New Models and Algorithm for Throughput Maximization in Broadcast Scheduling

Jian Li (University of Maryland)

Joint work with Chandra Chekuri, Sungjin Im, Benjamin Moseley (UIUC)

Samir Khuller, Richard McCutchen, Louiqa Raschid (UMD)

Avigdor Gal (Technion)

Broadcast Scheduling

Problem definition:

- Given a set of pages $P=\{p_1,p_2,\ldots,p_n)$
- Time is slotted, $T=\{1,2,\ldots,T\}$
- Each client sends a request r for page p, *with* release time a_r and deadline d_r
- The server broadcasts one page p in a time slot t, and all requests r of page p with $t \in [a_r, d_r]$ can be satisfied

Broadcast Scheduling

- Example:

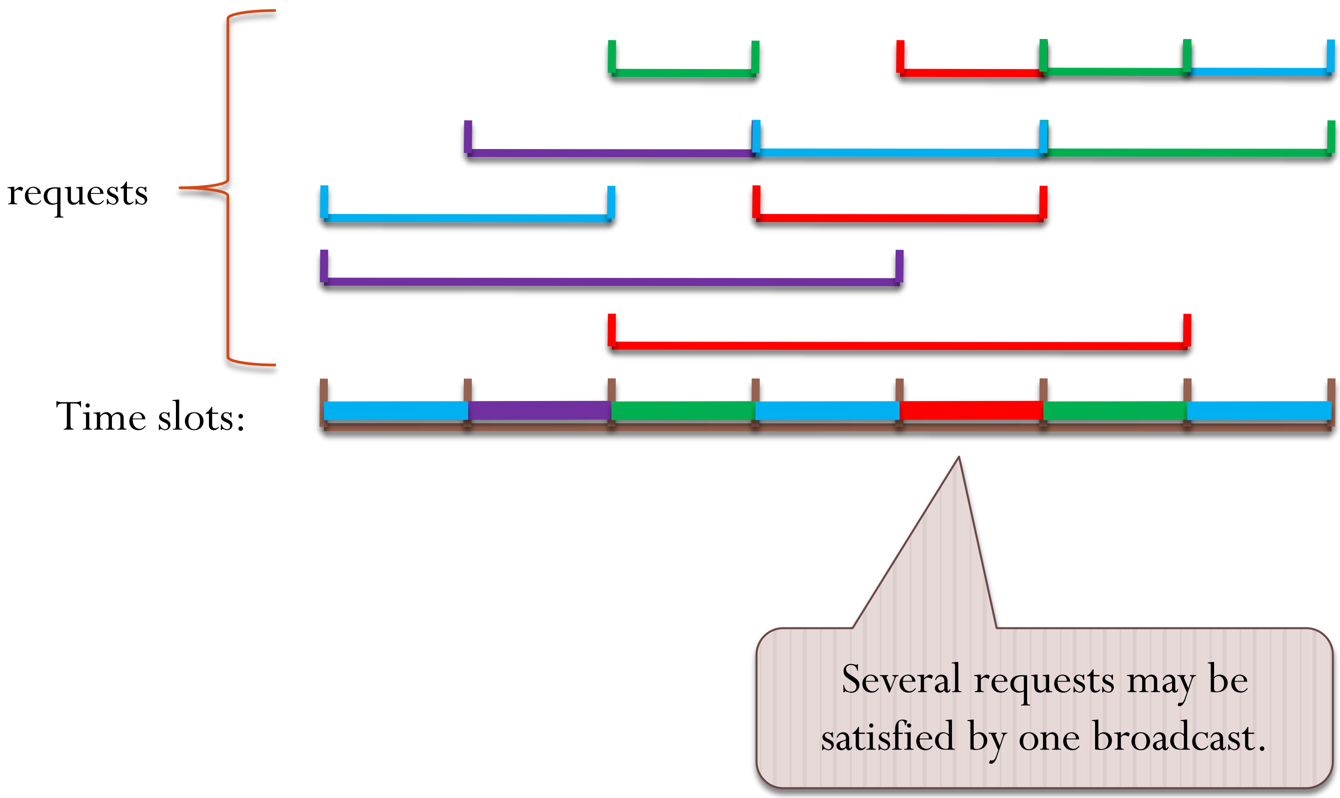

Broadcast Scheduling

- Traditional objectives.
  - Hard deadlines:
    - Throughput maximization (MAX-THP)
    - …
  - No deadlines:
    - Minimizing the max response time.
    - Minimizing the flow time (i.e., avg. response time).
    - ….
  - NP-hardness [Chang et al. 08].

Motivation

- Each client request the reading of some sensor at some time. The server can probe one sensor in a time slot.
  - A client requests the temperature reading at 5:30PM. She may be satisfied with a reading at 5:33PM. A reading at 5:40PM may be still useful, but not as much. But a reading in 6:00PM is useless.
- Traditional objectives are not sufficient in this example.
  - Minimizing response time ignores deadlines.
