Aalto University
School of Electrical Engineering

Optimizing Flow-level Performance by Scheduling in Wireless Cellular Systems

Pasi Lassila (joint work with Samuli Aalto, Aleksi Penttinen, Noé Falzon, Jari Melasniemi, Prajwal Osti)

COMNET Department

Aalto University, School of Electrical Engineering

Outline

- Introduction to flow-level models
- Overview of 3 problem settings
  - Size-based multiuser scheduling in non-channel aware setting
  - Optimal trade-off between channel-aware scheduling and size-based scheduling
  - Intercell coordination with 2 base stations

What are flow-level models?

- Flow-level models are used to model the system from the point of view of elastic data traffic
  - Elastic data traffic corresponds to, e.g., TCP file transfers
  - The transmission rate varies during the lifetime of the flow
- Flow-level models are dynamic
  - Flows arrive according to some stochastic process (Poisson)
  - Flows have random service requirements (file sizes)
  - Instantaneous service rate of a flow depends on state of the system (e.g., number of other flows)
  - How this dependence is modeled typically contains idealizations

What are flow-level models?

- Performance at the flow level
  - Performance is expressed as throughput or flow delay (file transfer delay)
  - For example, mean flow delay would describe how long file transfers on the average last
- Importance
  - Users do not care about delays of individual packets, but only about the total time to transmit a file of a given size

- Flow-level models try to characterize the system at the time-scale where users experience the performance

M/G/1-PS queue

- Consider a single link with constant capacity C
  - In a single server processor sharing (PS) queue the capacity C of the server is equally shared between the customers in system
  - If there are n flows in system each receives service at the rate C/n
  - Flows are served in parallel (there is no queueing!)
  - Round robin (RR) service can be approximated by PS when service quantum in RR is small relative to flow durations (cf. time slot length in 3G/LTE systems)

- M/G/1-PS queueing model
  - Arrival process is Poisson with intensity λ
  - Flow sizes X obey a general distribution with mean $\mathrm{E}[X]$
  - Service discipline is PS
  - Well-known model with many nice analytical properties (robust)

Optimal size-based scheduling

- Objective: minimize mean delay
- Idea: get rid of small flows as soon as possible
  - Minimizing number of flows in system also minimizes delay
- SRPT – Shortest Remaining Processing Time first
  - Preemptive policy that always serves the user with the least amount of service time left
- SRPT is optimal for M/G/1 queue (Schrage, 1962)
  - It minimizes sample path-wise the queue length for any policy and any service time distribution
  - Thus, it also minimizes the mean queue length and also, by Little's result, mean flow delay
  - However, SRPT may not be practical (remaining work is unknown)

General goal

- Model wireless systems at flow-level
  - Involves abstractions (simplifications)

- Objective
  - To learn what queueing/scheduling theory can teach us about resource allocation and dynamic traffic

Multiuser scheduling in HSDPA

Introduction

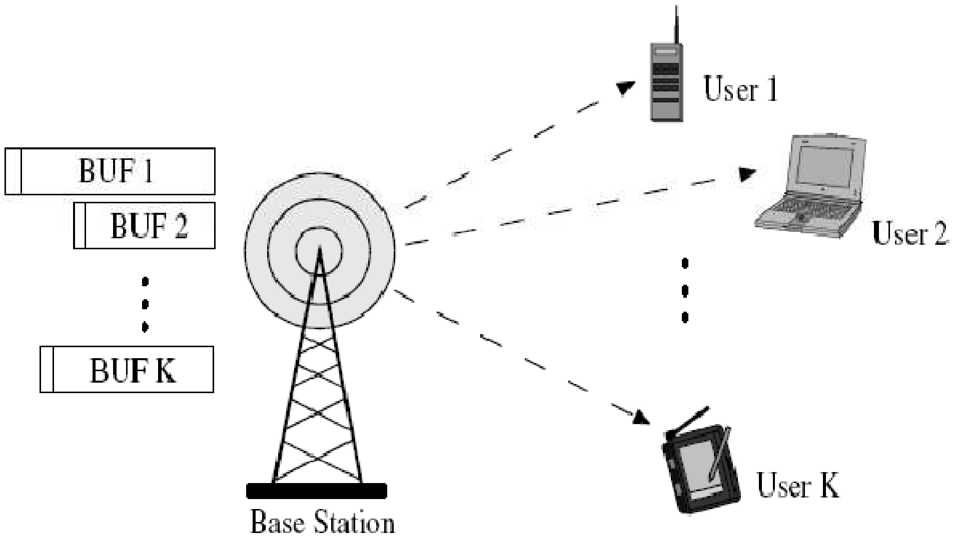

- Users randomly in cell area
- HSDPA
  - Base station schedules users in every time slot
  - Scheduler allocates CDMA codes
  - Each user has a constraint on number of codes it can use (processing power limitation) → multiuser scheduling
- Users are downloading files with random sizes
  - Objective: minimize mean file transfer delay using information on the rates of users (distances) and remaining file-sizes (cf. SRPT)

