

Stackelberg Scheduling Strategies


Tim Roughgarden
Cornell University

Machines and Latencies

The Model:

- m machines $1, 2, \dots, m$
- A rate 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


Job Assignments:

- x_i = amount of jobs assigned to machine i
- vector $x \Leftrightarrow$ assignment of all jobs

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

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

Solution #1: relax our notion of “near-optimal”

- **Theorem [RT00]:** cost of Nash assignment = cost of OPT at rate $2r$

Solution #2: restrict to linear latency functions ($l_i(x) = a_i x + b_i$)

- **Theorem [RT00]:** cost of Nash assignment = $4/3$ times cost of OPT (at same rate r)

Taming Selfishness through a Manager

Solution #3: change the model!


- 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]

Examples


Strategy #1:



1

x

Strategy




1

x

Induced Equilibrium
(same as **Nash eq**)


Strategy #2:



1

x

Strategy



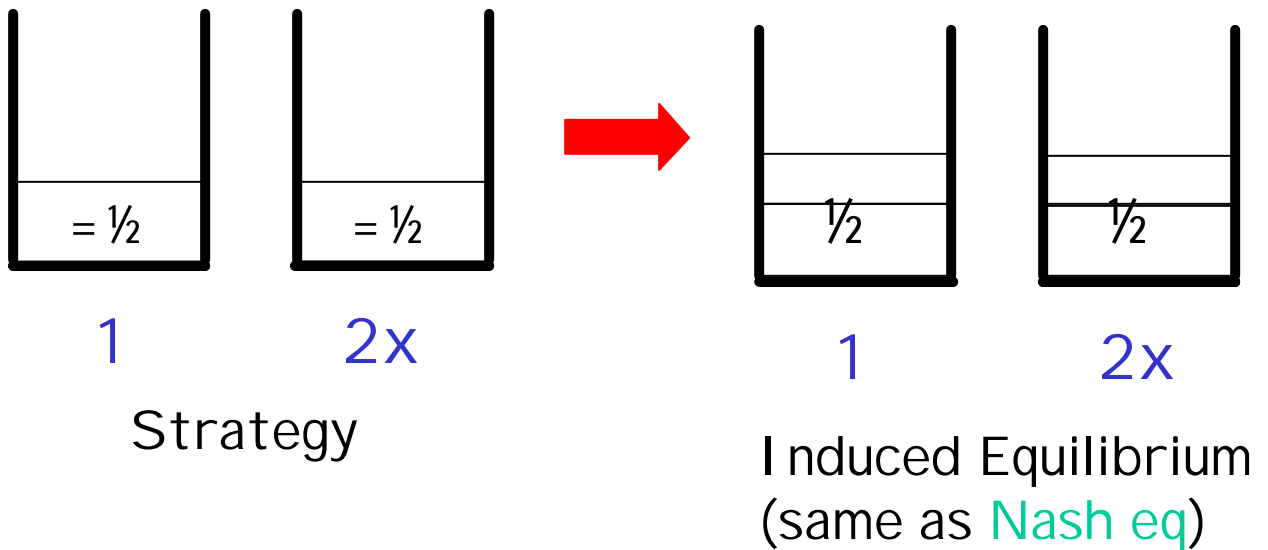
1

x


Induced Equilibrium
(same as **optimum!**)

Examples

Any Strategy:



Optimal Assignment (cost = ?):



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 = $(1/\beta) \times$ cost of optimal assignment.

β = fraction of centrally assigned jobs

Note:
 $(1/\beta) \times \text{OPT}$
is best possible



(Induced eq = Nash eq for any strategy)


Our Results - Linear Latency Functions

Theorem 2: Can efficiently compute a strategy inducing an equilibrium with cost = $[4/(3+\beta)]$ \times cost of optimal assignment.


β = fraction of centrally assigned jobs

Note:

