

AN OPTIMAL LOCAL APPROXIMATION ALGORITHM FOR MAX-MIN LINEAR PROGRAMS

PATRIK FLORÉEN AND JOEL KAASINEN AND PETTERI KASKI AND JUKKA SUOMELA

Helsinki Institute for Information Technology HIIT
Helsinki University of Technology and University of Helsinki
P.O. Box 68, FI-00014 University of Helsinki, Finland
E-mail addresses: `patrik.floreen@cs.helsinki.fi`, `joel.kaasinen@cs.helsinki.fi`,
`petteri.kaski@cs.helsinki.fi`, `jukka.suomela@cs.helsinki.fi`

ABSTRACT. We present a local algorithm (constant-time distributed algorithm) for approximating max-min LPs. The objective is to maximise ω subject to $A\mathbf{x} \leq \mathbf{1}$, $C\mathbf{x} \geq \omega\mathbf{1}$, and $\mathbf{x} \geq \mathbf{0}$ for nonnegative matrices A and C . The approximation ratio of our algorithm is the best possible for any local algorithm; there is a matching unconditional lower bound.

1. Introduction

In a *max-min linear program* (max-min LP), the objective is to

$$\begin{aligned} &\text{maximise} && \min_{k \in K} \mathbf{c}_k \mathbf{x} \\ &\text{subject to} && A\mathbf{x} \leq \mathbf{1}, \\ &&& \mathbf{x} \geq \mathbf{0} \end{aligned} \tag{1.1}$$

or, equivalently, to maximise ω subject to $A\mathbf{x} \leq \mathbf{1}$, $C\mathbf{x} \geq \omega\mathbf{1}$, and $\mathbf{x} \geq \mathbf{0}$. The matrices A and C are nonnegative and sparse: each row \mathbf{a}_i of A has at most Δ_I positive elements, and each row \mathbf{c}_k of C has at most Δ_K positive elements; here Δ_I and Δ_K are constants.

Our work studies *local algorithms* [9, 10], that is, distributed algorithms that complete in *constant time* (constant number of synchronous communication rounds), independent of the size of the network. We assume that there is a network with one node $v \in V$ for each variable x_v , one node $i \in I$ for each constraint $\mathbf{a}_i \mathbf{x} \leq 1$, and one node $k \in K$ for each objective $\mathbf{c}_k \mathbf{x}$. Nodes $v \in V$ and $i \in I$ are adjacent if a_{iv} is positive, and nodes $v \in V$ and $k \in K$ are adjacent if c_{kv} is positive. In a local algorithm for max-min LPs, the node $v \in V$ must choose the value x_v based on the information that is available within its constant-radius neighbourhood in the network.

Max-min LPs are a generalisation of packing LPs. Direct applications of max-min LPs include various tasks of fair resource allocation, such as fair bandwidth allocation in a communication network and balanced data gathering in a wireless sensor network. An algorithm for approximating max-min LPs allows one to solve approximate mixed packing

Key words and phrases: distributed algorithms, linear programs, local algorithms, networks.

This research was supported in part by the Academy of Finland, Grants 116547 and 117499, and by Helsinki Graduate School in Computer Science and Engineering (Hecse).

and covering problems [12] as well. Special cases of mixed packing and covering LPs include finding an (approximate) solution to a nonnegative system of linear equations.

1.1. Prior work

For any $\epsilon > 0$, there exists a local $(1 + \epsilon)$ -approximation algorithm for packing and covering LPs, assuming a bounded-degree graph and bounded coefficients [7, 8]. However, this is not the case with max-min LPs:

Theorem 1.1 (Floréen et al. [4]). *For any $\Delta_I \geq 2$ and $\Delta_K \geq 2$, there exists no local approximation algorithm for the max-min LP problem with the approximation ratio $\Delta_I(1 - 1/\Delta_K)$.*

This theorem holds even in the following special cases: (i) *bipartite* max-min LPs, where each column of A and each column of C contains only one nonzero element, and (ii) 0/1 max-min LPs, where each element of A and C is either 0 or 1.

The approximation ratio $\Delta_I(1 - 1/\Delta_K) + \epsilon$ for any positive ϵ can be achieved by a local algorithm for bipartite problems [4] and for a class of 0/1 problems [3]. However, for general max-min LPs, the best known local algorithm has been the *safe algorithm* that achieves the factor Δ_I approximation [5, 11].

1.2. Contribution

The following theorem summarises the main contribution of this paper.

Theorem 1.2. *For any $\Delta_I \geq 2$, $\Delta_K \geq 2$, and $\epsilon > 0$, there exists a local approximation algorithm for the max-min LP problem with the approximation ratio $\Delta_I(1 - 1/\Delta_K) + \epsilon$.*

By Theorem 1.1, the algorithm in Theorem 1.2 is optimal in the sense that no local algorithm can achieve a better approximation ratio. For the case $\Delta_I = \Delta_K = 2$, the approximation ratio is $1 + \epsilon$, that is, there is a local approximation scheme.

1.3. Definitions

We now give a detailed definition of a max-min LP in a distributed setting. Let $\mathcal{G} = (V \cup I \cup K, E)$ be a bipartite, undirected communication graph. The nodes $v \in V$ are called *agents*, the nodes $i \in I$ are called *constraints*, and the nodes $k \in K$ are called *objectives*; the sets V , I , and K are disjoint. Edges $e \in E$ are of the form $e = \{v, i\}$ or $e = \{v, k\}$ where $v \in V$, $i \in I$, and $k \in K$.

Let $V_i = \{v \in V : \{v, i\} \in E\}$, $V_k = \{v \in V : \{v, k\} \in E\}$, $I_v = \{i \in I : \{v, i\} \in E\}$, and $K_v = \{k \in K : \{v, k\} \in E\}$ for all $i \in I$, $k \in K$, and $v \in V$. We assume that $|V_i| \leq \Delta_I$ and $|V_k| \leq \Delta_K$ for all $i \in I$ and $k \in K$ for some constants Δ_I and Δ_K .

