Tight local approximation results for max-min linear programs

Patrik Floréen,
Marja Hassinen,
Petteri Kaski,
Jukka Suomela

HIIT,
University of Helsinki,
Finland

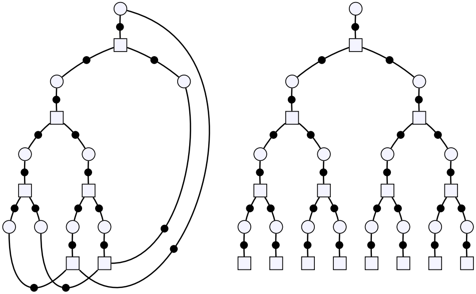

Algosensors
12 July 2008

Local algorithms

Local algorithm: output of a node is a function of input within its *constant-radius neighbourhood*

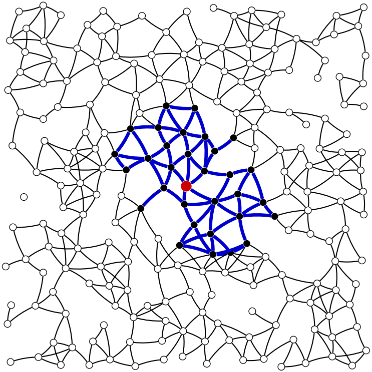

(Linial 1992; Naor and Stockmeyer 1995)

Local algorithms

Local algorithm: changes outside the *local horizon* of a node do not affect its output

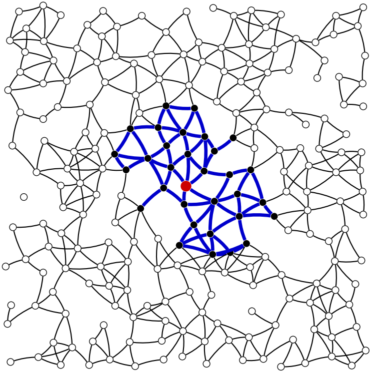

(Linial 1992; Naor and Stockmeyer 1995)

Max-min linear program

Let $A \geq 0$, $c_k \geq 0$

Objective:

$$\begin{aligned} \text{maximise} \quad & \min_{k \in K} c_k^\top x \\ \text{subject to} \quad & Ax \leq \mathbf{1}, \\ & x \geq \mathbf{0} \end{aligned}$$

Generalisation of packing LP:

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

Max-min linear program

Example: Data gathering in a sensor network

- circle = sensor
- square = relay
- edge = network connection

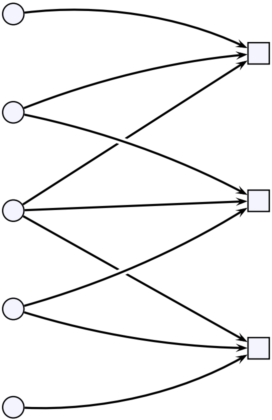

Max-min linear program

Example: Maximise the minimum amount of data gathered from each sensor

$$\text{maximise} \quad \min \{ x_1,\ \underline{x_2 + x_4},\ x_3 + x_5 + x_7,\ x_6 + x_8,\ x_9 \}$$

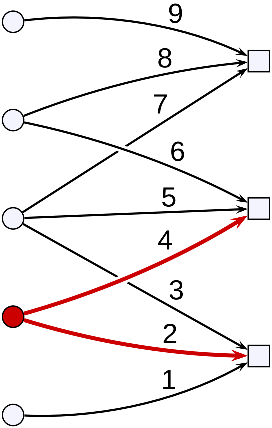

Max-min linear program

Example: Maximise the minimum amount of data gathered from each sensor; each relay has a limited battery capacity

$$
\begin{aligned}
\text{maximise} \quad & \min \big\{ \\
& \quad x_1,\ x_2 + x_4, \\
& \quad x_3 + x_5 + x_7, \\
& \quad x_6 + x_8,\ x_9 \\
& \big\} \\
\text{subject to} \quad & \underline{x_1 + x_2 + x_3 \le 1}, \\
& x_4 + x_5 + x_6 \le 1, \\
& x_7 + x_8 + x_9 \le 1, \\
& x_1, x_2, \dots, x_9 \ge 0
\end{aligned}
$$

