Aalto University
School of Electrical Engineering

Opportunistic scheduling with flow size information for Markovian time-varying channels

Samuli Aalto, Pasi Lassila, Prajwal Osti
Department of Communications and Networking
Aalto University, Finland

ECQT
18-20 July 2016
Toulouse, France

Outline

- Introduction
- Whittle index approach
- Opportunistic scheduling problem
- Our contribution
- Numerical illustrations
- Summary

Research problem

- Downlink data transmission in a cellular system
  - traffic = elastic flows
  - file transfers using TCP
  - remaining file sizes known
- Time-varying channels of users
  - channel states known
- Optimal scheduler for flow-level performance?
  - minimizing the mean file transfer time

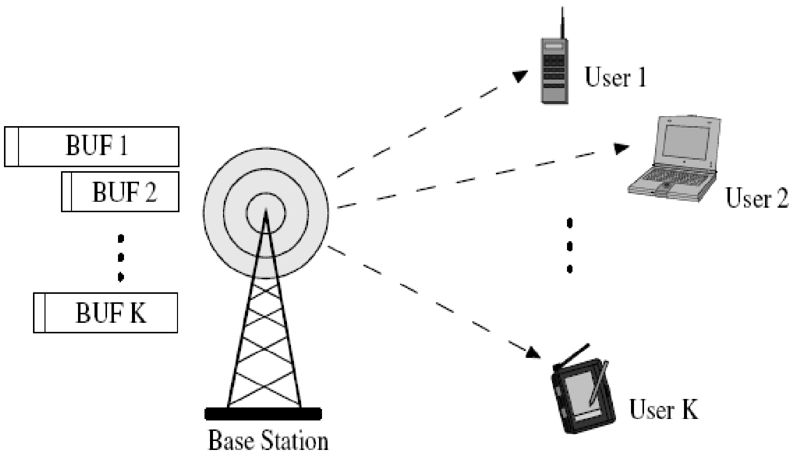

Source: Hu et al. (*ComNet* 2004)

Two approaches to solve the problem

- Time-scale separation
  - allows to solve the optimization problem exactly
  - applicable for the homogen. case
  - … but intractable in the general case with heterogeneous users
  - Sadiq and de Veciana (*ITC* 2010)
  - Aalto et al. (*Sigmetrics* 2011)
  - Aalto et al. (*QS* 2012)

- Whittle index approach
  - applies restless multi-armed bandits
  - tractable in the general case with heterogeneous users
  - … but solves the optimization problem just heuristically
  - Ayesta et al. (*PEVA* 2010)
  - Jacko (*PEVA* 2011)
  - Cecchi and Jacko (*Sigmetrics* 2013)
  - Taboada et al. (*ITC* 2014)
  - Taboada et al. (*PEVA* 2014)
  - Aalto et al. (*Sigmetrics* 2015)
  - Cecchi and Jacko (*PEVA* 2016)
  - Aalto et al. (*QS* 2016)

Outline

- Introduction
- Whittle index approach
- Opportunistic scheduling problem
- Our contribution
- Numerical illustrations
- Summary

Multi-armed bandit

"Las Vegas slot machines". Licensed under CC BY-SA 3.0 via Wikipedia - http://en.wikipedia.org/wiki/File:Las_Vegas_slot_machines.jpg#/media/File:Las_Vegas_slot_machines.jpg

Multi-armed bandit problem

- Problem:
  - Assume there are K discrete-time bandit processes
  - If chosen at time t, the bandit process evolves as a Markov process; otherwise its state is frozen until the next time slot $t+1$
  - If bandit i is chosen when in state x_i, a reward of $r_i(x_i)$ is earned
  - Given the states x_i, choose one bandit
- Answer:
  - Calculate the Gittins index $G_i(x_i)$ separately for each bandit i
  - Choose the bandit i^* with the highest Gittins index
  - Gittins and Jones (1974), Gittins (1989), Gittins & al. (2011)
- Note:
  - ”It was by no means evident that the optimal policy would take the form of such an *index policy*, and certainly not how the index should be calculated” Whittle (*JAP* 1988)