A *max-min linear program* associated with \mathcal{G} is defined as follows. Associate a variable x_v with each agent $v \in V$, associate a coefficient $a_{iv} > 0$ with each edge $\{i, v\} \in E$, $i \in I$, $v \in V$, and associate a coefficient $c_{kv} > 0$ with each edge $\{k, v\} \in E$, $k \in K$, $v \in V$. The task is to

$$\begin{aligned} & \text{maximise} && \omega(\mathbf{x}) = \min_{k \in K} \sum_{v \in V_k} c_{kv} x_v \\ & \text{subject to} && \sum_{v \in V_i} a_{iv} x_v \leq 1, \quad \forall i \in I, \\ & && x_v \geq 0, \quad \forall v \in V. \end{aligned} \tag{1.2}$$

The *local input* of an agent $v \in V$ consists of the sets I_v and K_v and the coefficients a_{iv}, c_{kv} for all $i \in I_v, k \in K_v$. The local input of a constraint $i \in I$ consists of V_i , and the

local input of an objective $k \in K$ consists of V_k . In a local algorithm, there is a constant D such that each agent $v \in V$ chooses the output x_v based on the local inputs of the nodes within distance D (in number of edges) from v in the communication graph \mathcal{G} ; the constant D is the *local horizon* of the algorithm.

Theorem 1.1 holds even if we assume that each node of the graph \mathcal{G} has a unique identifier. We show that Theorem 1.2 holds even if we do not have unique identifiers. We merely assume *port numbering* [1]: each edge $\{s, t\}$ in \mathcal{G} has two natural numbers associated with it, the port number in s and the port number in t .

1.4. Overview of the algorithm

We begin in §2 by reminding that, in the port numbering model, we can without loss of generality focus on the case where the communication graph \mathcal{G} is a (countably infinite but locally finite) tree [1, 4].

In §3, we present a series of local transformations that simplify the structure of the problem. We show that with the objective of establishing Theorem 1.2, it is sufficient to focus on the special case where $|V_i| = 2$, $|V_k| \geq 2$, $|K_v| = 1$, $|I_v| \geq 1$, and $c_{ku} = 1$ for all $i \in I$, $k \in K$, $v \in V$, $u \in V_k$. Figure 1 shows an example of a communication graph \mathcal{G} after the local transformations.

In §4, we present a local algorithm for this special case. In §5, we prove that the output of the algorithm is a factor $2(1 - 1/\Delta_K) + \epsilon'$ approximation. Put together, we have a local, factor $\Delta_I(1 - 1/\Delta_K) + \epsilon$ approximation for general max-min LPs, for any $\epsilon > 0$.

The intuition behind the algorithm in §4 is best understood if we study its analysis in §5. In the analysis, it is convenient to assume that we have assigned a one-dimensional coordinate, *layer*, to each node of the tree \mathcal{G} ; see Figure 1 for an example. When we assign the layers, we also partition the agents into *up-agents* and *down-agents*. We have alternatingly layers of up-agents, constraints, down-agents, and objectives. Each objective has exactly one adjacent up-agent “above” it, and at least one adjacent down-agent “below” it; hence the names “up” and “down”.

We now assign every R^{th} layer of objectives to be *passive*, including the adjacent agents that set $x_v = 0$. Each agent $v \in V$ computes an upper bound t_v of the optimum; see §4.2. Then we construct a solution for the active layers in a greedy manner, starting with a layer of passive up-agents and propagating information upwards until we reach the next layer of passive down-agents; see §4.3.

This is not yet an approximation for the original problem: while most objectives perform at least as well as in the global optimum, the passive objectives have utility 0. By applying ideas from the shifting strategy [2, 6], we could consider R possible choices for the locations of the passive layers. Then we could take averages over these to obtain a solution \mathbf{y} . In §5.2 we show that \mathbf{y} would indeed be a factor $R/(R - 1)$ approximation.

There is one difficulty, however. We cannot assign the layers by a local algorithm in a globally consistent manner; in particular, we do not know whether a given agent is a down-agent or an up-agent. To overcome this, we consider both possible roles for each agent, up and down. For both roles, we compute a candidate solution by applying the shifting strategy. Finally we take the average of both candidate solutions. This is the essence of (4.13).

In §5, we prove that this local approach yields a globally feasible solution, and the solution is within factor $2(1 - 1/\Delta_K) + \epsilon'$ of the optimum. The constant $\epsilon' > 0$ can be made arbitrarily small by choosing a sufficiently large R .




Figure 1: The graph \mathcal{G} (after the local transformations in §3) and the layers (see §5). We have chosen $R = 3$ and hence $r = 1$. Thick lines highlight the tree \mathcal{A}_u (see §4.1). There are several ways to choose the layers; nevertheless, if u is an up-agent on layer -1 then the levels in \mathcal{A}_u necessarily coincide with the layers.

Yet another basic hurdle implicitly overcome in the proof of Theorem 1.2 stems from *underconstrained instances*: if there are several equally good solutions, one needs to choose between them in a globally consistent manner. Our definition of the values g^+ and g^- in (4.8)–(4.10) addresses this by focusing on a particular extreme point. Each layer of down-agents chooses as large values g^+ as possible, without violating the constraints “below” them. Each layer of up-agents chooses as small values g^- as possible, as long as the objectives “below” them meet the smoothed upper bounds s_v .

2. Port numbering and unfolding

Our algorithm does not need to use any node identifiers; port numbering is sufficient. In the port numbering model, a local algorithm cannot distinguish between a short cycle and an infinitely long path. We can exploit this limitation to simplify the description of our local algorithm: we can assume that we have *unfolded* all cycles of the graph \mathcal{G} [1, 4].

The definition of unfolding requires some preliminaries. A *walk* of length ℓ in a graph \mathcal{G} is a nonempty tuple $(u_0, e_1, u_1, e_2, u_2, \dots, e_\ell, u_\ell)$ of alternating nodes and edges in \mathcal{G} such that, for all $j = 1, 2, \dots, \ell$, the edge e_j joins the nodes u_{j-1} and u_j . The walk is said to *start* at u_0 and *end* at u_ℓ . A walk is *non-backtracking* if $e_{j-1} \neq e_j$ holds for all $j = 1, 2, \dots, \ell$. A *path* is a walk with no repeated nodes.

Let \mathcal{G} be a finite connected graph and let r be a node of \mathcal{G} . The *unfolding* of \mathcal{G} rooted at r is the undirected simple graph \mathcal{G}' obtained as follows. The node set of \mathcal{G}' is the set of all non-backtracking walks in \mathcal{G} that start at r . Two nodes of \mathcal{G}' are joined by an edge iff one can be obtained from the other by appending exactly one edge and one node of \mathcal{G} .

We associate with each node of \mathcal{G}' a *parent* node of \mathcal{G} , namely the end-node of the walk. We also associate with each edge of \mathcal{G}' a *parent* edge of \mathcal{G} , namely the appended edge.

