


Stackelberg Scheduling Strategies

Tim Roughgarden
Cornell University

The Model


- m machines $1, 2, \dots, m$
- A quantity r of jobs
 - jobs are small (model in a **cts** way)
- For each machine i , a **load-dependent** latency function $l_i(\cdot)$
 - assume continuous, nondecreasing

Example: ($r=1$)




Equilibria

- which job assignments are “stable”?
 - jobs controlled by selfish, noncooperative agents
 - no job wants to switch machines (no job should be envious)



More on Equilibrium Assignments



Def: an assignment is at **Nash equilibrium** (is a **Nash assignment**) if:

- all used machines have equal latency
- unused machines have greater latency

Fact: always have existence, uniqueness

How Good is an Assignment?


The Cost of an Assignment:

- $C(x)$ = cost or total latency experienced by assignment x :


$$\sum_i x_i \cdot l_i(x_i)$$

- our notion of system performance
- can optimize in poly-time

Example:



1



x^2

cost =

$$\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{1}{4} = ?$$

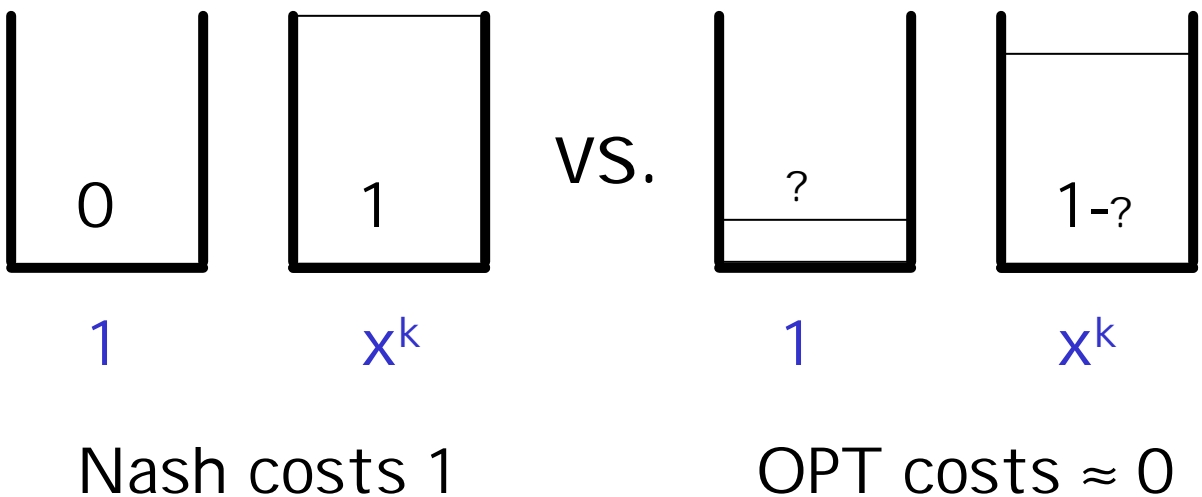
How Good are Nash Assignments?

Goal: prove that Nash assignments are **near-optimal**

- want a laissez faire approach to regulating users

Problem: false in general!

Example: ($r=1$, k large)



Near-Optimal Nash Assignments

Old Approach: weaken model,
compare Nash vs. OPT
(due to [Roughgarden/Tardos 00])

- general latency fns, weaker OPT

Thm 1: cost of Nash = cost of OPT
at rate $2r$

- stronger OPT, linear latency fns
($l_i(x) = a_i x + b_i$)

Thm 2: cost of Nash = $4/3$ cost of
OPT (at rate r)

Taming Selfishness through a Manager

New Approach:

- not all jobs need be controlled by selfish users
 - “centrally controlled” vs. “selfishly controlled” jobs
 - behavior of selfish users depends on assignment of managed jobs


Goal:

- assign centrally controlled jobs to induce “good” selfish behavior
 - see also [Korlis, Lazar, Orda 97]

Stackelberg Strategies

- **Stackelberg strategy** = assignment of centrally controlled jobs

↳ yields an **induced equilibrium**



- **Basic Questions:**
 - what's the best strategy?
 - can we compute/characterize it?
 - how **inefficient** is the best induced equilibrium?
 - are we provably near-optimal?

Our Results - General Latency Functions

Theorem 1: Can efficiently compute a strategy inducing an equilibrium with cost

$$= (1/\beta) \times \text{cost of opt assignment}$$


(β = fraction of centrally assigned jobs)

Fact:


$$(1/\beta) \times \text{OPT}$$

is best

possible



1



$$[x/(1-\beta)]^k$$

Our Results - Linear Latency Functions


Theorem 2: Can efficiently compute a strategy inducing an equilibrium with cost

$$= [4/(3+\beta)] \times \text{cost of OPT}$$


(β = fraction of centrally assigned jobs)

Fact:

$[4/(3+\beta)] \times$
OPT is best
possible



1




$x/(1-\beta)$

What makes a strategy (in)effective?