Restless bandit problem (1)

- Original problem:
  - Assume there are K discrete-time restless bandit processes
  - If chosen at time t, the bandit process evolves as a Markov process; otherwise its state evolves according to another Markov process
  - If bandit i is chosen when in state x_i, a reward of $r_{i,1}(x_i)$ is earned; otherwise another reward of $r_{i,2}(x_i)$ is earned
  - Given the states x_i, choose one bandit
- Relaxed problem:
  - Given the states x_i, choose bandits so that one bandit is chosen per time slot on average (in the long run)
  - Whittle (*JAP* 1988)

Restless bandit problem (2)

- Answer to the relaxed problem:
  - Consider the separable Lagrangian version of the relaxed problem
  - Show indexability separately for each bandit i
  - Calculate the Whittle index $W_i(x_i)$ separately for each bandit i
  - Choose all those bandits with the index greater than a threshold
  - Whittle (*JAP* 1988)
- Heuristic answer to the original problem:
  - Choose the bandit i^* with the highest Whittle index
  - Whittle (*JAP* 1988)
- Note:
  - In the multi-armed bandit problem: Whittle index = Gittins index

Outline

- Introduction
- Whittle index approach
- Opportunistic scheduling problem
- Our contribution
- Numerical illustrations
- Summary

Opportunistic scheduling with IID channels (1)

- Problem:
  - Assume there are K users with geometric file sizes X_i (prob. μ_i)
  - Channel states $R_i(t)$ are independent IID variables
  - Holding costs are accrued with rate c_i for any uncompleted flow i
  - Given the channel states r_i, choose one flow until all flows completed
- Heuristic answer:
  - Consider the separable Lagrangian version of the relaxed problem
  - Show indexability separately for each flow i
  - Calculate the Whittle index $W_i(r_i)$ separately for each flow i
  - Choose the flow i^* with the highest Whittle index
  - Ayesta et al. (*PEVA* 2010)
- Generalizations:
  - Taboada et al. (*ITC* 2014, *PEVA* 2014)
  - Aalto et al. (*Sigmetrics* 2015, *QS* 2016)

Opportunistic scheduling with IID channels (2)

- Result (for two-state channels):
  Primary Whittle index for a flow with channel state r is given by

$$W(r) = \begin{cases} \infty, & r = r^{\mathrm{g}} \text{ ("good" channel)} \\ \dfrac{c\, r^{\mathrm{b}}}{P\{R = r^{\mathrm{g}}\}(r^{\mathrm{g}} - r^{\mathrm{b}})}, & r = r^{\mathrm{b}} \text{ ("bad" channel)} \end{cases}$$

  Secondary (tie-breaking) Whittle index

$$\tilde{W}(r^{\mathrm{g}}) = c\mu\, r^{\mathrm{g}}$$

  – Ayesta et al. (*PEVA*, 2010)

Opportunistic scheduling with Markov channels (1)

- Problem:
  - Assume there are K users with geometric file sizes X_i (prob. μ_i)
  - Channel states $R_i(t)$ are two-state discrete-time Markov processes
  - Holding costs are accrued with rate c_i for any uncompleted flow i
  - Given the channel states r_i, choose one flow until all flows completed
- Heuristic answer:
  - Consider the separable Lagrangian version of the relaxed problem
  - Show indexability separately for each flow i
  - Calculate the Whittle index $W_i(r_i)$ separately for each flow i
  - Choose the flow i^* with the highest Whittle index
  - Jacko (*PEVA* 2011)
- Generalizations:
  - Cecchi and Jacko (*Sigmetrics* 2013, *PEVA* 2016)
  - Aalto et al. (*ECQT* 2016)

Opportunistic scheduling with Markov channels (2)