Remarks.

- (1) The unfolding \mathcal{G}' is a tree.
- (2) The unfolding \mathcal{G}' is finite iff \mathcal{G} is a tree; otherwise \mathcal{G}' is countably infinite.
- (3) Any two unfoldings of \mathcal{G} rooted at different nodes are isomorphic. In what follows we refer to “the” unfolding of \mathcal{G} without specifying a particular root node.
- (4) Assuming that the graph \mathcal{G} has port numbers associated with the ends of its edges, the unfolding \mathcal{G}' inherits the port numbering from the parent edges.
- (5) Assuming that the graph \mathcal{G} has a max-min LP associated with it, the max-min LP associated with the unfolding \mathcal{G}' is defined by inheritance from the parent nodes and edges. In particular, the type of each node (agent, constraint, objective) is the type of the parent node, and the coefficients associated with the edges (a_{iv} , c_{kv}) are inherited from the parent edges.
- (6) Any two nodes of \mathcal{G}' with the same parent are related by an automorphism of \mathcal{G}' . In particular, any deterministic local algorithm in the port numbering model must give the same output on any two nodes with the same parent. Any locally computable feasible solution of the max-min LP associated with \mathcal{G}' defines a feasible solution of the max-min LP associated with \mathcal{G} , with the same utility.
- (7) Any feasible solution of the max-min LP associated with \mathcal{G} defines, by inheritance, a feasible solution of the max-min LP associated with \mathcal{G}' , with the same utility.
- (8) Any locally computed feasible solution of \mathcal{G}' with utility at least $1/\alpha$ times the utility of any feasible solution of \mathcal{G} yields an α -approximation of the optimum of \mathcal{G} .

3. Local transformations

Consider an arbitrary max-min LP associated with the graph \mathcal{G} . In this section we carry out a sequence of locally computable transformations, with the goal of arriving at a more structured max-min LP. The transformations are applied in the order of presentation, from §3.1 to §3.5. We describe each individual transformation in three parts:

- (1) A description of the transformation.
- (2) Mapping a solution of the transformed instance back to the original instance.
- (3) Implications to approximability. We write $\omega(\cdot)$ for the utility of the original instance and $\omega'(\cdot)$ for the utility of the transformed instance.

Appendix A.1 presents the implementation of these transformations in the port numbering model. Figure 2 illustrates the transformations that modify the communication graph \mathcal{G} .

To avoid degenerate cases, we assume that each constraint and objective is adjacent to at least one agent, and every agent is adjacent to at least one constraint and at least one objective, that is, $|V_i| \geq 1$, $|V_k| \geq 1$, $|K_v| \geq 1$, and $|I_v| \geq 1$ for all $i \in I$, $k \in K$, $v \in V$. Indeed, isolated constraints can be deleted, isolated objectives force the optimum of (1.2) to zero, non-contributing agents can be set to zero, and unconstrained agents can be set to $+\infty$. Furthermore, we assume that \mathcal{G} is connected, as we can handle each connected component independently.




Figure 2: Local transformations in §3.1–§3.4.

3.1. Augmenting singleton constraints – forcing $|V_i| \geq 2$

The transformation. For each constraint $i \in I$ with $|V_i| = 1$, introduce three new agents, s , t , and u , two new objectives, h and ℓ , and one new constraint, j . Let $v \in V_i$ be the original agent adjacent to i , and let $k \in K_v$ be an objective adjacent to v . Set $a_{is} = a_{jt} = a_{ju} = 1$, $c_{hs} = c_{\ell s} = 1$, and $c_{ht} = c_{\ell u} = 2 \sum_{w \in V_k} c_{kw} \min_{i \in I_w} a_{iw}^{-1}$.

Mapping back. Let \mathbf{x}' be any feasible solution of the transformed instance. We map this back by setting $x_v = x'_v$ for all original agents $v \in V$.

Approximation ratio. Observe that we can always set $x'_s = 0$ and $x'_t = x'_u = 1/2$ without decreasing the objective value $\omega'(\mathbf{x}')$. Thus the optima of the original and transformed instances coincide, and any approximation ratio is preserved.

3.2. Reducing the degree of constraints – forcing $|V_i| = 2$

The transformation [3]. Replace each constraint $i \in I$ with $|V_i| > 2$ by the $\binom{|V_i|}{2}$ constraints

$$a_{iu}x_u + a_{iv}x_v \leq 1, \quad u, v \in V_i, \quad u < v. \quad (3.1)$$

Mapping back. Let \mathbf{x}' be an arbitrary feasible solution of the transformed instance. We map this back to a feasible solution \mathbf{x} of the original instance by setting

$$x_v = \frac{2x'_v}{\max_{i \in I_v} |V_i|}, \quad v \in V. \quad (3.2)$$

To verify that \mathbf{x} is feasible, consider an arbitrary original constraint $i \in I$. By the previous transformation, we have $|V_i| \geq 2$. Taking the sum over all the constraints (3.1) replacing i , or, if $|V_i| = 2$, considering the original constraint, we have

$$\sum_{v \in V_i} (|V_i| - 1) a_{iv} x'_v \leq |V_i| (|V_i| - 1) / 2.$$

By (3.2) we thus have

$$\sum_{v \in V_i} a_{iv} x_v \leq \sum_{v \in V_i} 2a_{iv} x'_v / |V_i| \leq 1.$$

Approximation ratio. Because the objectives are unchanged in the transformation, we have $\omega(\mathbf{x}) \geq 2\omega'(\mathbf{x}')/\Delta_I$ by linearity and (3.2). Furthermore, an optimal solution of the original instance is a feasible solution of the transformed instance. Therefore, if \mathbf{x}' is an α -approximate solution, then \mathbf{x} is a $\alpha\Delta_I/2$ -approximate solution.

3.3. Associating a unique objective with each agent – forcing $|K_v| = 1$

The transformation. For each agent $v \in V$ with $|K_v| > 1$, replace v with $|K_v|$ copies of v as follows. Associate each copy of v with a unique objective in K_v . Replace each constraint adjacent to v with $|K_v|$ copies of the constraint, with v replaced by a unique copy in each constraint. The coefficients are unchanged.

Mapping back. Let \mathbf{x}' be a feasible solution of the transformed instance. By symmetry we can assume that all copies u of v have the same value x'_u without decreasing the objective value $\omega'(\mathbf{x}')$; indeed, if the values are different, just set all copies to the maximum value. Mapping back is done simply by identifying the copies back to the original.