  - Minimizing throughput ignores the latency of satisfied requests.
- We capture this in two approaches.
  - A general time-dependent profit function.
  - Tradeoff between completeness and latency.

Profit Maximization

- A generalization of throughput maximization: Profit Maximization (MAX-PFT)
  - A time –dependent profit function $g_r(t)$ for each request *r*.
  - If a request is satisfied multiple times, we take the maximum one.
  - A more nuanced view of "satisfying " a request.

The profit function for a request:

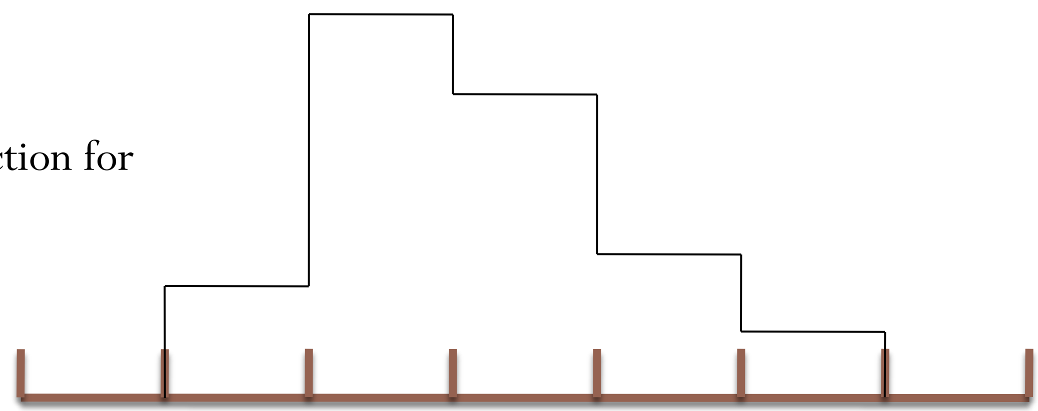

Our Results

- Offline setting.
  - A (1-1/e)-approximation for MAX-PFT.
  - A 3/4-approximation for MAX-PFT when the profit functions are unimodal.
    - MAX-THP offline: A 3/4 -approximation [Gandhi et al. '06].

- Online settings.
  - An *s*-speed (*1+1/s*)–competitive algorithm for MAX-PFT.
    - MAX-THP online: A 1/2-competitive algorithm [Kim et al. '04].

Our Results

- Minimizing latency subject to completeness requirement.
  - A (3/4, 1)-approximation for the (completeness, latency) pair.
  - Note that both ratios are *in expectation*.
- Throughput Maximization with Relaxed Time Windows.
  - Suppose there is a fractional solution that satisfies all requests. We can find a schedule in polynomial time such that each request can be satisfied by right (or left) shifting the window by at most its length.