- Result (for two-state channels):
  Primary Whittle index for a flow with channel state r is given by

$$W(r) = \begin{cases} \infty, & r = r^{\mathrm{g}} \text{ ("good" channel)} \\ \dfrac{c\, r^{\mathrm{b}}}{q^{*}_{\mathrm{b,g}}(r^{\mathrm{g}} - r^{\mathrm{b}})}, & r = r^{\mathrm{b}} \text{ ("bad" channel)} \end{cases}$$

  Secondary (tie-breaking) Whittle index

$$\tilde{W}(r^{\mathrm{g}}) = c\mu\, r^{\mathrm{g}}$$

  – Jacko (*PEVA* 2011)

Outline

- Introduction
- Whittle index approach
- Opportunistic scheduling problem
- Our contribution
- Numerical illustrations
- Summary

Size-aware opportunistic scheduling problem

- Problem:
  - Assume there are K users with known remaining file sizes x_i
  - Channel states $R_i(t)$ are two-state continuous-time Markov processes
  - Holding costs are accrued with rate c_i for any uncompleted flow i
  - Given the remaining sizes x_i and the channel states r_i, choose one flow until all flows completed
- Our approach:
  - Approximate the known size by a phase-type distribution (Erlang distribution)
  - Asymptotically exact when the number of phases increased without limits

Phase-type approximation

- Definition: Erlang distribution with J phases and intensities $J\mu$

$$X = X_1 + \ldots + X_J \qquad (X_j \text{ IID})$$

$$P\{X_j > t\} = e^{-J\mu t}$$

$$E[X] = \frac{1}{\mu}, \qquad \mathrm{Var}[X] = \frac{1}{J\mu^2}$$

- Deterministic size x approximated by a random variable X with such an Erlang distribution ($\mu = 1/x$)

$$E[X] = x, \qquad \mathrm{Var}[X] = \frac{x^2}{J} \to 0 \quad (J \to \infty)$$

Approximate opportunistic scheduling problem

- Problem:
  - Assume there are K users with Erlang(J, $J\mu_i$) file sizes X_i
  - Channel states $R_i(t)$ are two-state continuous-time Markov processes
  - Holding costs are accrued with rate c_i for any uncompleted flow i
  - Given the remaining number of phases j_i and the channel states r_i, choose one flow until all flows completed
- Heuristic answer:
  - Consider the separable Lagrangian version of the relaxed problem
  - Show indexability separately for each flow i
  - Calculate the Whittle index $W_i(j_i,r_i)$ separately for each flow i
  - Choose the flow i^* with the highest Whittle index

Relaxed opportunistic scheduling problem

- Separable Lagrangian version of the relaxed problem:

$$f_i^{\pi_i} + \nu\, g_i^{\pi_i} = \min_{\pi_i}! \qquad (*)$$

where

$$f_i^{\pi_i} \triangleq E\left[\int_0^{\infty} c_i 1_{\{Z_i^{\pi_i}(t)>0\}}\, dt\right], \qquad g_i^{\pi_i} \triangleq E\left[\int_0^{\infty} A_i^{\pi_i}(t)\, dt\right]$$

- Definition:
  Optimization problem (*) is indexable if
  for any *j* and *r* there is $W_i(j,r)$ such that
  – it is optimal to schedule flow *i* in state (j,r) if $\nu \le W_i(j,r)$
  – it is optimal *not* to schedule flow *i* in state (j,r) if $\nu \ge W_i(j,r)$

Whittle index for Erlang file sizes (1)

- Result:
  Primary Whittle index
  for a flow with j remaining phases and channel state r is given by

$$W(j,r)=\begin{cases}\infty, & r=r^{\mathrm{g}}\ (\text{"good" channel})\\ \dfrac{\dfrac{c\,r^{\mathrm{b}}}{q^{*}_{\mathrm{b,g}}(r^{\mathrm{g}}-r^{\mathrm{b}})}-cd(j-1)}{1+d(j)}, & r=r^{\mathrm{b}}\ (\text{"bad" channel})\end{cases}$$

  Secondary (tie-breaking) Whittle index

$$\tilde{W}(j,r^{\mathrm{g}})=\frac{Jc\mu\,r^{\mathrm{g}}}{j}$$

Whittle index for Erlang file sizes (2)