Approximation ratio. Preserved. The optima of the original and the transformed instance coincide.

3.4. Augmenting singleton objectives – forcing $|V_k| \geq 2$

The transformation. For each objective $k \in K$ with $|V_k| = 1$, let v be the unique agent adjacent to k . Replace v with two copies, t and u , and replace each constraint adjacent to v with two copies of the constraint, one containing t in place of v , and the other containing u in place of v . Let $c_{kt} = c_{ku} = c_{kv}/2$. The coefficients are otherwise unchanged.

Mapping back. Let \mathbf{x}' be a feasible solution of the transformed instance. By symmetry we can assume that the copies of v have the same value $x'_t = x'_u$ without decreasing the objective value $\omega'(\mathbf{x}')$; indeed, if the values are different, just set all copies to the maximum value. Mapping back is done simply by identifying the copies back to the original.

Approximation ratio. Preserved. The optima of the original and the transformed instance coincide.

3.5. Normalising coefficients – forcing $c_{kv} = 1$ for adjacent $k \in K$ and $v \in V$

For each $v \in V$, let $k(v)$ be the unique objective in K_v .

The transformation. For each $v \in V$, $i \in I$, $k \in K$, divide a_{iv} and c_{kv} by $c_{k(v)v}$.

Mapping back. For each $v \in V$, multiply x_v by $c_{k(v)v}$.

Approximation ratio. Preserved.

4. Local approximation algorithm

Throughout this section we consider a max-min LP associated with a bipartite graph $\mathcal{G} = (V \cup I \cup K, E)$ with these properties that follow from §2 and §3:

- (1) the graph \mathcal{G} is an unfolding of a finite graph,
- (2) every agent $v \in V$ is adjacent to exactly one objective and at least one constraint,
- (3) every constraint $i \in I$ is adjacent to exactly two agents,
- (4) every objective $k \in K$ is adjacent to at least two agents,
- (5) for adjacent $k \in K$ and $v \in V$, the coefficient c_{kv} equals 1.

It follows that \mathcal{G} is countably infinite.

Recall that $k(v)$ is the unique objective adjacent to $v \in V$. Let $N(v) = V_{k(v)} \setminus \{v\}$ be the set of other agents adjacent to this objective. For a constraint $i \in I$ and an agent $v \in V_i$, denote by $n(v, i)$ the unique agent other than v in V_i . See Figure 1 for an illustration.

Let $R = 2, 3, \dots$ be a fixed parameter that will determine the local horizon and the approximation ratio. Let $r = R - 2$.

4.1. An upper bound via alternating trees

A walk $W = (u_0, e_1, u_1, \dots, e_\ell, u_\ell)$ in \mathcal{G} is *alternating* if (i) for all $1 \leq j < j' \leq \ell$ with $u_j \in K$ and $u_{j'} \in K$, there exists a $j < j'' < j'$ with $u_{j''} \in I$; and (ii) for all $1 \leq j < j' \leq \ell$ with $u_j \in I$ and $u_{j'} \in I$, there exists a $j < j'' < j'$ with $u_{j''} \in K$.

Let $u \in V$ be an arbitrary agent, and consider the subgraph \mathcal{A}_u of \mathcal{G} induced by the nodes reachable via alternating paths starting at u that (i) traverse the constraint $k(u)$ and have length at most $4r + 3$; or (ii) have length at most 1. The *level* of a node of \mathcal{A}_u is its distance to $k(u)$, with the exception of u , which we define to have level -1 , and the constraints adjacent to u , which we define to have level -2 . See Figure 1 for an illustration.

Associate with \mathcal{A}_u a max-min LP by restriction from the max-min LP associated with \mathcal{G} .

Lemma 4.1. *The graph \mathcal{A}_u is a finite tree. Moreover,*

- (1) *every objective in \mathcal{A}_u is at level 0 (mod 4),*
- (2) *every agent in \mathcal{A}_u is at level either 1 (mod 4) or 3 (mod 4),*
- (3) *every constraint in \mathcal{A}_u is at level 2 (mod 4),*
- (4) *the leaves of \mathcal{A}_u are constraints at levels -2 and $4r + 2$,*
- (5) *for any objective k in \mathcal{A}_u and any agent v adjacent to k in \mathcal{G} , the agent v , occurs in \mathcal{A}_u and is adjacent to k in \mathcal{A}_u .*

Proof. By induction on R using the assumptions on the structure of \mathcal{G} . ■

Lemma 4.2. *The optimum value of the max-min LP associated with \mathcal{A}_u is an upper bound on the value of any feasible solution of the max-min LP associated with \mathcal{G} .*

Proof. Because \mathcal{A}_u is finite, the optimum is well defined. It suffices to show that any feasible solution of \mathcal{G} is (by restriction) a feasible solution of \mathcal{A}_u . This follows from Lemma 4.1: the objectives in \mathcal{A}_u are identical to those in \mathcal{G} , and the constraints on variables in \mathcal{A}_u are either identical or relaxed from \mathcal{G} (at leaves or non-alternating constraints). ■

4.2. The optimum of \mathcal{A}_u

The tree structure of \mathcal{A}_u enables a recursive characterisation of the optimum that proceeds level-wise towards u . Denote by $L(u, \ell)$ the set of all nodes at level ℓ in \mathcal{A}_u .

Let t_u be the maximum value $\omega \geq 0$ such that for all $d = 0, 1, \dots, r$ it holds that the values

$$f_{u,v,0}^+(\omega) = \min_{i \in I_v} a_{iv}^{-1}, \quad v \in L(u, 4r + 1), \quad (4.1)$$

$$f_{u,v,d}^-(\omega) = \max(0, \omega - \sum_{w \in N(v)} f_{u,w,d}^+(\omega)), \quad v \in L(u, 4(r - d) - 1), \quad (4.2)$$

$$f_{u,v,d}^+(\omega) = \min_{i \in I_v} a_{iv}^{-1} (1 - a_{i,n(v,i)} f_{u,n(v,i),d-1}^-(\omega)), \quad d \geq 1, \quad v \in L(u, 4(r - d) + 1) \quad (4.3)$$

satisfy the constraints

$$f_{u,v,d}^+(\omega) \geq 0, \quad 0 \leq d \leq r, \quad v \in L(u, 4(r - d) + 1), \quad (4.4)$$

$$f_{u,u,r}^-(\omega) \leq \min_{i \in I_u} a_{iu}^{-1}. \quad (4.5)$$

Note that the maximum exists because (4.2) and (4.3) can be expressed using linear inequalities (by introducing additional variables), and $\omega = 0$ is a feasible value. In a practical

implementation of our algorithm, we do not need to invoke an LP solver; a simple binary search for an *approximation* of t_u is sufficient.