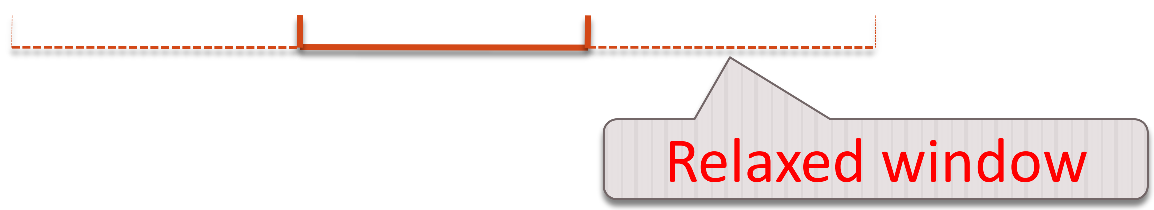

Our Results

- Offline MAX-THP:
  - 2-speed 1-approximation.
    - Such a result was known only if all request can be scheduled in a fractional solution [Chang et al. 08].
    - This directly implies a 2-approximation for MAX-THP.
- Minimizing the max response time
  - A (2-ε)-lower bound for randomized algorithms in the oblivious adversary model.
    - The same bound was only known for deterministic algorithm [Bartal et al. 00, Chang et al. 08].
    - Note that FIFO is 2-competitive [Bartal et al. 00, Chang et al. 08, Chekuri et al. 09].