The Scale Strategy

- compute optimal assignment x of **all** jobs, assign centrally controlled jobs via $\beta \cdot x$

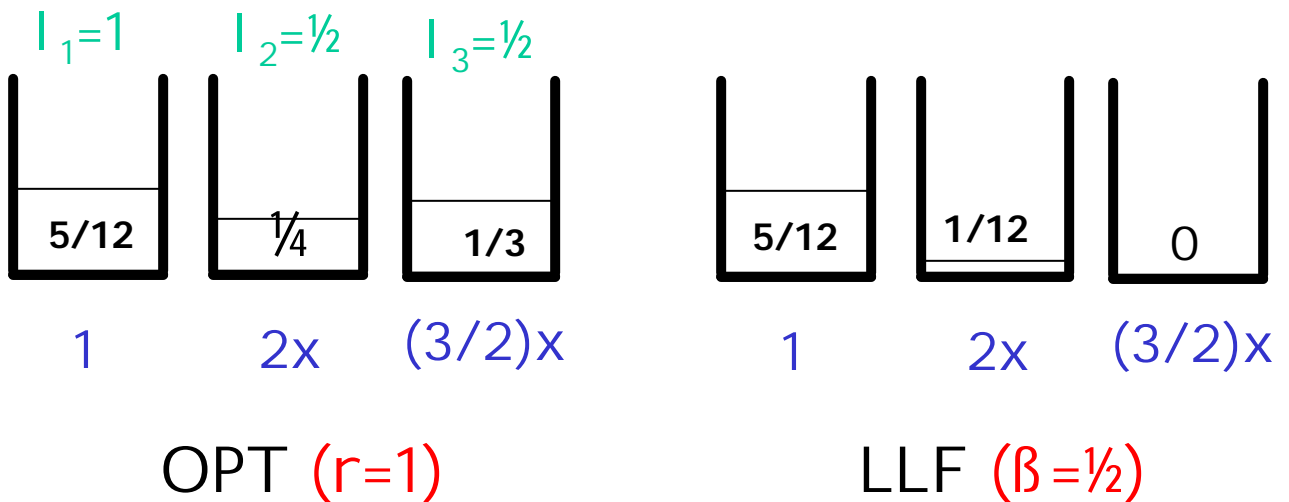


Moral: avoid machines that selfish users will (over)use anyways

The LLF Strategy

Largest Latency First (LLF):

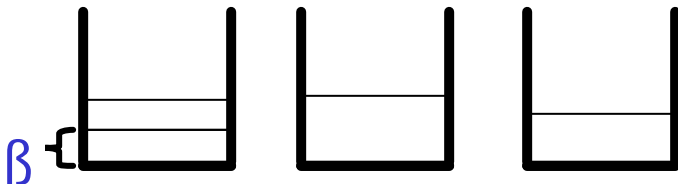
- compute opt, x_i , for all jobs
- assign x_i jobs to i in order of decreasing $l_i(x_i)$'s (until no managed jobs remain)



LLF with General Latency Functions

Theorem 1: The LLF strategy induces an equilibrium with cost $= (1/\beta) \times$ cost of opt assignment.

Proof idea: Exploit iterative structure of LLF to proceed by induction on # of machines.

Base case:  common latency = L


- LLF \Rightarrow Machine 1's latency $\geq L$
- OPT has $\geq \beta$ jobs on machine 1
 - \Rightarrow OPT pays $\geq \beta L$, we pay L

LLF with Linear Latency Functions

Theorem 2: The LLF strategy induces an equilibrium with cost $= [4/(3+\beta)] \times \text{cost of OPT}$.

Main difficulty:

- previous argument too weak for:



- need detailed study of Nash, OPT when all latencies are linear

Computing Optimal Strategies

We've seen: LLF has the best possible **worst-case** guarantee

Question: is the LLF strategy optimal **on all instances**?

Bad news: no, in fact:

Theorem 3: Computing the optimal strategy is **NP-hard** (even for linear latency fns).

- **Compare to:** Opt, Nash assignments

Open Questions

Approximating the optimal strategy:

- LLF is best possible using OPT as a lower bound
- better guarantees for LLF via a better lower bound?
- more sophisticated algorithms?
 - Thm 3 is reduction from Partition
 - existence of a (F)PTAS?

General Graphs

Open:

- for general latency fns, fixed β : a strategy inducing an equilibrium w/cost = $f(\beta) \times \text{opt}$
 - $1/\beta$ not achievable in general graphs (!)
 - maybe $2/\beta$? (or $O(n)$)
- for linear latency fns, a strategy w/cost $< 4/3 \times \text{opt}$
 - e.g., is $8/7$ achievable for $\beta=1/2$?