- Result:
  Primary Whittle index
  for a flow with J phases satisfies

$$\lim_{J\to\infty} W(J, r^{\mathrm{b}}) = \lim_{J\to\infty} \frac{\dfrac{c\, r^{\mathrm{b}}}{q^{*}_{\mathrm{b,g}}(r^{\mathrm{g}} - r^{\mathrm{b}})} - cd(J-1)}{1 + d(J)} = \frac{c\, r^{\mathrm{b}}}{P\{R = r^{\mathrm{g}}\}(r^{\mathrm{g}} - r^{\mathrm{b}})}$$

  Secondary Whittle index satisfies

$$\lim_{J\to\infty} \tilde{W}(J, r^{\mathrm{g}}) = \lim_{J\to\infty} c\mu\, r^{\mathrm{g}} = c\mu\, r^{\mathrm{g}} = \frac{c\, r^{\mathrm{g}}}{x}$$

Approximate size-aware Whittle index

- Result:
  Primary approximate Whittle index
  for a flow with remaining size x and channel state r is given by

$$W(x,r) = \begin{cases} \infty, & r = r^{\mathrm{g}} \text{ ("good" channel)} \\ \dfrac{c\, r^{\mathrm{b}}}{P\{R = r^{\mathrm{g}}\}(r^{\mathrm{g}} - r^{\mathrm{b}})}, & r = r^{\mathrm{b}} \text{ ("bad" channel)} \end{cases}$$

  Secondary approximate Whittle index

$$\tilde{W}(x, r^{\mathrm{g}}) = \frac{c\, r^{\mathrm{g}}}{x}$$

Outline

- Introduction
- Whittle index approach
- Opportunistic scheduling problem
- Our contribution
- Numerical illustrations
- Summary

Performance in Scenario 1: Homogeneous users

- 1 class
- Poisson flow arrivals
- Pareto file sizes
- 2 channel states

- **PF** = Proportional Fair scheduler
- **PI*** = Potential Improv.* Jacko (2011)
- **ASPI** = Attained Service dependent PI Taboada et al. (2014)
- **SW** = Size-aware Whittle index policy