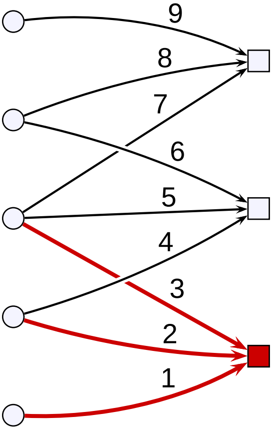

Max-min linear program

Example: Maximise the minimum amount of data gathered from each sensor; each relay has a limited battery capacity

An optimal solution:

$$\begin{aligned} x_1 = x_5 = x_9 &= 3/5, \\ x_2 = x_8 &= 2/5, \\ x_4 = x_6 &= 1/5, \\ x_3 = x_7 &= 0 \end{aligned}$$

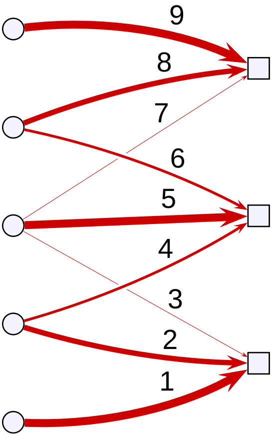

Max-min linear program

Communication graph: $\mathcal{G} = (V \cup I \cup K, E)$

$$
\begin{aligned}
\text{maximise} \quad & \min \{ \\
& \quad x_1,\ x_2 + x_4, \\
& \quad x_3 + x_5 + x_7, \\
& \quad x_6 + x_8,\ x_9 \\
& \} \\
\text{subject to} \quad & x_1 + x_2 + x_3 \le 1, \\
& x_4 + x_5 + x_6 \le 1, \\
& x_7 + x_8 + x_9 \le 1, \\
& \underline{x_1}, x_2, \ldots, x_9 \ge 0
\end{aligned}
$$

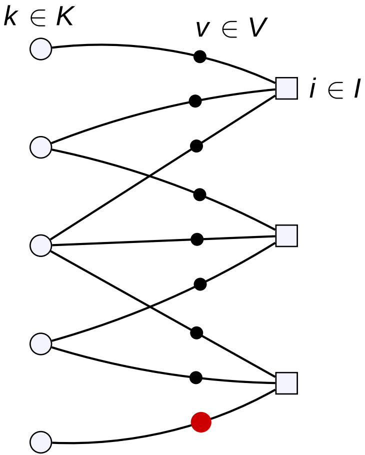

Max-min linear program

Communication graph: $\mathcal{G} = (V \cup I \cup K, E)$

$$
\begin{aligned}
\text{maximise} \quad & \min \Big\{ \\
& \quad x_1,\ \underline{x_2 + x_4}, \\
& \quad x_3 + x_5 + x_7, \\
& \quad x_6 + x_8,\ x_9 \\
& \Big\} \\
\text{subject to} \quad & x_1 + x_2 + x_3 \leq 1, \\
& x_4 + x_5 + x_6 \leq 1, \\
& x_7 + x_8 + x_9 \leq 1, \\
& x_1, x_2, \ldots, x_9 \geq 0
\end{aligned}
$$