User model

- Single cell (unit circle) with users uniformly in cell area
- Users are classified based on their terminal category
  - Terminal category characterizes how many codes user can use
  - Class-i users are able to use k_i codes
- BS has altogether K codes
  - Typically, $k_i < K \rightarrow$ BS has to select multiple users
  - In practice, $K = 30$ and $k_i = \{5, 10, 15\}$
- Users/flows in class i arrive according to a Poisson process with rate λ_i, each flow has random size X (bits)

Channel model

- Time-slot duration very short compared with time scale of flow-level dynamics (time scale separation)

- Rate $c(r)$ represents the mean throughput of a user at distance r if all codes are allocated to him, for example

$$c(r) = \begin{cases} c_0, & r \leq r_0, \\ c_0 \left(\frac{r_0}{r}\right)^\alpha, & r > r_0. \end{cases}$$

  – Service time distribution depends on $c(R)$ and X (distance R and size X are independent)

- Codes are perfectly orthogonal
  – Rate per code is $c(r)/K$

Scheduling policies

- Codes correspond to servers
  - If any user can use all codes → M/G/1 queue
  - Due to terminal constraints, the system is a multiserver queue with heterogeneous server allocation contraints (k_i)
- Processor sharing, PS (= round-robin at time-slot level)
  - Fair baseline policy
  - No information about sizes or rates
- SRPT for multiserver models
  - Requires exact knowledge of both size and rate
  - Surprisingly, optimal policy even for M/G/2 is not known!
  - SRPT-FM optimal for standard multiserver queue (one server per flow) **with no new arrivals** (Pinedo, 2008)
  - Our model is not even a standard multiserver queue ⇒ modify SRPT-FM principle for our model in the dynamic setting

Results

- Compared SRPT scheduling against PS (using simulations)
  - Modified SRPT-FM: dynamic ranking of flows depending on amount of unallocated resources and resource requirements

- SRPT decreases significantly (~ 30%) mean delays at moderate loads (~70%) with reasonable system parameter values
  - Also, fairness between sizes (large vs. small) and classes remained good

- Publications
  - J. Melasniemi, "Size-based scheduling under terminal constraints in cellular systems", M.Sc. thesis, Aalto University School of Science and Technology, August, 2010.
  - J. Melasniemi, P. Lassila and S. Aalto, "Minimizing file transfer delays using SRPT in HSDPA with terminal constraints", in Proc. of 4th Workshop on Network Control and Optimization (Net-Coop), Ghent, Belgium, 2010.

Optimal trade-off between channel-aware scheduling and size-based scheduling

The fundamental trade-off

- Opportunistic (channel-aware) scheduling:
  - In each time-slot, select the user that has instantaneously good channel
  - Aggregate mean service rate increases with number of users (opportunistic gain, mutiuser diversity gain)

- Size-based scheduling (SRPT)
  - By favoring small flows we minimize number of flows
  - However, we lose some of opportunistic gain

The general problem

- Queuing model with state dependent service rates
  - Service rates can be ”tuned” by the time-slot level scheduler
  - Given n users, the set of feasible rate vectors is known as the capacity region
  - To minimize delay we must select the optimal rate vector for every state
  - Very difficult optimization problem

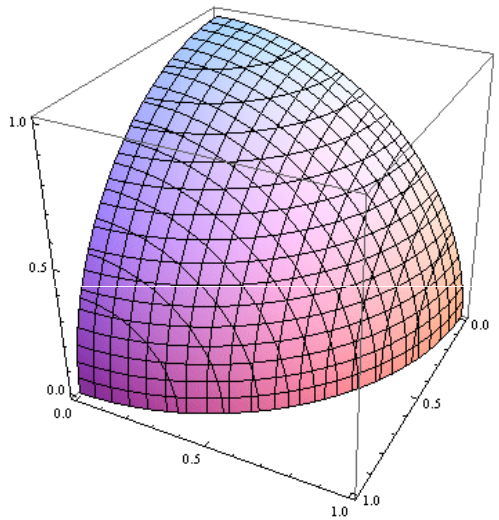

- Previous work
  - Dynamic setting: no optimality results, only stability has been studied
  - Transient setting: optimality results for symmetric upper bounded capacity regions