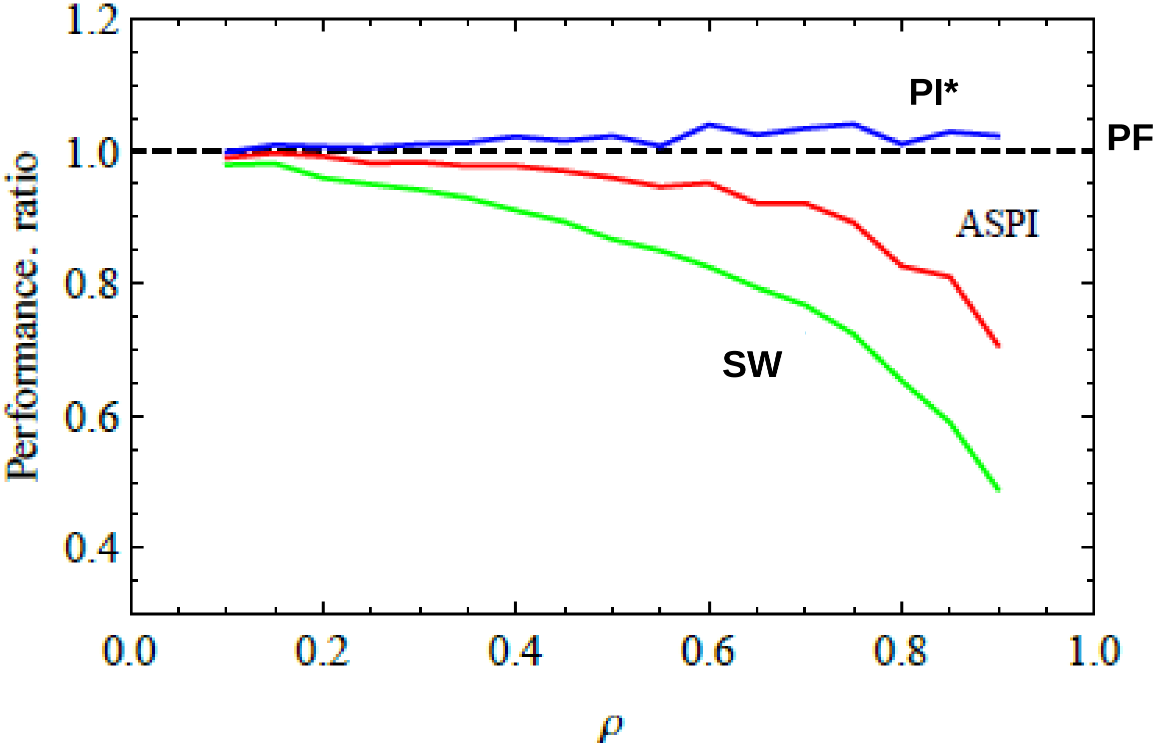

Performance in Scenario 2: Heterogeneous users

- 2 classes
- Poisson flow arrivals
- Pareto file sizes
- 2 channel states

- **PF** = Proportional Fair scheduler
- **PI*** = Potential Improv.* Jacko (2011)
- **ASPI** = Attained Service dependent PI Taboada et al. (2014)
- **SW** = Size-aware Whittle index policy

