An optimal local approximation algorithm for max-min linear programs

$$\begin{aligned} \text{max } & \omega \\ \text{s.t. } & A\mathbf{x} \le \mathbf{1}, \\ & C\mathbf{x} \ge \omega\mathbf{1}, \\ & \mathbf{x} \ge \mathbf{0} \end{aligned}$$

Patrik Floréen
Joel Kaasinen
Petteri Kaski
Jukka Suomela

Helsinki Institute for Information Technology HIIT
University of Helsinki, Finland

SPAA, Calgary, Canada, 13 August 2009

Result on one slide

$$\begin{aligned} \text{max} \quad & \omega \\ \text{s.t.} \quad & A\mathbf{x} \le \mathbf{1}, \\ & C\mathbf{x} \ge \omega\mathbf{1}, \\ & \mathbf{x} \ge \mathbf{0} \end{aligned}$$

A, C: nonnegative matrices

edge $\sim$ positive coefficient

Distributed setting:

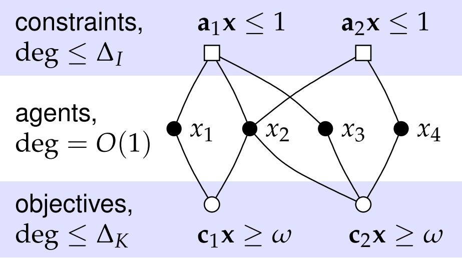

Approximability with constant-time distributed algorithms:

– new positive result: $\Delta_I(1 - 1/\Delta_K) + \epsilon$

– earlier negative result: $\Delta_I(1 - 1/\Delta_K)$

Max-min linear programs

General form:

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

Equivalent form:

$$\begin{aligned} \text{maximise} \quad & \omega \\ \text{subject to} \quad & A\mathbf{x} \leq \mathbf{1}, \\ & C\mathbf{x} \geq \omega\mathbf{1}, \\ & \mathbf{x} \geq \mathbf{0} \end{aligned}$$

A and C are nonnegative matrices

Intuition: solution $\mathbf{x}$ uses $\mathbf{a}_i\mathbf{x}$ units of resource $i \in I$, and provides $\mathbf{c}_k\mathbf{x}$ units of service to customer $k \in K$

Max-min LPs vs. packing LPs

Max-min LP:

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

Packing LP:

$$\begin{aligned} \text{maximise} \quad & \mathbf{c}\mathbf{x} \\ \text{subject to} \quad & A\mathbf{x} \leq \mathbf{1}, \\ & \mathbf{x} \geq \mathbf{0} \end{aligned}$$

A, C, and $\mathbf{c}$ are nonnegative

Applications of max-min LPs

Maximising the lifetime of a wireless sensor network:

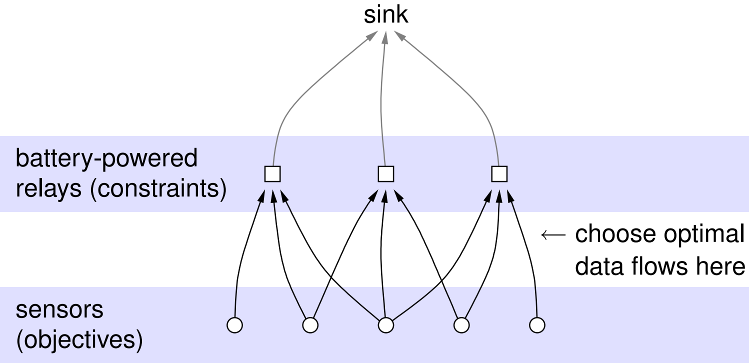

Applications of max-min LPs

Maximising the lifetime of a wireless sensor network:

Abstraction that we study here:

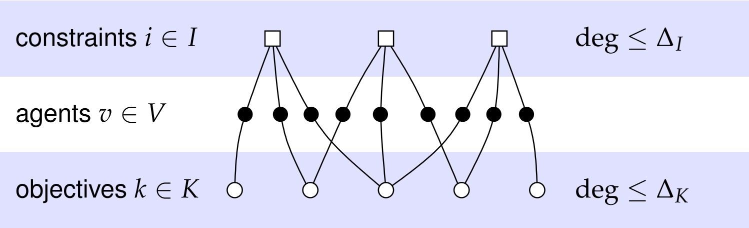

Applications of max-min LPs

Max-min linear program:

$$\begin{aligned} \text{maximise} \quad & \omega \\ \text{subject to} \quad A\mathbf{x} &\le \mathbf{1}, \\ C\mathbf{x} &\ge \omega\mathbf{1}, \\ \mathbf{x} &\ge \mathbf{0} \end{aligned}$$

Mixed packing and covering problem:

$$\begin{aligned} \text{find} \quad & \mathbf{x} \\ \text{such that} \quad A\mathbf{x} &\le \mathbf{1}, \\ C\mathbf{x} &\ge \mathbf{1}, \\ \mathbf{x} &\ge \mathbf{0} \end{aligned}$$