Lemma 4.3. *Let \mathbf{x} be any feasible solution that achieves the objective value ω in the max-min LP associated with \mathcal{A}_u . Then, for all $d = 0, 1, \dots, r$ it holds that*

$$0 \leq x_v \leq f_{u,v,d}^+(\omega) \quad \text{for all } v \in L(u, 4(r-d) + 1); \text{ and} \quad (4.6)$$

$$f_{u,v,d}^-(\omega) \leq x_v \leq \min_{i \in I_v} a_{iv}^{-1} \quad \text{for all } v \in L(u, 4(r-d) - 1). \quad (4.7)$$

In particular, t_u is the optimum utility of the max-min LP associated with \mathcal{A}_u .

Proof. By induction on d and in order of evaluation of the recursive steps (4.2) and (4.3); see Appendix A.2 for a full proof. ■

In what follows we use the shorthand notation $f_{u,v,d}^+ = f_{u,v,d}^+(t_u)$ and $f_{u,v,d}^- = f_{u,v,d}^-(t_u)$.

4.3. Smoothing

For each agent $v \in V$ in \mathcal{G} , let s_v be the minimum of the values t_u over all agents $u \in V$ at distance at most $4r + 2$ from v in \mathcal{G} . For all $v \in V$ in \mathcal{G} and all $d = 0, 1, \dots, r$, define

$$g_{v,0}^+ = \min_{i \in I_v} a_{iv}^{-1}, \quad (4.8)$$

$$g_{v,d}^- = \max(0, s_v - \sum_{w \in N(v)} g_{w,d}^+), \quad (4.9)$$

$$g_{v,d}^+ = \min_{i \in I_v} a_{iv}^{-1} (1 - a_{i,n(v,i)} g_{n(v,i),d-1}^-), \quad d \geq 1. \quad (4.10)$$

Lemma 4.4. *For all $u \in V$ and all $d = 0, 1, \dots, r$ it holds that*

$$g_{v,d}^- \leq f_{u,v,d}^- \quad \text{for all } v \in L(u, 4(r-d) - 1), \quad (4.11)$$

$$f_{u,v,d}^+ \leq g_{v,d}^+ \quad \text{for all } v \in L(u, 4(r-d) + 1). \quad (4.12)$$

Proof. By induction on d , using the definition of s_v ; see Appendix A.3 for a full proof. ■

Lemma 4.5. *For all $v \in V$ it holds that $g_{v,r}^+ \geq 0$ and $g_{v,r}^- \leq \min_{i \in I_v} a_{iv}^{-1}$.*

Proof. Let $v \in V$ be arbitrary and choose $u \in N(v)$; such a u exists because every objective is adjacent to at least two agents. We have $g_{v,r}^+ \geq f_{u,v,r}^+ \geq 0$ by (4.12) and (4.4). Similarly, let $u = v$ to obtain

$$g_{v,r}^- \leq f_{v,v,r}^- \leq \min_{i \in I_v} a_{iv}^{-1}$$

by (4.11) and (4.5). ■

Lemma 4.6. *For all $v \in V$ and $d = 1, 2, \dots, r$ it holds that $g_{v,d-1}^- \leq g_{v,d}^-$ and $g_{v,d}^+ \leq g_{v,d-1}^+$.*

Proof. By induction on d ; see Appendix A.4 for a full proof. ■

Lemma 4.7. *For all $v \in V$ and $d = 0, 1, \dots, r$ it holds that $g_{v,d}^+ \geq 0$.*

Proof. By Lemma 4.5 and Lemma 4.6. ■




Figure 3: The weights used to assign the layers.

Finally, each agent v outputs the value

$$x_v = \frac{1}{2R} \sum_{d=0}^r (g_{v,d}^+ + g_{v,d}^-). \quad (4.13)$$

This completes the description of the algorithm. The algorithm is local, with the local horizon $\Theta(R)$. We proceed to show the vector \mathbf{x} is a feasible solution, and within factor $2(1 - 1/\Delta_K) + \epsilon'$ of the optimum.

5. Analysis

We start by partitioning the set of agents V into *up-agents* and *down-agents* such that (i) every constraint is adjacent to exactly one up-agent and exactly one down-agent; and (ii) every objective is adjacent to exactly one up-agent.

Associate an integer *layer* to each node of \mathcal{G} as follows. First, fix an arbitrary objective $k \in K$ to be at layer 0. Then, determine the layer of every other node u by considering the unique directed path connecting k to u . The layer of u is determined by taking the sum of the weights of the directed edges in the path, where the weights are displayed in Figure 3. Figure 1 displays an example of a partition into up- and down-agents together with an assignment of layers.

Lemma 5.1. *The layers of the nodes of \mathcal{G} satisfy the following four properties:*

- (1) *every objective has layer 0 (mod 4),*
- (2) *every down-agent has layer 1 (mod 4),*
- (3) *every constraint has layer 2 (mod 4),*
- (4) *every up-agent has layer 3 (mod 4).*

Proof. Immediate from Figure 3. ■

5.1. Shifting strategy

Let $j = 0, 1, \dots, R-1$ be a shift parameter. For each agent $v \in V$, represent the layer of v uniquely as $4(Rc + j) + 4d + e$ for integers c, d, e with $0 \leq d \leq R-1$ and $e \in \{-1, 1\}$. Recall that $r = R-2$. Associate with v the value

$$y_v(j) = \begin{cases} 0 & \text{if } d = R-1; \\ g_{v,r-d}^- & \text{if } d \leq R-2 \text{ and } e = -1; \\ g_{v,r-d}^+ & \text{if } d \leq R-2 \text{ and } e = 1. \end{cases} \quad (5.1)$$

Observe that a down-agent has $e = 1$ and an up-agent has $e = -1$, regardless of j .

For an objective k and a vector \mathbf{z} indexed by the agents, let $\omega_k(\mathbf{z}) = \sum_{v \in V_k} z_v$.