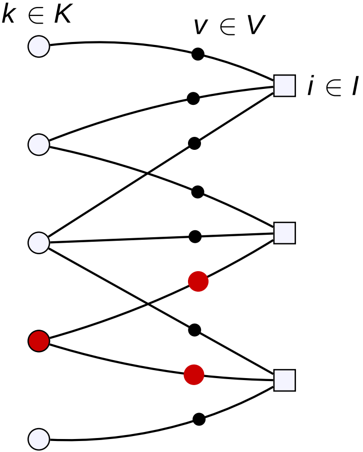

Max-min linear program

Communication graph: $\mathcal{G} = (V \cup I \cup K, E)$

$$
\begin{aligned}
\text{maximise} \quad & \min\Big\{ \\
& \quad x_1,\ x_2 + x_4, \\
& \quad x_3 + x_5 + x_7, \\
& \quad x_6 + x_8,\ x_9 \\
& \Big\} \\
\text{subject to} \quad & \underline{x_1 + x_2 + x_3 \le 1}, \\
& x_4 + x_5 + x_6 \le 1, \\
& x_7 + x_8 + x_9 \le 1, \\
& x_1, x_2, \dots, x_9 \ge 0
\end{aligned}
$$

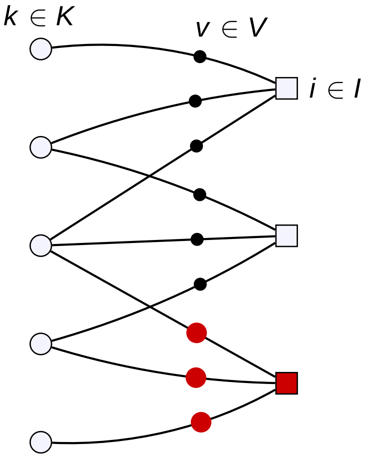

Max-min linear program

Communication graph: $\mathcal{G} = (V \cup I \cup K, E)$

Key parameters:

Δ_I = max. degree of $i \in I$
Δ_K = max. degree of $k \in K$

Problem is *bipartite* if each $v \in V$ adjacent to exactly one $i \in I$ and exactly one $k \in K$

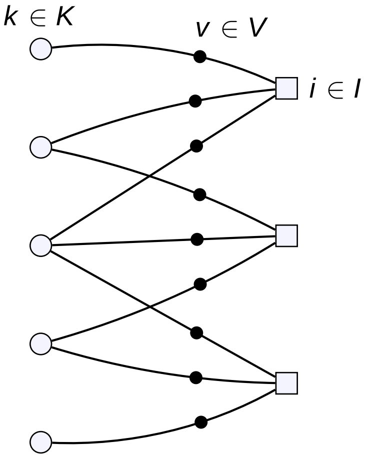

Old results

“Safe algorithm”:

Node v chooses

$$x_v = \min_{i\,:\,a_{iv}>0} \frac{1}{a_{iv}\,\left|\{u : a_{iu} > 0\}\right|}$$

(Papadimitriou and Yannakakis 1993)

Factor Δ_I approximation

Uses information only in radius 1 neighbourhood of v

A better approximation ratio with a larger radius?

New results: bipartite problems

The safe algorithm is factor Δ_I approximation

Theorem

For any $\epsilon > 0$, there is a local algorithm for bipartite max-min LPs with approximation ratio $\Delta_I(1 - 1/\Delta_K) + \epsilon$

Theorem

There is no local algorithm for bipartite max-min LPs with approximation ratio $\Delta_I(1 - 1/\Delta_K)$

Degree of a constraint $i \in I$ is at most Δ_I

Degree of an objective $k \in K$ is at most Δ_K

Bipartite: each $v \in V$ adjacent to one $i \in I$ and one $k \in K$

New results: bounded growth

Assume *bounded relative growth* beyond radius R:

$$\frac{|B(v, r+2)|}{|B(v, r)|} \leq 1 + \delta \qquad \text{for all } v \in V,\ r \geq R$$

where $B(v, r)$ = agents in radius r neighbourhood of v

There is a local algorithm for max-min LPs with approximation ratio $1 + 2\delta + o(\delta)$ (Floréen et al. 2008)

Theorem

There is no local algorithm for max-min LPs with approximation ratio $1 + \delta/2$
(assuming $\Delta_I \geq 3$, $\Delta_K \geq 3$, $0.0 < \delta < 0.1$)

Inapproximability

Regular high-girth graph or regular tree?

