufficient to cover the needs of all other gadgets.
- (c) As  $v$  satisfies all clauses  $c_i \in C$ , each clause has at least one literal ( $x_i$  or  $\overline{x_i}$ ) which evaluates to True. Hence, each clause has one input  $3_{iR}$  or  $3_{iB}$  edge which will be scheduled in red slots. By Lemma 5.11, there therefore exists a valid schedule for each  $OR$  gadget in the clause layer such that the  $12_O$  output edge of that gadget is scheduled in blue slots.
- (d) Lemma 5.12 shows that this implies a valid schedule for each Tension gadget.

These valid schedules for each layer of gadgets combine to form a global schedule  $S$  for  $\mathcal{P}_{d\varphi}$ .

Now, assume that we are given a schedule  $S$  for  $\mathcal{P}_{d\varphi}$ . By applying Lemma 5.5 to the True Clock, we can assign coloured slots to  $S$ . It then follows from the construction of  $\mathcal{P}_{d\varphi}$  that  $S$  must be slot-respecting. Set  $v(x_i) = \text{True}$  if  $S$  schedules  $3_{iR}$  in red slots and  $v(x_i) = \text{False}$  otherwise. Lemma 5.11 shows that for each  $c_i \in C$ , at least one input edge must be red for the output edge to be blue; Lemma 5.12 shows that all such output edges must be blue given the existence of  $S$ . Thus, all clauses  $c_i \in C$  evaluate to True under  $v$ , and the claim follows.

It remains to argue that our reduction can be realised in polynomial time. The size of the polycule  $\mathcal{P}_{d\varphi}$  is clearly polynomial in the size of the formula, leaving the gadgets themselves. Variable,  $OR$ , and Tension gadgets have both constant size and constant number. D6 duplicator gadgets are linear in size with respect to the number of  $6_P$  and  $6_G$  edges consumed by the system as a whole, bounded by the number of Tension gadgets, hence by the number of clauses. Blue D3 duplicator gadgets are linear with respect to the number of  $3_B$  edges consumed by  $OR$  gadgets, i.e. the number of clauses.  $3_R$  edges are consumed by both Tension gadgets and  $OR$  gadgets, but this still leaves the size of their D3 duplicator gadgets bounded by the number of clauses and hence by the size of the formula. Thus, it is easy to implement our reduction in polynomial time.  $\square$

\*   \*   \*   \*   \*

## 6 Conclusion and Open Problems

In this paper, we revisited the Polyamorous Scheduling problem. Via the observation that for an approximation, the hardest local Bamboo Garden Trimming instances around a single person dominate the solution, we found simple generalizations of existing algorithms for Bamboo Garden Trimming that yield constant-factor approximations for Poly Scheduling.

This paper opens up several avenues for future work. The most obvious open problem concerns algorithms with better approximation ratios and tighter lower bounds. We now know that a 5.24-approximation is possible and that  $(\frac{4}{3} - \varepsilon)$ -approximations are NP-hard, leaving a sizable gap. Apart from progress on this general question, a natural research direction is to look at restricted classes of polycules. For example, can we obtain better approximation guarantees on graphs with bounded maximum degree  $\Delta$  (beyond the simple results from [GSW24])? Can we exploit bipartite polycules towards a better approximation?

For the poly density of Decision Poly Scheduling, we showed that  $d \leq \frac{1}{4}$  implies feasibility; as for Pinwheel Scheduling, there are infeasible instances with density  $d = \frac{5}{6} + \varepsilon$  for any  $\varepsilon > 0$ . Can we say anything for densities between  $\frac{1}{4}$  and  $\frac{5}{6}$ ? In particular, we pose the following question.

**Question 6.1 (5/6-Poly-Density Question):** *Does every Decision Poly Scheduling instance with poly density  $d \leq \frac{5}{6}$  admit a schedule?* ◀

Given the tight connection between the problems, it is natural to ask whether Bamboo Garden Trimming can be arbitrarily well approximated, or whether there, too, a hardness-of-approximation result can be shown.

There are also several generalizations of Poly scheduling introduced by [GSW24] to which our results could be extended: Fungible Polyamorous Scheduling allows for each person  $P$  to have to  $s$  meetings per day, while Secure Polyamorous Scheduling allows for meetings to contain cliques of any size instead of only pairs. Can poly density be generalized to these problems? Can it identify classes of instances which do or do not permit solutions? And finally, how well do Reduce-Fastest and other known heuristic perform on them?

## References

- [AdBB<sup>+</sup>22] Peyman Afshani, Mark de Berg, Kevin Buchin, Jie Gao, Maarten Löffler, Amir Nayyeri, Benjamin Raichel, Rik Sarkar, Haotian Wang, and Hao-Tsung Yang. On cyclic solutions to the min-max latency multi-robot patrolling problem. In *38th International Symposium on Computational Geometry*, volume 224 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 2, 14. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2022. doi:10.4230/lipics.socg.2022.2.
- [Als15] Sultan Alshamrani. How reduce max algorithm behaves with symptoms appearance on virtual machines in clouds. In *2015 International Conference on Cloud Computing (ICCC)*, pages 1–4. IEEE, 2015.
- [BFCK19] Michael A. Bender, Martín Farach-Colton, and William Kuszmaul. Achieving optimal backlog in multi-processor cup games. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, page 1148–1157, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3313276.3316342.
- [BGL<sup>+</sup>22] Davide Bilò, Luciano Gualà, Stefano Leucci, Guido Proietti, and Giacomo Scornavacca. Cutting bamboo down to size. *Theoret. Comput. Sci.*, 909:54–67, 2022. doi:10.1016/j.tcs.2022.01.027.
- [BvEJ<sup>+</sup>22] Thomas Bosman, Martijn van Ee, Yang Jiao, Alberto Marchetti-Spaccamela, R. Ravi, and Leen Stougie. Approximation algorithms for replenishment problems with fixed turnover times. *Algorithmica*, 84(9):2597–2621, 2022. doi:10.1007/s00453-022-00974-4.
- [CC92] Mee Yee Chan and Francis Y. L. Chin. General schedulers for the pinwheel problem based on double-integer reduction. *IEEE Transactions on Computers*, 41(06):755–768, 1992.
- [CC93] M. Y. Chan and Francis Chin. Schedulers for larger classes of pinwheel instances. *Algorithmica*, 9(5):425–462, 1993. doi:10.1007/BF01187034.