Lemma 5.2. *The vector $\mathbf{y}(j)$ is a feasible solution of the max-min LP associated with \mathcal{G} . For every objective $k \in K$, it holds that $\omega_k(\mathbf{y}(j)) = 0$ if k is at layer $4j - 4 \pmod{4R}$, and $\omega_k(\mathbf{y}(j)) \geq \min_{v \in V_k} s_v$ otherwise.*

Proof. From (4.8), (4.9), (4.10), and (5.1); see Appendix A.5 for a full proof. ■

Let us now average over all values of the shift parameter $j = 0, 1, \dots, R - 1$ to obtain

$$y_v = \frac{1}{R} \sum_{j=0}^{R-1} y_v(j) = \begin{cases} \frac{1}{R} \sum_{d=0}^r g_{v,d}^- & \text{if } v \text{ is an up-agent} \\ \frac{1}{R} \sum_{d=0}^r g_{v,d}^+ & \text{if } v \text{ is a down-agent.} \end{cases} \quad (5.2)$$

Lemma 5.3. *The vector \mathbf{y} is a feasible solution of the max-min LP associated with \mathcal{G} . For every objective $k \in K$, it holds that $\omega_k(\mathbf{y}) \geq (1 - 1/R) \min_{v \in V_k} s_v$.*

Proof. Follows from Lemma 5.2. ■

5.2. Averaging

Associate with each agent $v \in V$ a solution $\mathbf{y}^{\uparrow v}$ defined as follows. Choose the layers so that v is an up-agent; this is always possible. Let $\mathbf{y}^{\uparrow v}$ be the value of (5.2).

Lemma 5.4. *The vector \mathbf{x} is a feasible solution of the max-min LP associated with \mathcal{G} .*

Proof. Consider an arbitrary constraint $i \in I$; let $V_i = \{v, w\}$. Note that whenever v is an up-agent w is a down-agent and vice versa. Let

$$\mathbf{z} = (\mathbf{y}^{\uparrow v} + \mathbf{y}^{\uparrow w})/2.$$

By (5.2) and (4.13) we have

$$z_v = \frac{y_v^{\uparrow v} + y_v^{\uparrow w}}{2} = \frac{1}{2} \left(\frac{1}{R} \sum_{d=0}^r g_{v,d}^- + \frac{1}{R} \sum_{d=0}^r g_{v,d}^+ \right) = x_v$$

and $z_w = x_w$. By Lemma 5.3, the solutions $\mathbf{y}^{\uparrow v}$ and $\mathbf{y}^{\uparrow w}$ do not violate the constraint i . Therefore

$$a_v x_v + a_w x_w = a_v z_v + a_w z_w = ((a_v y_v^{\uparrow v} + a_w y_w^{\uparrow v}) + (a_v y_v^{\uparrow w} + a_w y_w^{\uparrow w}))/2 \leq (1 + 1)/2 = 1.$$

We conclude that the solution \mathbf{x} does not violate the constraint i . ■

Lemma 5.5. *For every objective $k \in K$,*

$$\omega_k(\mathbf{x}) \geq \frac{1}{2} \left(1 - \frac{1}{R} \right) \frac{|V_k|}{|V_k| - 1} \min_{v \in V_k} s_v. \quad (5.3)$$

Proof. Consider an arbitrary objective k in \mathcal{G} . Note that whenever $v \in V_k$ is an up-agent, each $w \in N(v)$ is a down-agent. Let

$$\mathbf{z} = \frac{1}{|V_k|} \sum_{v \in V_k} \mathbf{y}^{\uparrow v}.$$

By (5.2) and (4.13) we have

$$\frac{|V_k|}{2(|V_k| - 1)} z_u = \frac{1}{2(|V_k| - 1)} \sum_{v \in V_k} y_u^{\uparrow v} = \frac{1}{|V_k| - 1} \left(\frac{1}{2R} \sum_{d=0}^r g_{u,d}^- + \frac{|V_k| - 1}{2R} \sum_{d=0}^r g_{u,d}^+ \right) \leq x_u$$

for all $u \in V_k$. By Lemma 5.3,

$$\omega_k(\mathbf{x}) = \sum_{u \in V_k} x_u \geq \sum_{u \in V_k} \frac{|V_k|}{2(|V_k| - 1)} z_u \geq \frac{|V_k|}{2(|V_k| - 1)} \left(1 - \frac{1}{R}\right) \min_{v \in V_k} s_v.$$

The claim follows. ■

5.3. Completing the analysis

Lemmata 4.2 and 4.3 show that for any $v \in V$, the value t_v is an upper bound for the utility of any feasible solution of the max-min LP instance associated with \mathcal{G} , and so is s_v . Lemma 5.5 therefore shows that our local approximation algorithm achieves the approximation ratio of $2(1 - 1/\Delta_K)(1 + 1/(R - 1))$, for the special case studied in §4.

Together with the local transformations of §3, taking into account the increase of the approximation in §3.2, we conclude that the max-min LP problem admits a local algorithm with the approximation ratio of $\Delta_I(1 - 1/\Delta_K)(1 + 1/(R - 1))$; the local horizon is $\Theta(R)$. Theorem 1.2 follows by choosing a sufficiently large R . In particular, the constants Δ_I and Δ_K are not required to run the algorithm.

References

- [1] D. Angluin. Local and global properties in networks of processors. In *Proc. 12th Annual ACM Symposium on Theory of Computing (STOC, Los Angeles, CA, USA, April 1980)*, pages 82–93, New York, NY, USA, 1980. ACM Press.
- [2] B. S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *Journal of the ACM*, 41(1):153–180, 1994.
- [3] P. Floréen, M. Hassinen, P. Kaski, and J. Suomela. Local approximation algorithms for a class of 0/1 max-min linear programs, 2008. Manuscript, arXiv:0806.0282 [cs.DC].
- [4] P. Floréen, M. Hassinen, P. Kaski, and J. Suomela. Tight local approximation results for max-min linear programs. In *Proc. 4th International Workshop on Algorithmic Aspects of Wireless Sensor Networks (Algosensors, Reykjavík, Iceland, July 2008)*, Lecture Notes in Computer Science, Berlin, Germany, 2008. Springer-Verlag. To appear.
- [5] P. Floréen, P. Kaski, T. Musto, and J. Suomela. Approximating max-min linear programs with local algorithms. In *Proc. 22nd IEEE International Parallel and Distributed Processing Symposium (IPDPS, Miami, FL, USA, April 2008)*, Piscataway, NJ, USA, 2008. IEEE.
- [6] D. S. Hochbaum and W. Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *Journal of the ACM*, 32(1):130–136, 1985.
- [7] F. Kuhn. *The Price of Locality: Exploring the Complexity of Distributed Coordination Primitives*. PhD thesis, ETH Zürich, Dec. 2005.
- [8] F. Kuhn, T. Moscibroda, and R. Wattenhofer. The price of being near-sighted. In *Proc. 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA, Miami, FL, USA, January 2006)*, pages 980–989, New York, NY, USA, 2006. ACM Press.
- [9] N. Linial. Locality in distributed graph algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992.
- [10] M. Naor and L. Stockmeyer. What can be computed locally? *SIAM Journal on Computing*, 24(6):1259–1277, 1995.
- [11] C. H. Papadimitriou and M. Yannakakis. Linear programming without the matrix. In *Proc. 25th Annual ACM Symposium on Theory of Computing (STOC, San Diego, CA, USA, May 1993)*, pages 121–129, New York, NY, USA, 1993. ACM Press.
- [12] N. E. Young. Sequential and parallel algorithms for mixed packing and covering. In *Proc. 42nd Annual Symposium on Foundations of Computer Science (FOCS, Las Vegas, NV, USA, October 2001)*, pages 538–546, Los Alamitos, CA, USA, 2001. IEEE Computer Society Press.