Near-optimal solution to max-min LP $\Longrightarrow$
near-feasible solution to mixed packing and covering
(or proof that there is no feasible solution)

Problem

$$\begin{aligned} \max \quad & \omega \\ \text{s.t.} \quad A\mathbf{x} &\le \mathbf{1}, \\ C\mathbf{x} &\ge \omega\mathbf{1}, \\ \mathbf{x} &\ge \mathbf{0} \end{aligned}$$

Focus: distributed algorithms that run in constant time (*local algorithms*)

Running time may depend on parameters Δ_I, Δ_K, etc., but must be *independent* of the number of nodes

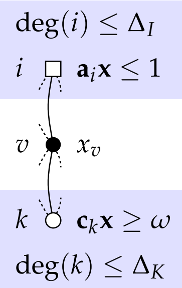

Old results

$$\begin{aligned} \max \ & \omega \\ \text{s.t.} \ & A\mathbf{x} \le \mathbf{1}, \\ & C\mathbf{x} \ge \omega\mathbf{1}, \\ & \mathbf{x} \ge \mathbf{0} \end{aligned}$$

Old negative result:

- Approximation factor $\Delta_I(1 - 1/\Delta_K)$ impossible

Old positive results:

- Approximation factor Δ_I easy (Papadimitriou–Yannakakis 1993)
- Factor $\Delta_I(1 - 1/\Delta_K) + \epsilon$ possible in some special cases

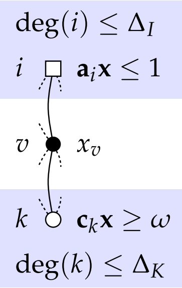

New results

$$\begin{aligned} \max \quad & \omega \\ \text{s.t.} \quad & A\mathbf{x} \le \mathbf{1}, \\ & C\mathbf{x} \ge \omega\mathbf{1}, \\ & \mathbf{x} \ge \mathbf{0} \end{aligned}$$

Old negative result:

- Approximation factor $\Delta_I(1 - 1/\Delta_K)$ impossible

New positive result:

- Approximation factor $\Delta_I(1 - 1/\Delta_K) + \epsilon$ possible for any constant $\epsilon > 0$

Matching upper and lower bounds!

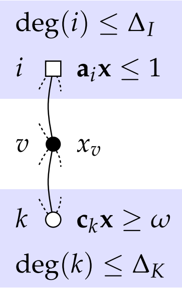

New results

$$\begin{aligned} \text{max} \quad & \omega \\ \text{s.t.} \quad A\mathbf{x} &\le \mathbf{1}, \\ C\mathbf{x} &\ge \omega\mathbf{1}, \\ \mathbf{x} &\ge \mathbf{0} \end{aligned}$$

Tight bound $\Delta_I(1 - 1/\Delta_K) + \epsilon$ holds for any combination of these assumptions:

- anonymous networks
  or unique identifiers
- $0/1$ coefficients in A, C
  or arbitrary nonnegative numbers
- one nonzero per column in A, C
  or arbitrary structure

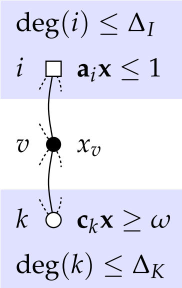

Local reductions

It is enough to solve the following special case:

- Communication graph is (infinite) tree
- Degree of each constraint is 2
- Degree of each objective is at least 2
- Each agent is adjacent to at least one constraint
- Each agent is adjacent to exactly one objective

General result then follows by a series of *local reductions*

Local reductions

Hence we focus on instances with the following structure:

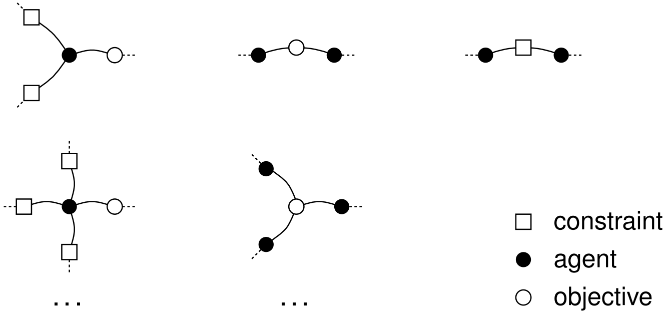

□ Local reductions

An example:

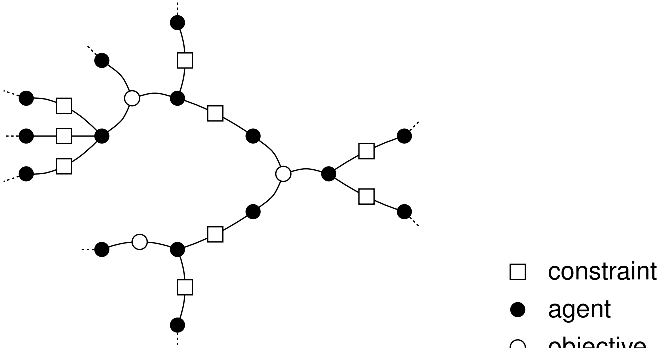

Algorithm

How to solve it? We begin with a thought experiment...