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



1



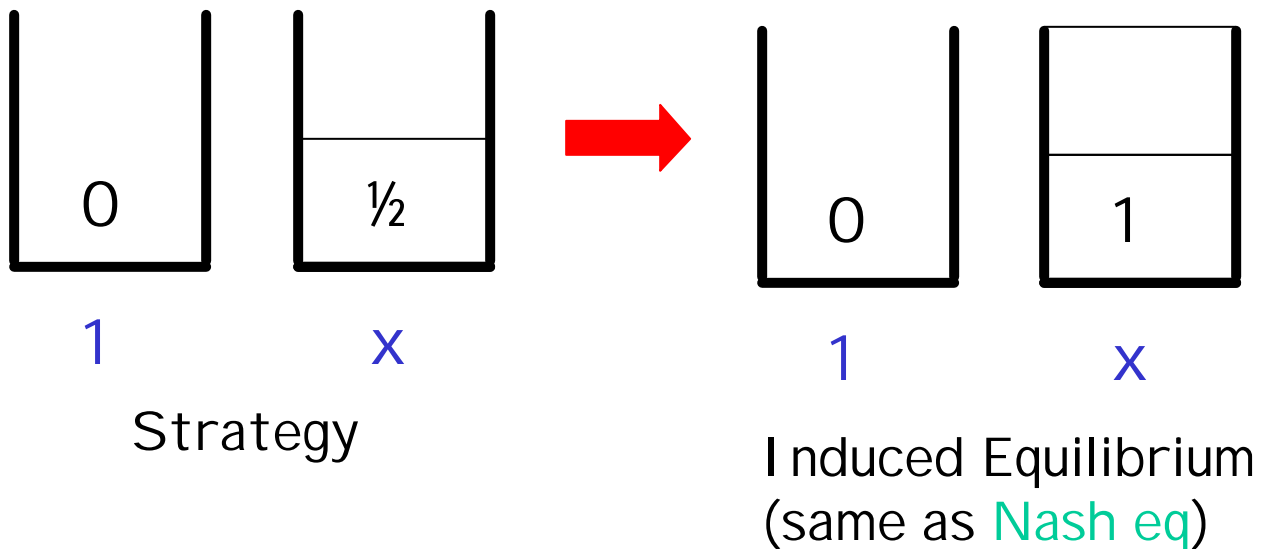
$x/(1-\beta)$

(Induced eq = Nash eq for any strategy)

What makes a strategy (in)effective?

The Aloof Strategy

- ignore selfish users, assign centrally controlled jobs in cheapest possible way




- induced eq costs 1, OPT = $\frac{3}{4}$

What makes a strategy (in)effective? (con'd)

The **Scale** Strategy

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



- **β** induced eq costs 1, OPT = $5/6$

What have we learned?

Morals:

- **avoid** machines that selfish users will (over)use anyways
- assign centrally controlled jobs to machines **least attractive** to the selfish users
(\Leftrightarrow edges with large latency)
- don't **oversaturate** any machines

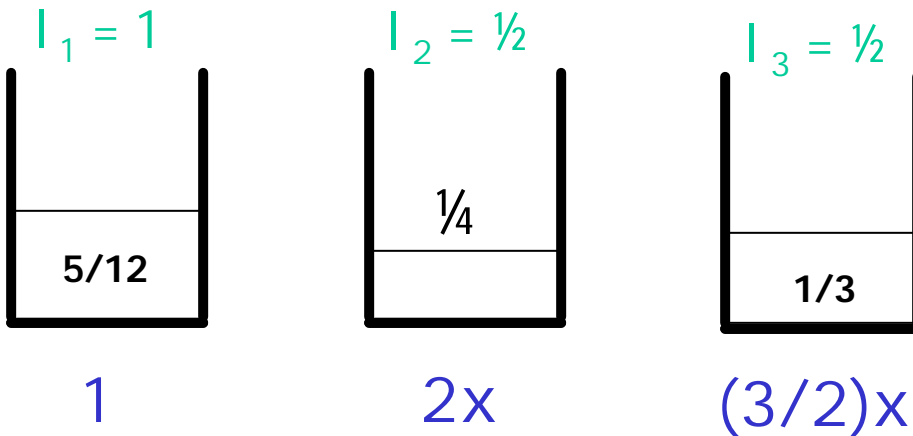
The LLF Strategy

Largest Latency First (LLF):


- compute opt assignment x of all jobs
- for each machine i , in order of decreasing $l_i(x_i)$ values:
 - if $\geq x_i$ centrally controlled jobs remain, assign x_i jobs to i (saturate machine i)
 - else assign remaining jobs (if any) to i

Example

Optimal Assignment: ($r=1$)



The LLF Strategy: ($\beta = 1/2$)




LLF with General Latency Functions

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

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

Proof Idea

Induced Equilibrium (w.r.t. LLF):




cost of induced eq and OPT are equal here

will apply IH on these machines


\$64K Question: Is our inductive guarantee strong enough?