Appendix A.

A.1. Local computability of the transformations in §3

Even though a description of a local algorithm often involves interleaved steps of communication and computation, it should be noted that any local algorithm with local horizon D can always be implemented as follows:

- (1) Each node gathers full information about its radius D neighbourhood; this is the *local view* of the agent.
- (2) Each node simulates the algorithm in its local view to determine its output.

As pointed out in §2, we can assume that the local view is a tree. Furthermore, in our case it is sufficient that each *agent* performs these steps – constraints and objectives do not need to produce any output.

Therefore we can implement each transformation presented in this section as follows (with a small increase of the local horizon):

- (1) Each agent gathers its local view, up to some constant distance. This is a tree.
- (2) Each agent performs the transformation in its local view. The result is a graph, possibly with cycles.
- (3) Each agent unfolds the graph to obtain a tree, discarding parts that are beyond its local horizon.
- (4) Each agent simulates the rest of the local algorithm in this tree, and applies the back-mapping to determine its output.

In some of the transformations, new agent nodes are created. In §3.1, the output of the new agents is not needed. In §3.3 and §3.4, existing agents can simulate their copies and compute the back-mapping.

The nontrivial part is to make sure that the transformations can be performed deterministically: if the local views of agents u and v partially overlap, and both agents perform a transformation in the common part, the common parts must be identical after the transformation. In particular, the port numbers must be identical. In the following, we show how to achieve this.

§3.1: Node k is chosen to be the node in K_v that has the smallest port number in v . The port number in i for $\{i, s\}$ is one larger than the port number for $\{i, v\}$. Within the gadget defined by s, t, u, h, ℓ, j , we can choose some fixed port numbering.

§3.2, §3.3, §3.4: These transformations can be summarised as follows: pick a subgraph, take n copies of it, and discard some edges (e.g., in §3.3, we take copies of a subgraph induced by an agent $v \in V$ and the adjacent constraints I_v). We can choose the port numbers of the transformed instance deterministically as follows. While we are taking copies of subgraphs, copy the port numbers as well. This may create duplicates. However, in each of these transformations, a set of duplicate port numbers in a node corresponds to a set of copies of one subgraph (e.g., in §3.3, each copy of the subgraph induced by v and I_v creates exactly one duplicate port number for each agent $u \neq v$ adjacent to any $i \in I_v$). Furthermore, we can impose an ordering for the copies: in §3.2 by using pairs of port number in i ; in §3.3 by using port numbers in v ; and in §3.4 by arbitrarily ordering otherwise identical copies. Therefore we can re-number the ports in each node: first, order the ports by existing port numbers, using the ordering of the copies to break the ties; second, assign new port numbers $1, 2, \dots$ in this order.

§3.5: The graph is not changed, and port numbering is preserved.

A.2. Proof of Lemma 4.3

To set up the base case at $d = 0$, observe that by the feasibility of \mathbf{x} and (4.1) we have

$$x_v \leq \min_{i \in I_v} a_{iv}^{-1} = f_{u,v,0}^+(\omega) \quad \text{for all } v \in L(u, 4r + 1).$$

Next consider any $0 \leq d \leq r$ and assume inductively that

$$x_w \leq f_{u,w,d}^+(\omega) \quad \text{for all } w \in L(u, 4(r - d) + 1).$$

Consider an arbitrary $v \in L(u, 4(r - d) - 1)$. Observe that for all $w \in V_{k(v)}$ it holds that either $v = w$ or $w \in L(u, 4(r - d) + 1)$. If we have

$$0 \geq \omega - \sum_{w \in N(v)} f_{u,w,d}^+(\omega),$$

then (4.2) implies $f_{u,v,d}^-(\omega) = 0 \leq x_v$; otherwise (4.2) implies

$$f_{u,v,d}^-(\omega) = \omega - \sum_{w \in N(v)} f_{u,w,d}^+(\omega) \leq \omega - \sum_{w \in N(v)} x_w \leq x_v.$$

Here the first inequality follows by the induction hypothesis, and the second inequality follows by assumption that \mathbf{x} achieves the objective value ω ; in particular,

$$\sum_{w \in V_{k(v)}} x_w = x_v + \sum_{w \in N(v)} x_w \geq \omega.$$

To complete the induction, consider any $1 \leq d \leq r$ and assume inductively that

$$f_{u,w,d-1}^-(\omega) \leq x_w \quad \text{for all } w \in L(u, 4(r - (d - 1)) - 1).$$

Consider an arbitrary $v \in L(u, 4(r - d) + 1)$. Observe that $n(v, i) \in L(u, 4(r - (d - 1)) - 1)$. Because \mathbf{x} is feasible, we have $a_{iv}x_v + a_{i,n(v,i)}x_{n(v,i)} \leq 1$ for all $i \in I_v$. Thus, the inductive hypothesis and (4.3) imply

$$x_v \leq \min_{i \in I_v} a_{iv}^{-1} \left(1 - a_{i,n(v,i)}x_{n(v,i)} \right) \leq \min_{i \in I_v} a_{iv}^{-1} \left(1 - a_{i,n(v,i)}f_{u,n(v,i),d-1}^-(\omega) \right) = f_{u,v,d}^+(\omega).$$

To conclude that t_u is the optimum utility of the max-min LP associated with \mathcal{A}_u , consider an optimal solution \mathbf{x} with the objective value ω^* . Observe that (4.6) and (4.7) imply (4.4) and (4.5). Therefore $t_u \geq \omega^*$. Furthermore, $t_u > \omega^*$ would contradict with the assumption that \mathbf{x} is optimal. \blacksquare

A.3. Proof of Lemma 4.4

Consider an arbitrary $u \in V$. For all $w \in V$ at distance at most $4r + 2$ from u we have, by the definition of s_w ,

$$0 \leq s_w \leq t_u. \quad (\text{A.1})$$

We proceed by induction on d . To establish the base case at $d = 0$, observe by (4.1) and (4.8) that $f_{u,v,0}^+ = g_{v,0}^+$ for all $v \in L(u, 4r + 1)$.