Offline – Profit Maximization

- The slicing trick: Convert MAX-PFT to weighted MAX-THP

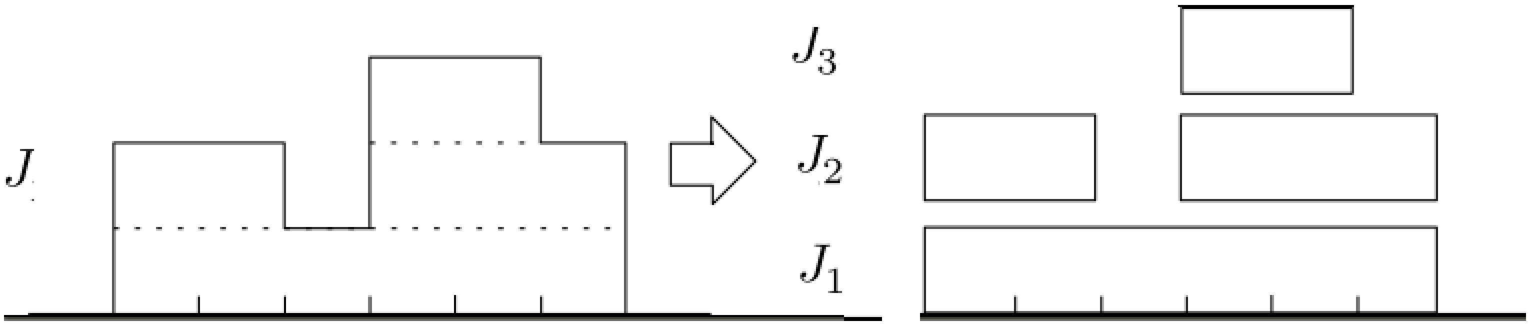

A general profit function

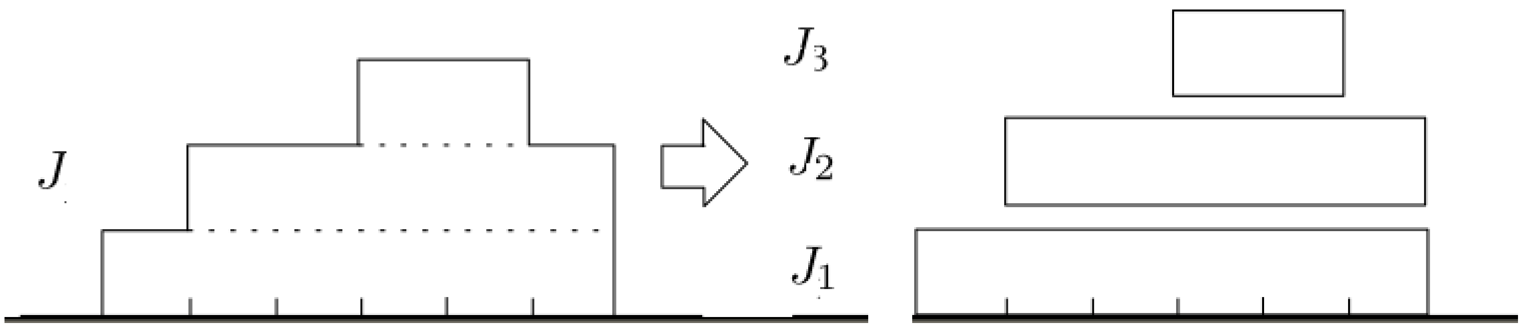

A unimodal profit function

Offline – Profit Maximization

THM: A 3/4 –approximation for MAX-PFT when the profit functions are unimodal.

Proof: The slicing trick and the 3/4-approximation for weighted MAX-THP.

THM: A (1-1/e) –approximation for MAX-PFT with general profit functions.

Proof 1: A simple independent rounding schema.

Offline – Profit Maximization

Proof 2:

(submodular maximization subject to a matroid constraint)

- f: $2^N \rightarrow R$ is a submodular function if

$$f(A+x) - f(A) \cdot \; f(B+x) - f(B) \;\; \forall B \subseteq A, x \in N$$

- Let N be $\{(p,t)\}_{p,t}$. The set of feasible solutions is a partition matroid.
- Let *Profit(S)* be the profit obtained by schedule *S* (*S*$\subseteq$*N*). *Profit(.)* is submodular.
- Submodular function maximization subject to a matroid constraint: (1-1/e)-approximation [Calinescu et al. ’07, Vondrak ’08, Chekuri et al. ’10].

Online – Profit Maximization

- **M**aximum **A**dditional **P**rofit **F**irst (MAPF):
  - At any time *t*, broadcast *s* pages which give the maximum additional profits.

THM: MAPF is an *s*-speed (*1+1/s*)-competitive online algorithm for MAX-PFT.

The analysis is tight.

THM: For any $\varepsilon>0$ and $s\geq 1$, MAPF is not *s*-speed $(1+1/s-\varepsilon)$-competitive.

Online – Profit Maximization

Proof (sketch):

$\Delta(t)$: the increase of the so-far-gained profit by OPT over the final profit by MAPF.