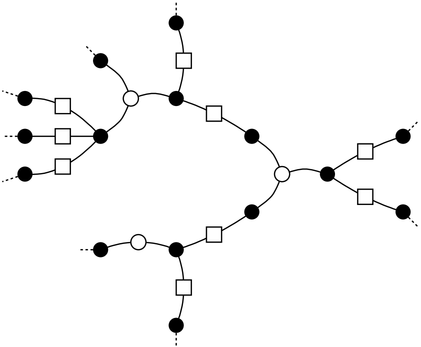

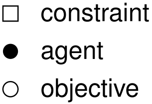

□ Two roles: “up” and “down”

What if we could partition agents in two sets so that there is exactly one up-agent adjacent to each constraint or objective?

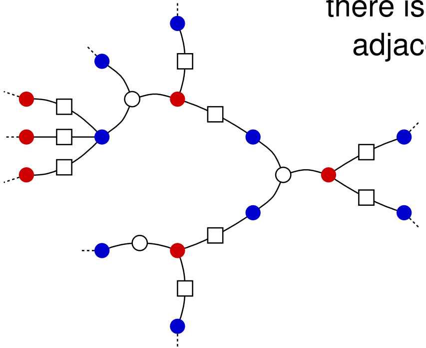

□ constraint

● up-agent

● down-agent

○ objective

☐ Layers

Then we could also organise the graph in layers

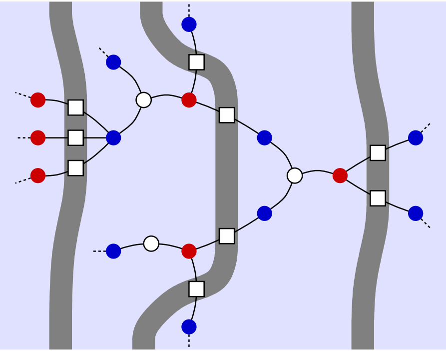

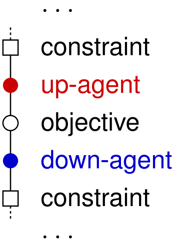

Layers

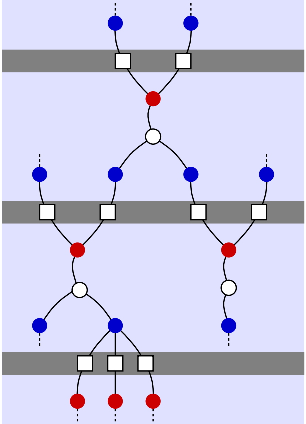

Solve by using layers:

- Message propagation upwards
- Use the *shifting strategy*
- Remove slack:
  down-agents choose large values, up-agents choose small values

Layers

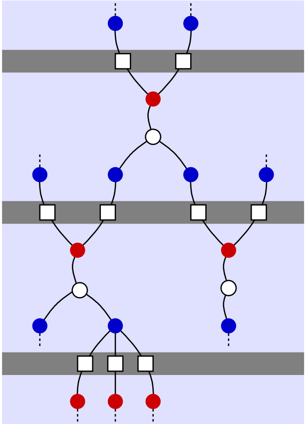

Solve by using layers:

. . .

Globally consistent solution, $(1+\epsilon)$-approximation

But we had to assume that the agents are partitioned in two sets, “up” and “down”!

Trick

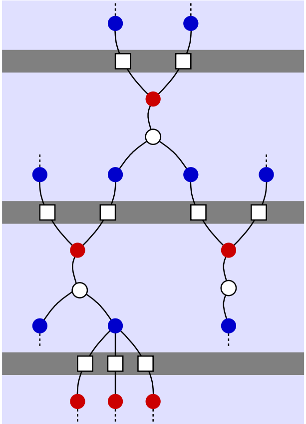

Useful property: the output of a node depends only on its *own* role (up or down)

Consider both roles, take the average!

A lucky coincidence: approximation guarantee weakens only by factor $\Delta_I(1 - 1/\Delta_K)$

Summary

$$\begin{aligned} \max \ & \omega \\ \text{s.t.} \ & A\mathbf{x} \leq \mathbf{1}, \\ & C\mathbf{x} \geq \omega\mathbf{1}, \\ & \mathbf{x} \geq \mathbf{0} \end{aligned}$$

Distributed setting:

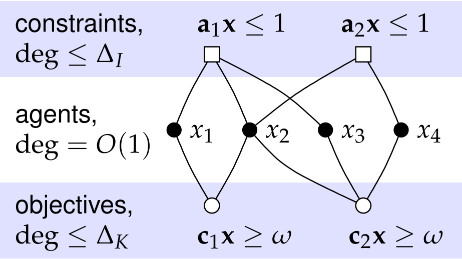

A, C: nonnegative matrices

edge $\sim$ positive coefficient

Approximability with constant-time distributed algorithms:

– new positive result: $\Delta_I(1 - 1/\Delta_K) + \epsilon$

– earlier negative result: $\Delta_I(1 - 1/\Delta_K)$