Next consider any $0 \leq d \leq r$ and assume inductively that

$$f_{u,w,d}^+ \leq g_{w,d}^+ \quad \text{for all } w \in L(u, 4(r - d) + 1).$$

Consider an arbitrary $v \in L(u, 4(r - d) - 1)$. Observe that for all $w \in V_{k(v)}$ either $w = v$ or $w \in L(u, 4(r - d) + 1)$. Apply (4.9), (A.1), the inductive hypothesis, and (4.2) to obtain

$$g_{v,d}^- = \max\left(0, s_v - \sum_{w \in N(v)} g_{w,d}^+\right) \leq \max\left(0, t_u - \sum_{w \in N(v)} f_{u,w,d}^+\right) = f_{u,v,d}^-.$$

Finally, consider any $1 \leq d \leq r$ and assume inductively that

$$g_{w,d-1}^- \leq f_{u,w,d-1}^- \quad \text{for all } w \in L(u, 4(r - (d - 1)) - 1).$$

Consider an arbitrary $v \in L(u, 4(r - d) + 1)$. Observe that for all $i \in I_v$ it holds that $n(v, i) \in L(u, 4(r - (d - 1)) - 1)$. Apply (4.3), the inductive hypothesis, and (4.10), to obtain

$$f_{u,v,d}^+ = \min_{i \in I_v} a_{iv}^{-1} \left(1 - a_{i,n(v,i)} f_{u,n(v,i),d-1}^-\right) \leq \min_{i \in I_v} a_{iv}^{-1} \left(1 - a_{i,n(v,i)} g_{n(v,i),d-1}^-\right) = g_{v,d}^+.$$

The induction is now complete. \blacksquare

A.4. Proof of Lemma 4.6

To set up the base case at $d = 1$, observe first that $g_{v,0}^- \geq 0$ by (4.9). By (4.10) and (4.8) thus

$$g_{v,1}^+ = \min_{i \in I_v} a_{iv}^{-1} \left(1 - a_{i,n(v,i)} g_{n(v,i),0}^-\right) \leq \min_{i \in I_v} a_{iv}^{-1} = g_{v,0}^+.$$

Next consider any $1 \leq d \leq r$ and assume inductively that $g_{v,d}^+ \leq g_{v,d-1}^+$. Apply (4.9) and the inductive hypothesis to obtain

$$g_{v,d}^- = \max\left(0, s_v - \sum_{w \in N(v)} g_{w,d}^+\right) \geq \max\left(0, s_v - \sum_{w \in N(v)} g_{w,d-1}^+\right) = g_{v,d-1}^-.$$

Finally, consider any $2 \leq d \leq r$ and assume inductively that $g_{v,d-2}^- \leq g_{v,d-1}^-$. Apply (4.10) and the inductive hypothesis to obtain

$$g_{v,d}^+ = \min_{i \in I_v} a_{iv}^{-1} \left(1 - a_{i,n(v,i)} g_{n(v,i),d-1}^-\right) \leq \min_{i \in I_v} a_{iv}^{-1} \left(1 - a_{i,n(v,i)} g_{n(v,i),d-2}^-\right) = g_{v,d-1}^+.$$

This completes the induction. \blacksquare

A.5. Proof of Lemma 5.2

Feasibility. The vector is nonnegative by (4.9) and Lemma 4.7. Consider an arbitrary constraint $i \in I$. By Lemma 5.1 we can represent the layer of i uniquely as $4(Rc+j)+4d+2$ for integers c, d with $0 \leq d \leq R-1$. Let $V_i = \{v, w\}$ with v at layer $4(Rc+j)+4d+1$ and w at layer $4(Rc+j)+4(d+1)-1$.

First consider the case $d = R-1$. Note that the layer of w is actually $4(R(c+1)+j)-1$, that is, $d = 0$ and $e = -1$ for w . By (5.1) and Lemma 4.5, we have

$$a_{iv}y_v(j) + a_{iw}y_w(j) = a_{iw}g_{w,r}^- \leq 1.$$

Next consider the case $d = R-2$. By (5.1) and (4.8), we have

$$a_{iv}y_v(j) + a_{iw}y_w(j) = a_{iv}g_{v,0}^+ \leq 1.$$

Finally consider the case $d < R-2$. By (5.1) and (4.10), we have

$$\begin{aligned} a_{iv}y_v(j) + a_{iw}y_w(j) &= a_{iv}g_{v,r-d}^+ + a_{iw}g_{w,r-d-1}^- \\ &\leq a_{iv}a_{iv}^{-1}(1 - a_{iw}g_{w,r-d-1}^-) + a_{iw}g_{w,r-d-1}^- = 1. \end{aligned}$$

The claim follows since i was arbitrary.

Objectives. Consider an arbitrary objective k in \mathcal{G} . By Lemma 5.1 we can represent the layer of k uniquely as $4(Rc+j)+4d$ for integers c, d with $0 \leq d \leq R-1$. There is a unique up-agent in V_k at layer $4(Rc+j)+4d-1$. Denote this agent by v . The other agents $w \in N(v)$ are down-agents at layer $4(Rc+j)+4d+1$.

First consider the case $d = R-1$. By (5.1) we have

$$\omega_k(\mathbf{y}(j)) = y_v(j) + \sum_{w \in N(v)} y_w(j) = 0.$$

Then consider the case $d \leq R-2$. By (5.1) and (4.9), we have

$$\begin{aligned} \omega_k(\mathbf{y}(j)) &= y_v(j) + \sum_{w \in N(v)} y_w(j) = g_{v,r-d}^- + \sum_{w \in N(v)} g_{w,r-d}^+ \\ &\geq s_v - \sum_{w \in N(v)} g_{w,r-d}^+ + \sum_{w \in N(v)} g_{w,r-d}^+ = s_v \geq \min_{u \in V_k} s_u. \end{aligned}$$

The claim follows because k was arbitrary. ■