$$\Delta(t)=\Sigma_{p,i}\delta_{p,i}(t).$$

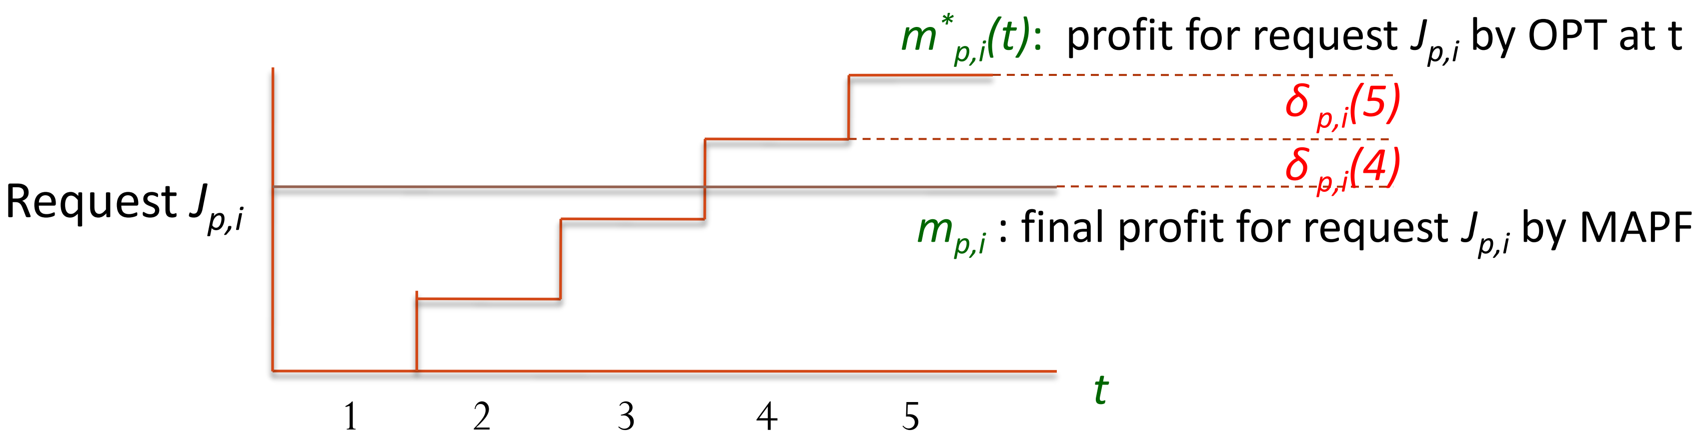

Online – Profit Maximization

Proof (sketch):

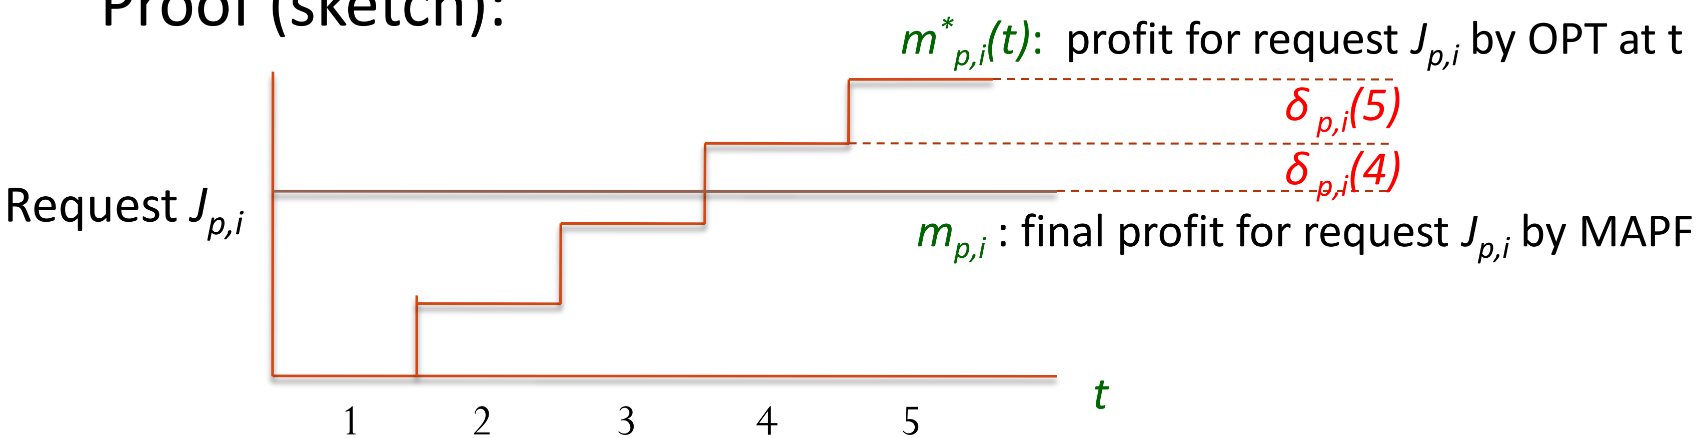

$\Delta(t)=\Sigma_{p,i}\delta_{p,i}(t).$

$OPT \le MAFP + \Sigma_t \Delta(t)$ -- *by definition of* Δ.