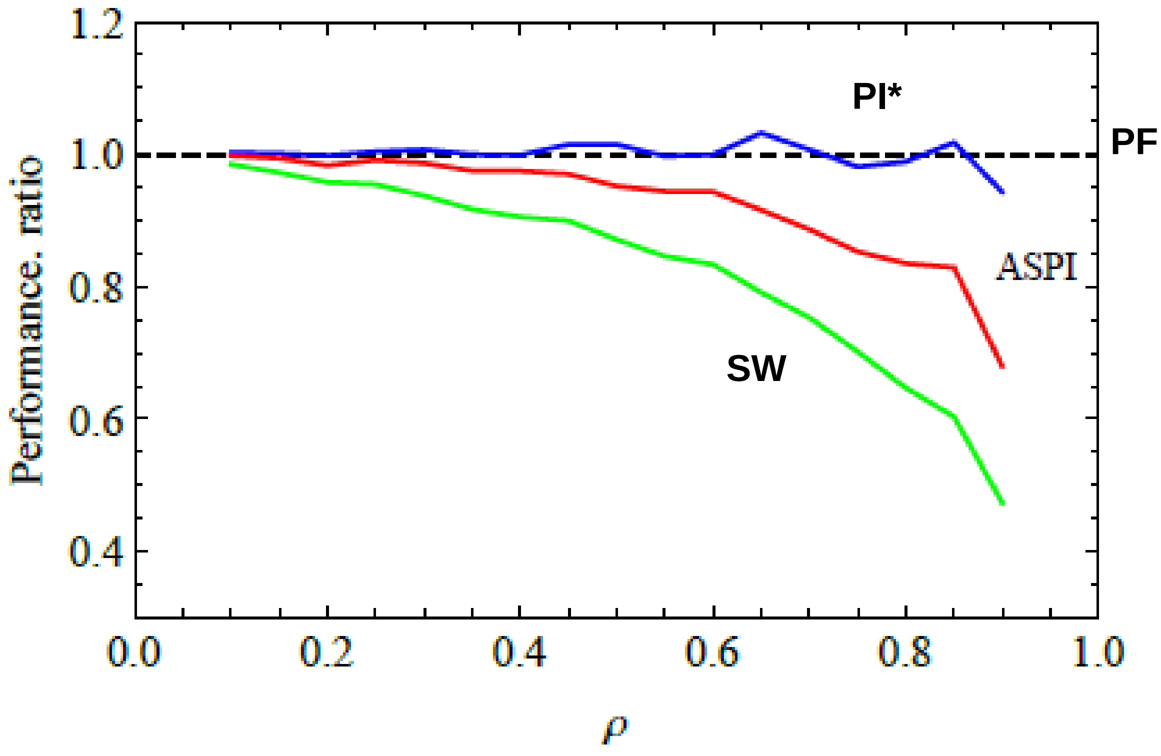

Fairness in Scenario 2: Heterogeneous users

- 2 classes
- Poisson flow arrivals
- Pareto file sizes
- 2 channel states

- **PF** = Proportional Fair scheduler
- **PI*** = Potential Improv.* Jacko (2011)
- **ASPI** = Attained Service dependent PI Taboada et al. (2014)
- **SW** = Size-aware Whittle index policy

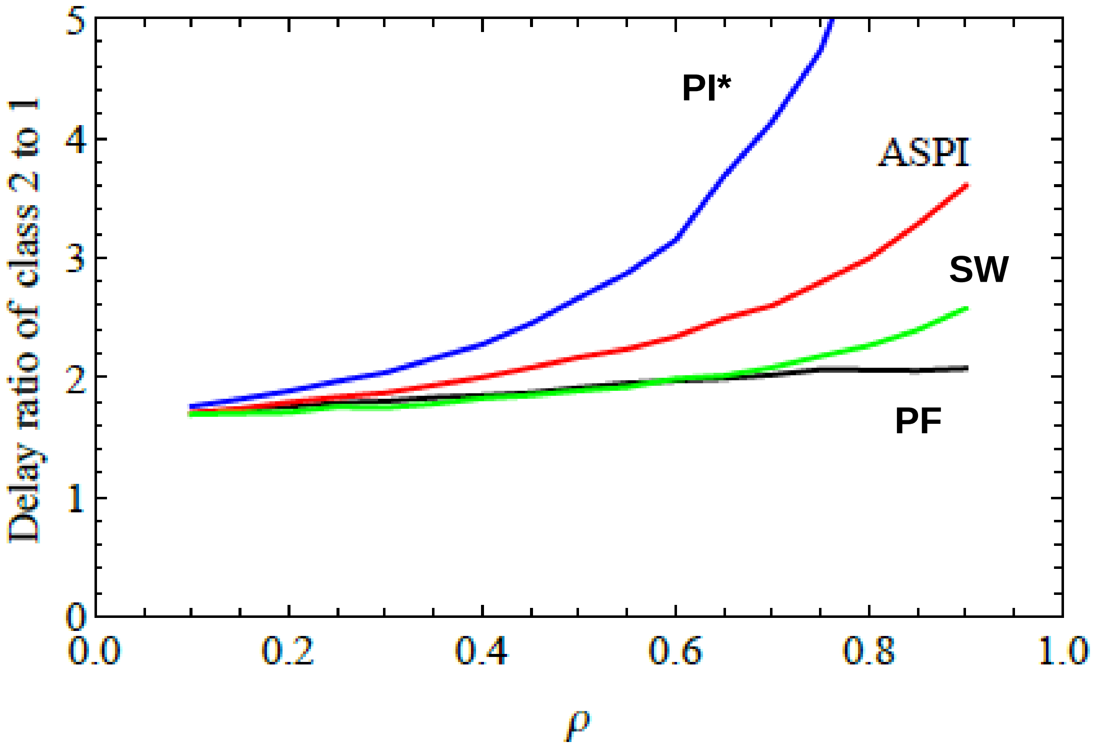

Outline

- Introduction
- Whittle index approach
- Opportunistic scheduling problem
- Our contribution
- Numerical illustrations
- Summary

Summary

- We considered the size-aware opportunistic scheduling problem for elastic downlink data traffic with heterogeneous two-state Markovian time-varying channels
- By the Whittle index approach and a phase-type approximation, we were able to derive an approximative size-aware Whittle index
- Primary index:
  - infinite for the good channel state
  - independent of the remaining size for the bad channel state
- Secondary index:
  - inversely proportional to the remaining size for the good channel state

The End