- [DDSN19] Mattia D’Emidio, Gabriele Di Stefano, and Alfredo Navarra. Bamboo garden trimming problem: priority schedulings. *Algorithms (Basel)*, 12(4):Paper No. 74, 23, 2019. doi:10.3390/a12040074.
- [Din20] Wei Ding. A branch-and-cut approach to examining the maximum density guarantee for pinwheel schedulability of low-dimensional vectors. *Real-Time Systems*, 56:293–314, 2020.
- [FC05] Eugene A. Feinberg and Michael T. Curry. Generalized pinwheel problem. *Math. Methods Oper. Res.*, 62(1):99–122, 2005. doi:10.1007/s00186-005-0443-4.
- [FL02] Fishburn and Lagarias. Pinwheel scheduling: Achievable densities. *Algorithmica*, 34:14–38, 2002.
- [GJK<sup>+</sup>24] Leszek Gąsieniec, Tomasz Jurdziński, Ralf Klasing, Christos Levcopoulos, Andrzej Lingas, Jie Min, and Tomasz Radzik. Perpetual maintenance of machines with different urgency requirements. *Journal of Computer and System Sciences*, 139:103476, 2024.
- [GKL<sup>+</sup>17] Leszek Gąsieniec, Ralf Klasing, Christos Levcopoulos, Andrzej Lingas, Jie Min, and Tomasz Radzik. Bamboo garden trimming problem (perpetual maintenance of machines with different attendance urgency factors). In *SOFSEM 2017: theory and practice of computer science*, volume 10139 of *Lecture Notes in Comput. Sci.*, pages 229–240. Springer, Cham, 2017.
- [GSW22] Leszek Gąsieniec, Benjamin Smith, and Sebastian Wild. Towards the 5/6-density conjecture of pinwheel scheduling. In *2022 Proceedings of the Symposium on Algorithm Engineering and Experiments (ALENEX)*, pages 91–103. SIAM, 2022.
- [GSW24] Leszek Gąsieniec, Benjamin Smith, and Sebastian Wild. Polyamorous scheduling. In *12th International Conference on Fun with Algorithms*, volume 291 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. No. 15, 18. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2024. doi:10.4230/lipics.fun.2024.15.
- [HL92] C-C Han and K-J Lin. Scheduling distance-constrained real-time tasks. In *Proceedings Real-Time Systems Symposium*, pages 300–301. IEEE Computer Society, 1992.
- [HMR<sup>+</sup>89] Robert Holte, Aloysius Mok, Louis Rosier, Igor Tulchinsky, and Donald Varvel. The pinwheel: A real-time scheduling problem. In *Proceedings of the 22nd Hawaii International Conference of System Science*, pages 693–702, 1989.
- [HRTV92] Robert Holte, Louis Rosier, Igor Tulchinsky, and Donald Varvel. Pinwheel scheduling with two distinct numbers. *Theoret. Comput. Sci.*, 100(1):105–135, 1992. doi:10.1016/0304-3975(92)90365-M.
- [HvS23] Felix Höhne and Rob van Stee. A 10/7-approximation for discrete bamboo garden trimming and continuous trimming on star graphs. In *Approximation, randomization, and combinatorial optimization. Algorithms and techniques*, volume 275 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. No. 16, 19. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2023. doi:10.4230/lipics.approx/random.2023.16.

- [Kaw24] Akitoshi Kawamura. Proof of the Density Threshold Conjecture for Pinwheel Scheduling. In *STOC'24—Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1816–1819. ACM, New York, 2024. doi:10.1145/3618260.3649757.
- [KS20] Akitoshi Kawamura and Makoto Soejima. Simple strategies versus optimal schedules in multi-agent patrolling. *Theoret. Comput. Sci.*, 839:195–206, 2020. doi:10.1016/j.tcs.2020.07.037.
- [Kus20] William Kuszmaul. Achieving optimal backlog in the vanilla multi-processor cup game. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1558–1577, 2020. doi:10.1137/1.9781611975994.96.
- [Kus22] John Kuszmaul. Bamboo trimming revisited: Simple algorithms can do well too. In *Proceedings of the 34th ACM Symposium on Parallelism in Algorithms and Architectures*, pages 411–417, 2022.
- [LL73] C. L. Liu and James W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *J. Assoc. Comput. Mach.*, 20:46–61, 1973. doi:10.1145/321738.321743.
- [LL97] Shun-Shii Lin and Kwei-Jay Lin. A pinwheel scheduler for three distinct numbers with a tight schedulability bound. *Algorithmica*, 19(4):411–426, 1997. doi:10.1007/PL00009181.
- [Sha49] Claude E Shannon. A theorem on coloring the lines of a network. *Journal of Mathematics and Physics*, 28(1-4):148–152, 1949.
- [Viz64] V. G. Vizing. On an estimate of the chromatic class of a  $p$ -graph. *Diskret. Analiz*, pages 25–30, 1964.