$\le MAFP + \Sigma_t (1/s)\, \underline{\Sigma_{p,i}(m_{p,i}(t)-m_{p,i}(t-1))}$ --*next slides.*

$\le (1+1/s)\, MAFP$

Additional profit obtained by MAFP

Online – Profit Maximization

Proof (sketch):

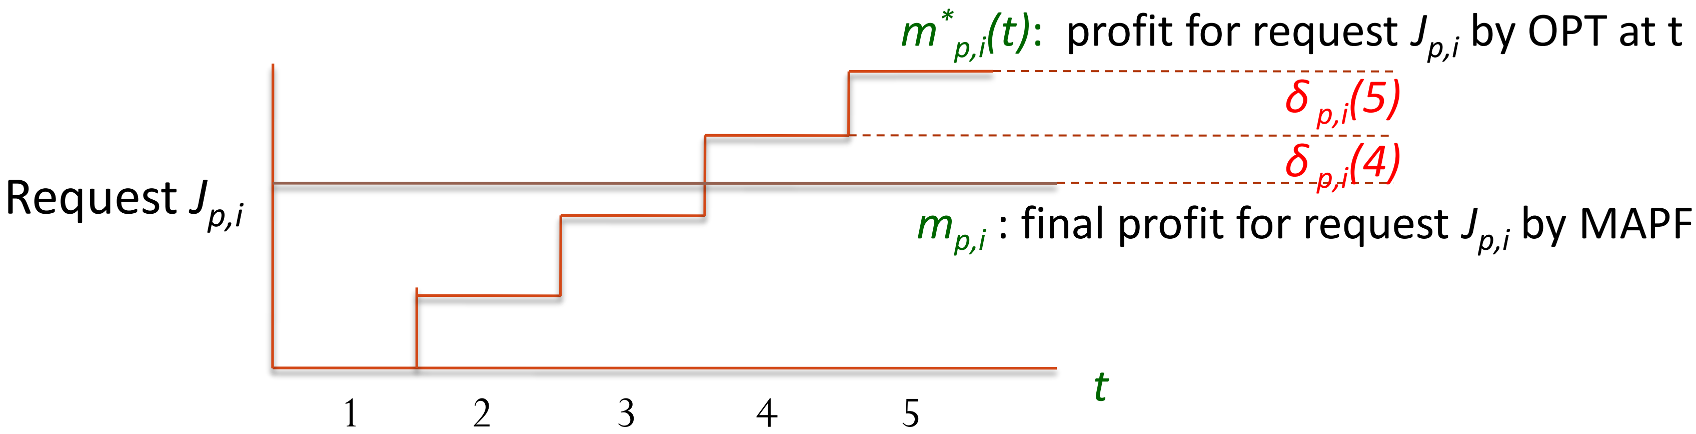

It suffices to show $\Delta(t) \le (1/s)\Sigma_{p,i}(m_{p,i}(t)-m_{p,i}(t-1))$

- Assume OPT broadcast q at time t and $\Delta(t)>0$.
- We can show MAPF does not broadcast q. O.w. $\Delta(t)=0$.
- $\Delta(t) \le \Sigma_i (m_{q,i}(t)-m_{q,i}(t-1))$.
- $(m_{q,i}(t)-m_{q,i}(t-1)) \le (m_{p,i}(t)-m_{p,i}(t-1))$ if MAPF broadcast p.

Open Problems

- Is it possible to get a 4/3-speed 1-approximation for MAX-THP. Note this would imply a 3/4-approximation for MAX-THP (matching the bound by Gandhi et al. ’06).
- Derandomizing the (3/4,1)-approximation for the (completeness, latency) pair.
- A better understanding of the completeness-latency tradeoff.

Thanks

Offline – Profit Maximization

Proof 1: (LP rounding)

$Y_p^{(t)}=1$: The server broadcasts *p* at time *t*.
$X_{p,i}=1$: The *i*th request of *p* is satisfied.

$$\text{maximize} \quad \sum_{p,i} w_{p,i} X_{p,i}$$

$$\text{subject to} \quad \sum_{t \in \mathcal{T}_{p,i}} Y_p^{(t)} \geq X_{p,i} \ \forall p, t,$$

$$\sum_p Y_p^{(t)} \cdot \ 1, \ \forall t,$$

$$X_{p,i} \in \{0,1\}, \forall p, t, \quad Y_p^{(t)} \in \{0,1\}, \forall p, t$$

Offline – Profit Maximization

Proof 1: (LP rounding)

Let $x_{p,i}$, $y_p^{(t)}$ be the optimal fractional solution.