The transient model

- At time 0
  - We have n flows with sizes $s_1 > \ldots > s_n$
  - No new flows arrive after time 0

- Operation proceeds in stages
  - At each stage $k = 1,\ldots, n$, we select a service rate vector $\mathbf{c}_k$ from capacity region C_k
  - C_k are given and we assume that they are compact and symmetric (all users have statistically identical channels)

- Objective: Select optimal service rate vector $\mathbf{c}^*_k$ at each stage k so that mean delay is minimized

Main result

- Recursive solution for $\mathbf{c}^*_k$, $k = 1,\ldots, n$

$$g_1(c_1) = \frac{1}{c_1},$$
$$G_1^* = g_1(c_1^*) = \min_{c_1 \in \mathcal{C}_1} g_1(c_1),$$

$$g_k(\mathbf{c}_k) = \frac{1}{c_{kk}}\left(k - \sum_{j=1}^{k-1} c_{kj} G_j^*\right),$$
$$G_k^* = g_k(\mathbf{c}_k^*) = \min_{\mathbf{c}_k \in \mathcal{C}_k} g_k(\mathbf{c}_k), \quad k = 2, \ldots, n.$$

if $G^*_1 < \ldots < G^*_n$

- Properties
  - Independent of sizes!
  - Obeys SRPT-FM principle (shortest job served with highest rate, and so on)

Example with 2 jobs

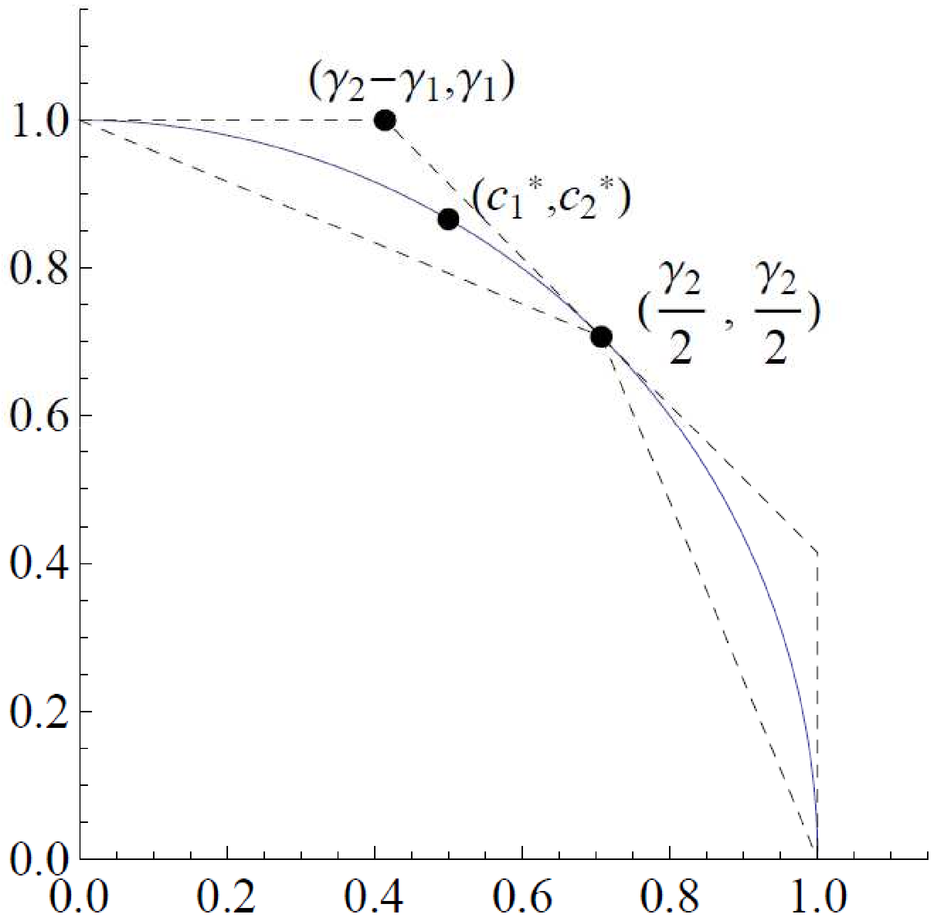

Results

- First optimality result for this difficult trade-off
  - Additional results on (upper/lower) bounds for the delay

- New result
  - Above capacity region was assumed to be known
  - Optimal result holds (i.e., G^*_i are increasing) for practically any channel assuming symmetric users
  - Characterization of the time-slot level opportunistic scheduler to realize the optimal rate vectors

- Publications

  - S. Aalto, A. Penttinen, P. Lassila and P. Osti, “On the optimal trade-off between SRPT and opportunistic scheduling”, in Proceedings of ACM SIGMETRICS, 2011, to appear.
  - S. Aalto, A. Penttinen, P. Lassila and P. Osti, “Optimal size-based opportunistic scheduler for wireless systems”, submitted.
  - S. Aalto and P. Lassila, “Flow-level stability and performance of channel-aware priority-based schedulers”, in Proc. of NGI 2010, Paris, France, 2010.