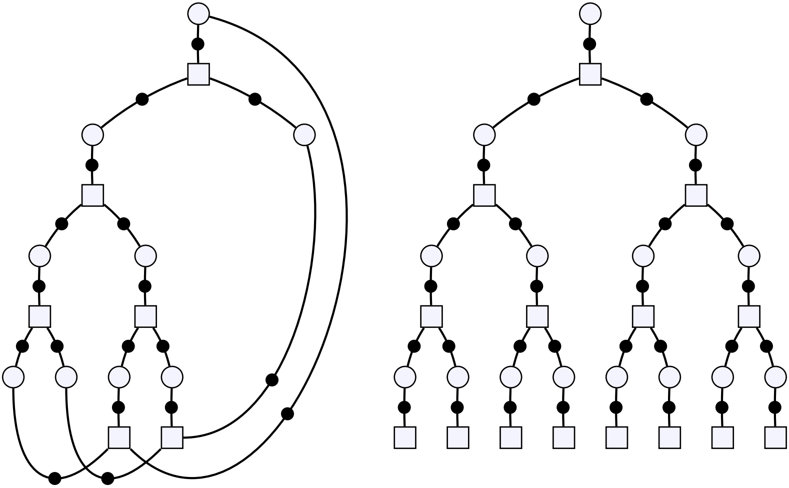

Inapproximability

Locally indistinguishable

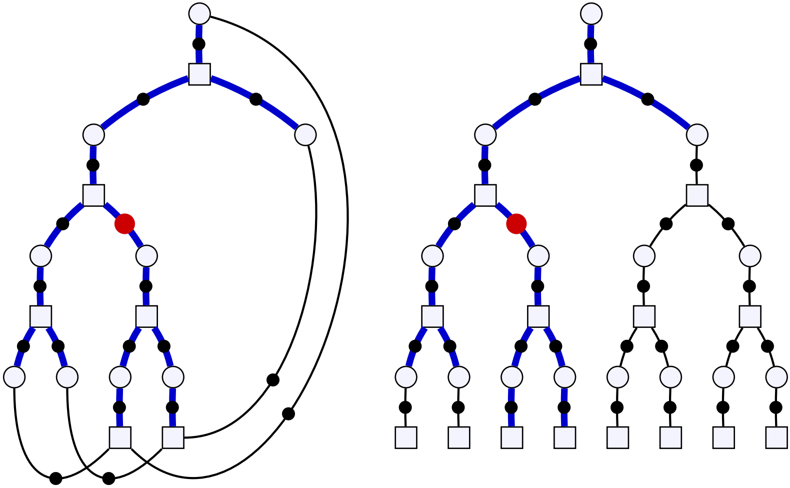

Inapproximability

Optimum $\leq 2/3$ vs. optimum ≥ 1

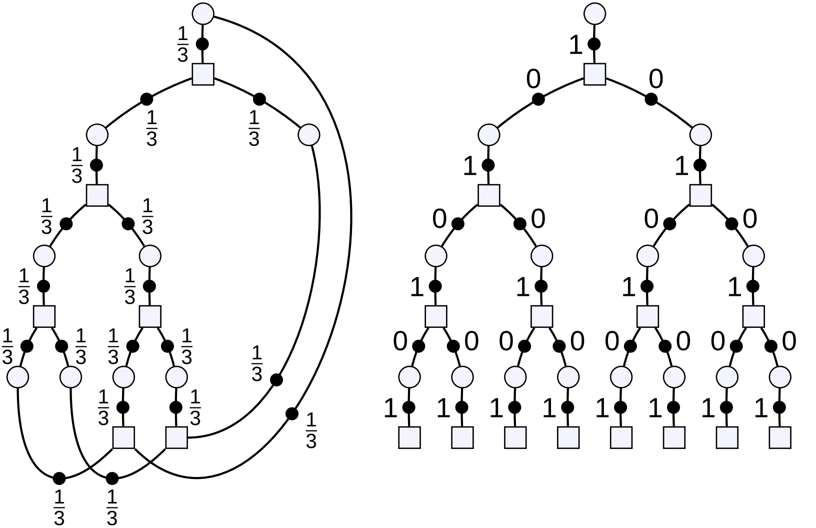

Inapproximability

Approx. ratio $\geq 1/(2/3) = 3\,(1 - 1/2) = \Delta_I\,(1 - 1/\Delta_K)$

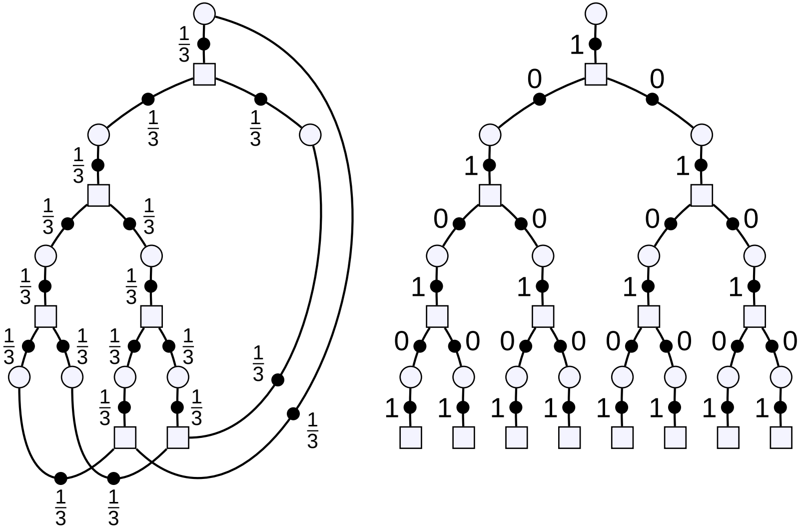

Approximability

Step 1: Unfold the graph into an infinite tree

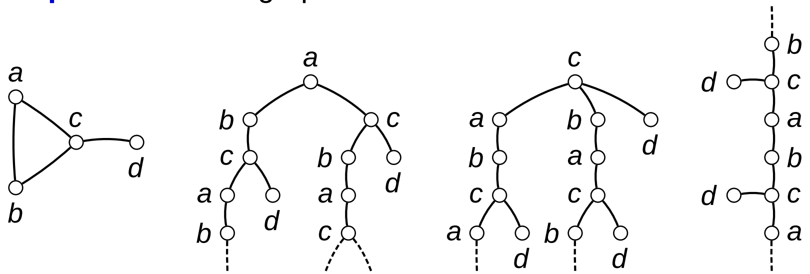

Step 2: Regularise the tree:
each constraint $i \in I$ has degree exactly Δ_I, etc.

Step 3: Construct local subproblems,
solve them optimally, take averages

Summary

Max-min linear programs: given $A, c_k \geq 0$,

$$\begin{aligned} \text{maximise} \quad & \min_{k \in K} c_k^\top x \\ \text{subject to} \quad & Ax \leq \mathbf{1},\ x \geq \mathbf{0} \end{aligned}$$

Local algorithms: output of a node is a function of input within its constant-radius neighbourhood

Results:

- Bipartite max-min LPs: tight upper and lower bound
- Bounded relative growth: near-tight lower bound

http://www.hiit.fi/ada/geru — jukka.suomela@cs.helsinki.fi

References

P. Floréen, P. Kaski, T. Musto, and J. Suomela (2008).
Approximating max-min linear programs with local algorithms.
Proc. IPDPS 2008. [DOI]

P. Floréen, M. Hassinen, P. Kaski, and J. Suomela (2008b).
Tight local approximation results for max-min linear programs.
Proc. Algosensors 2008.

N. Linial (1992). Locality in distributed graph algorithms.
SIAM Journal on Computing, 21(1):193–201. [DOI]

M. Naor and L. Stockmeyer (1995). What can be computed locally?
SIAM Journal on Computing, 24(6):1259–1277. [DOI]

C. H. Papadimitriou and M. Yannakakis (1993).
Linear programming without the matrix.
Proc. STOC 1993. [DOI]