Algorithm: (Independent rounding)

- At time *t*, choose *p* to broadcast with prob. $y_p^{(t)}$

It is not hard to show that $Pr(\text{request } J_{p,i} \text{ is satisfied}) \geq (1-1/e)\, x_{p,i}$

Throughput Maximization with Relaxed Time Windows

- The fractional solution corresponds to a flow.
- There is integral flow.

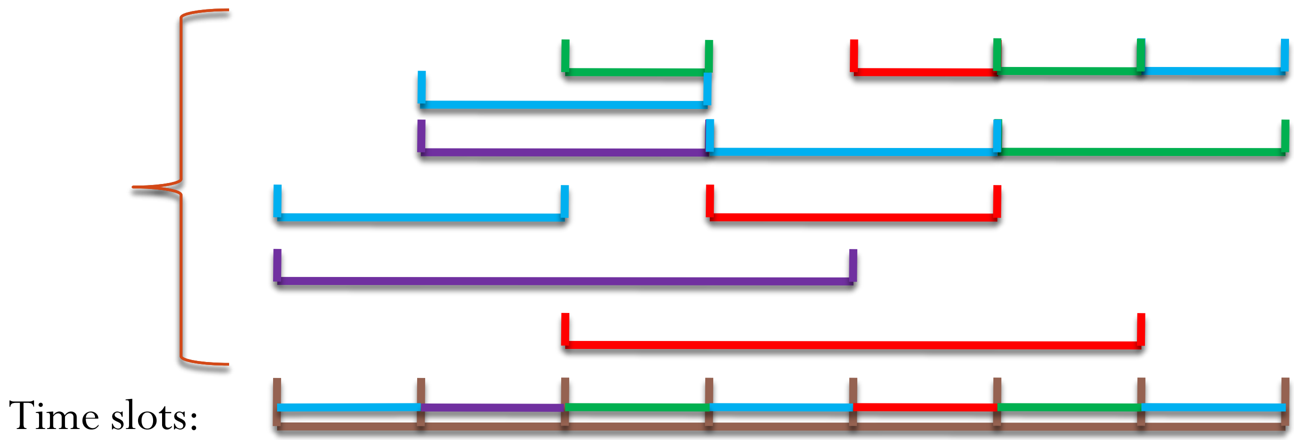

Throughput Maximization with Relaxed Time Windows

- Suppose there is a fractional solution that satisfies all requests.