Intercell coordination with 2 base stations

Intercell coordination

- Idea: A base station does not interfere when it is turned off
  - Depending on state, it may be useful to turn off a base station
- Flow-level analysis
  - Dynamic setting
  - Base stations correspond to servers (multiserver models)
  - Switching base stations on/off makes service rates coupled
- Recent work
  - T. Bonald, S. Borst, and A. Proutiere, “Inter-cell coordination in wireless data networks,” Euro. Trans. Telecoms., vol. 17, 2006.
  - I. Verloop and R. Nuñez-Queija, “Asymptotically optimal parallel resource assignment with interference,” Queuing Syst., vol. 65, 2010.
  - Our work is based on (Verloop& Nuñez-Queija, 2010)

Queueing model

- 2 base stations

- One class of users per BS
  - Same service requirements, exponentially distributed
  - Different Poisson arrival rates λ_i

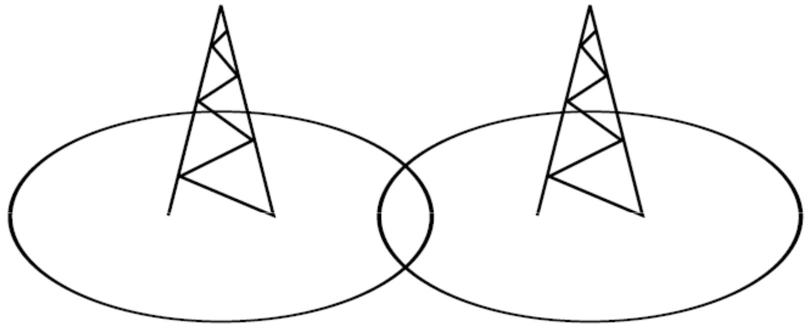

- 2 service rates
  - BS alone: μ_0
  - Both BS's on: $\mu_1 \, (\leq \mu_0)$

- Markovian model
  - System state: number of users in both classes (n_1, n_2)

Objective

- Service modes
  - 3 modes: $(\mu_0,0)$, $(0, \mu_0)$, (μ_1,μ_1)
  - (i) low interference: $2\mu_0 \le \mu_1$
  - (ii) high interference: $2\mu_0 > \mu_1$

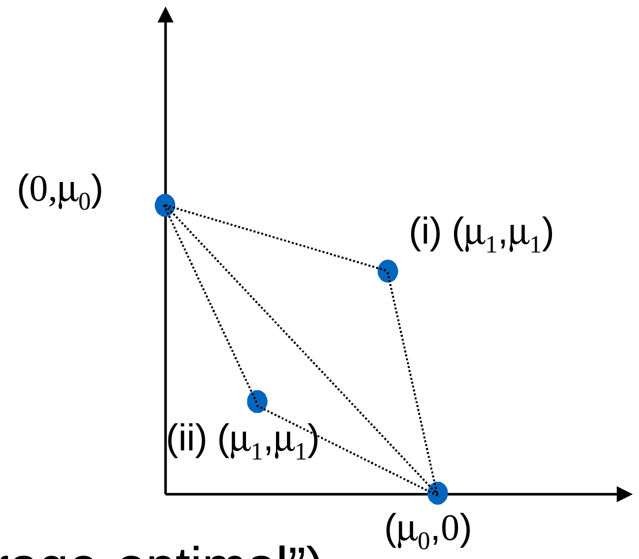

- Objective
  - Find a policy that minimizes the average total number of users ("average-optimal")

- Policy:
  - Which mode to use in each state (n_1, n_2) of the system?
  - Can be solved using theory of Markov decision processes

Results

- Optimal policy
  - Low interference: always keep both stations on, except when one queue is empty, in which case, serve only the other one
  - High interference: never serve both queues at the same time, and if one queue is empty, serve the other one

- Publications
  - N. Falzon, “Optimal inter-cell coordination for elastic data traffic”, M.Sc. thesis, Aalto University School of Science and Technology, September, 2010.
  - Thesis contains much more work …