Proof of Theorem 1



machines have
latency $\geq L$

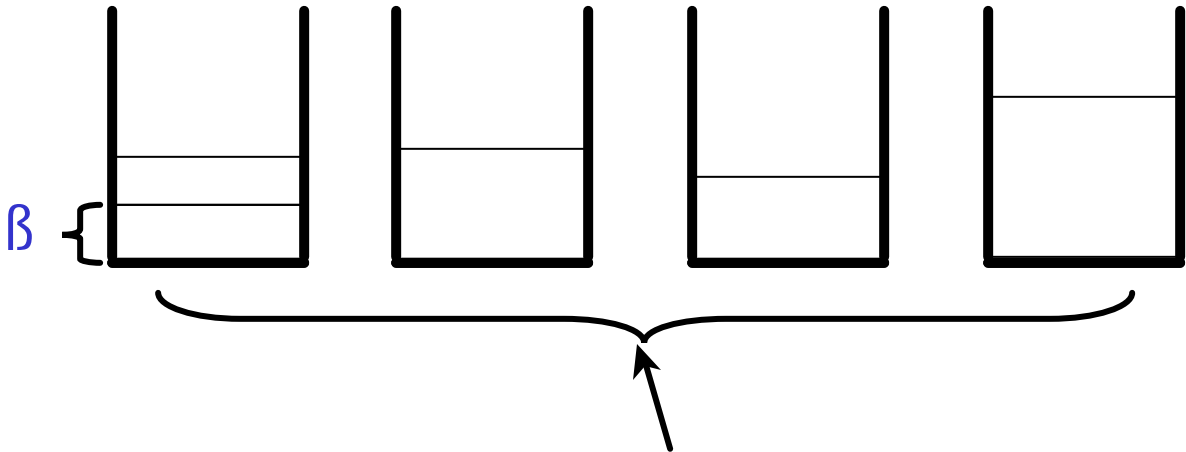
- $\mu = \#$ of jobs
- OPT pays
 \geq we pay
 $\geq \mu L$



all machines same
latency = L

- we pay $(1-\mu)L$
- $\beta' = \frac{(\beta - \mu)}{(1-\mu)}$
- I H \mathbb{P} OPT pays
 $\geq \beta' \times (1-\mu)L$
 $= (\beta - \mu)L$

Bug Fix



all machines same latency = L

Problem: Can't recurse!

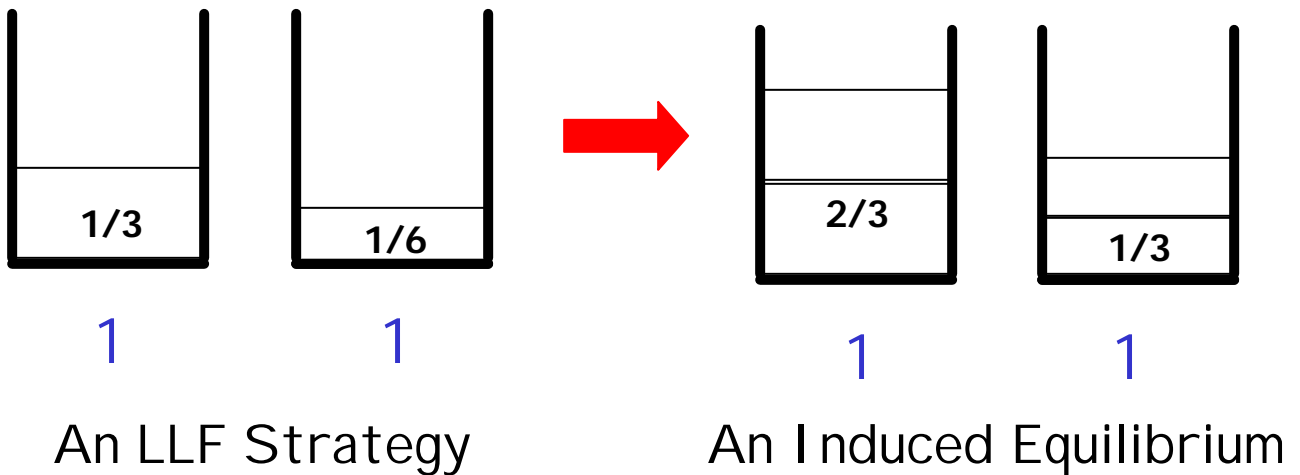
Fix: LLF failed to saturate machine 1

- Machine 1 is high-latency ($\geq L$)
- our solution costs L (assume $r=1$)
- OPT puts $\geq \beta$ jobs on machine 1

\mathbb{P} OPT pays $\geq \beta L$

Bug Fix (con'd)

Problem: have we really covered all the cases?



Lemma: Only a problem with locally constant latency functions.


- can reassign selfish users, get a “nicer” induced equilibrium

LLF with Linear Latency Functions

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

Main difficulty:

- previous argument too weak for:



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

Computing Optimal Strategies

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:** optimal, Nash assignments

Sketch of Reduction


Idea: view the best SS as solving the optimization problem **P**:

$$\max \quad S_i \max (0, x_i - \text{Nash}_i)$$

$$\text{s.t.} \quad S_i x_i = \beta$$


i.e., don't oversaturate $\rightarrow 0 = x_i = \text{OPT}_i \quad \text{all } i$

- **Intuition:** Strategy should look more like OPT than Nash



Reduction (con'd)

Partition: given n numbers a_1, \dots, a_n summing to A , is there a subset with sum $A/2$?



- amount of controlled jobs = $A/2$
- saturate machine $i \iff$ pick a_i
- "yes" Instance \mathcal{P} has value $A/4$
- "no" Instance \mathcal{P} has value $< A/4$

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?
 - existence of a (F)PTAS?

Also:

- $\min \beta$ sufficient to recover OPT
 - see also [KLO 97]

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 $\beta=1/2$?