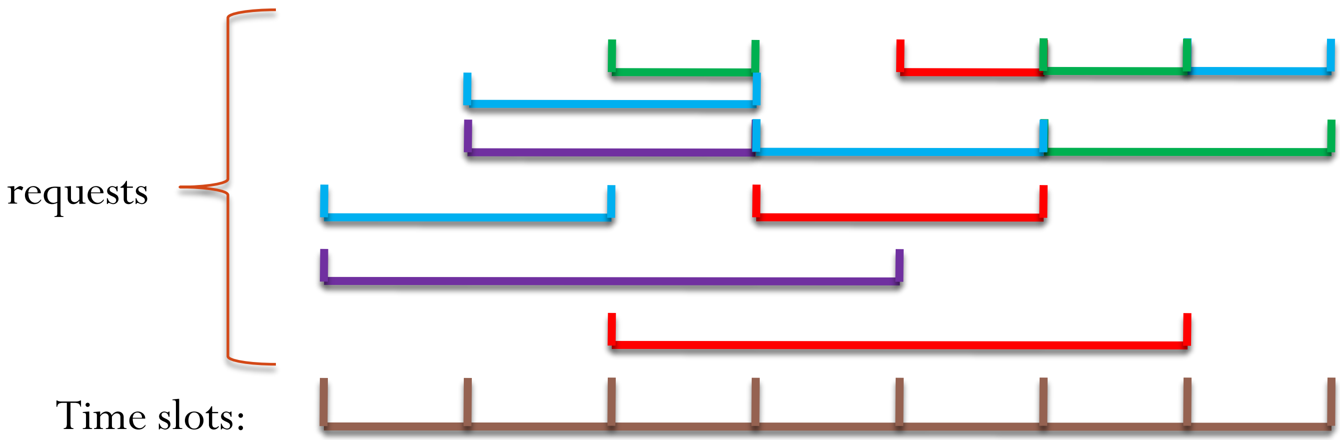

Throughput Maximization with Relaxed Time Windows

- For each page
  - Order the requests for page p in non-decreasing window length.
  - Insert each request as long as there is no overlap

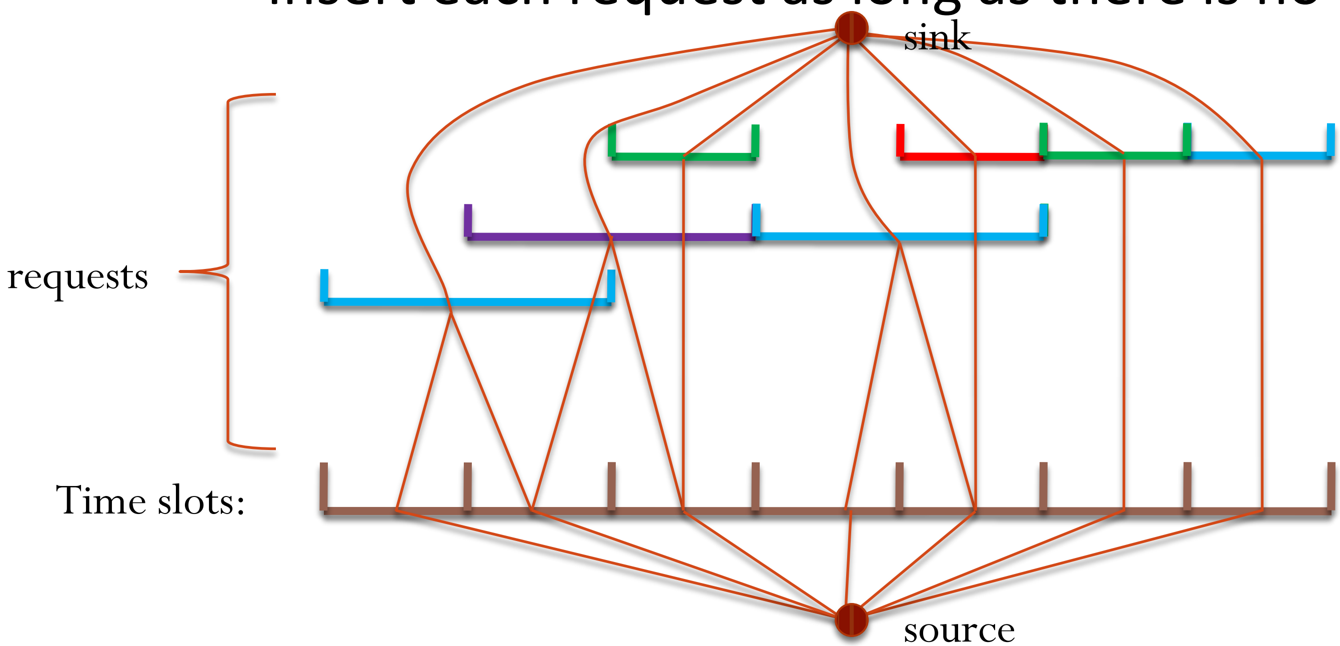

  - The fractional solution corresponds to a flow.

Throughput Maximization with Relaxed Time Windows

- The fractional solution corresponds to a flow.
- There is integral flow.