Aalto University
School of Electrical Engineering

Optimal Trade-off between Size-based and Opportunistic Scheduling: Whittle Index Approach

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

INFORMS APS
5-8 July 2015
Istanbul, Turkey

Outline

- Introduction
- Whittle index approach
- Our contribution
- Numerical illustrations
- Summary

Research problem

- Downlink data transmission in a cellular system
  - traffic = elastic flows
  - file transfers using TCP
  - file sizes known
- Traffic dynamics
  - time scale of seconds+
- Time-varying channels of users
  - time scale of milliseconds
  - channel states known
- Scheduling decisions
  - time scale of milliseconds
- Optimal scheduler for flow-level performance?
  - time scale of seconds+

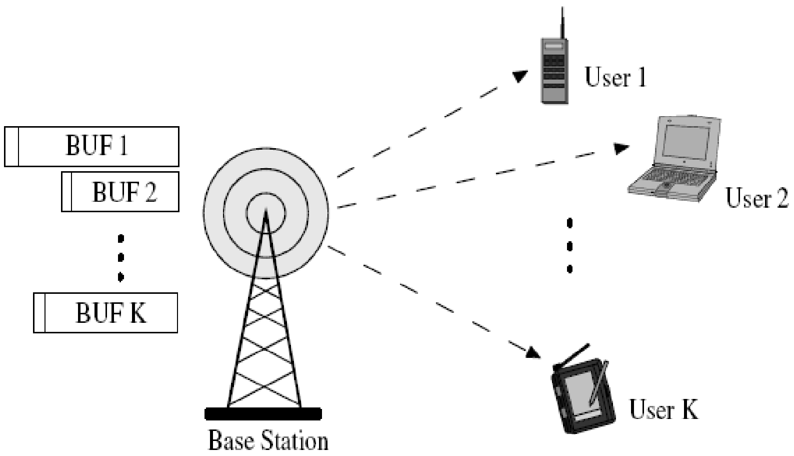

Source: Hu et al. (*ComNet* 2004)

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. (*QUESTA* 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)

Outline

- Introduction
- Whittle index approach
- 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 process i is chosen when in state x_i, a reward of $r_i(x_i)$ is earned
  - Given the states x_i of the bandit processes, choose the optimal bandit i^*
- Answer:
  - Calculate the Gittins index $G_i(x_i)$ separately for each process i
  - Choose the bandit i^* with the highest Gittins index
  - Gittins and Jones (1974), Gittins (1989)
- 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)

Gittins index

- Definition:
  Gittins index for a job with attained service a is given by

$$G(a) \triangleq c \cdot \sup_{\Delta \geq 0} J(a, \Delta)$$

where

$$J(a, \Delta) \triangleq \frac{P\{S - a \leq \Delta \mid S > a\}}{E[\min\{S - a, \Delta\} \mid S > a]} = \frac{\int_a^{a+\Delta} f(y)dy}{\int_a^{a+\Delta} (1 - F(y))dy}$$

- Note: For deterministic service time s,

$$G(a) = \frac{1}{s - a}$$

  – Optimality of SRPT!

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 process 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 of the bandit processes, choose the optimal bandit i^*
- Relaxed problem:
  - Given the states x_i of the bandit processes, choose the optimal bandits so that at most one process is chosen per time slot 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 process i
  - Calculate the Whittle index $W_i(x_i)$ separately for each process 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

Opportunistic scheduling problem

- Problem:
  - Assume there are K jobs with geometric sizes X_i (prob. μ_i)
  - Channel states $R_i(t)$ are independent two-state IID variables (good/bad)
  - If job i with channel state r_i is chosen, it completes with prob. $\mu_i \cdot r_i$
  - Holding costs are accrued with rate c_i for any uncompleted job i
  - Given the channel states r_i of the jobs, choose the optimal job i^*
- Heuristic answer:
  - Show indexability separately for each process i
  - Calculate the Whittle index $W_i(r_i)$ separately for each process i
  - Choose the job i^* with the highest Whittle index
  - Ayesta et al. (*PEVA* 2010)
- Generalizations:
  - Jacko (*PEVA* 2011), Cecchi and Jacko (*Sigmetrics* 2013)
  - Taboada et al. (*ITC* 2014), Taboada et al. (*PEVA* 2014)

Whittle index for geometric job sizes

- Result:
  Primary Whittle index for a job 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 Whittle index:

$$\tilde{W}(r) = \begin{cases} c\mu\, r^{\mathrm{g}}, & r = r^{\mathrm{g}} \text{ ("good" channel)} \\ 0, & r = r^{\mathrm{b}} \text{ ("bad" channel)} \end{cases}$$

  - Ayesta et al. (*PEVA*, 2010)

Outline

- Introduction
- Whittle index approach
- Our contribution
- Numerical illustrations
- Summary

Size-aware opportunistic scheduling problem

- Problem:
  - Assume there are K jobs with known sizes x_i
  - Channel states $R_i(t)$ are independent two-state IID variables (good/bad)
  - If job i with channel state r_i is chosen, it completes if $x_i < r_i$
  - Holding costs are accrued with rate c_i for any uncompleted job i
  - Given the job sizes x_i and the channel states r_i, choose the optimal job i^*
- Our approach:
  - Approximate the known size with a discrete-time phase-type distribution (i.e., shifted Pascal distribution)

Phase-type approximation

- Definition: Shifted Pascal distribution with J phases and succ. prob. p

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

$$P\{X_j = n\} = (1-p)^{n-1} p, \quad n = 1,2,\ldots$$

$$E[X] = \frac{J}{p}, \qquad \mathrm{Var}[X] = \frac{J(1-p)}{p^2}$$

- Deterministic job size x approximated by a random variable X with shifted Pascal distribution (J phases, success prob. $p = J/x$)

$$E[X] = x, \qquad C[X] = \sqrt{\frac{1}{J} - \frac{1}{x}}$$

  – For large x and J, the relative variance is small!

Approx. opportunistic scheduling problem

- Problem:
  - Assume there are K jobs with shifted Pascal sizes X_i (J, p_i)
  - Channel states $R_i(t)$ are independent two-state IID variables (good/bad)
  - If job i with channel state r_i is chosen, the job completes its phase with probability $p_i \cdot r_i$
  - Holding costs are accrued with rate c_i for any uncompleted job i
  - Given the phases j_i and the channel states r_i of the jobs, choose the optimal job i^*
- Heuristic answer:
  - Consider the separable Lagrangian version of the relaxed problem
  - Show indexability separately for each process i
  - Calculate the Whittle index $W_i(j_i,r_i)$ separately for each process i
  - Choose the job 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

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

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

Whittle index for shifted Pascal job sizes

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

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

Secondary Whittle index:

$$\tilde{W}(j,r)=\begin{cases}\dfrac{c\,p\,r^{\mathrm{g}}}{j}, & r=r^{\mathrm{g}}\ (\text{"good" channel})\\ 0, & r=r^{\mathrm{b}}\ (\text{"bad" channel})\end{cases}$$

Approximative size-aware Whittle index

- Result:
  Primary approximative Whittle index
  for a job with remaining size y and channel state r is given by

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

  Secondary approximative Whittle index:

$$\tilde{W}(y,r) = \begin{cases} \dfrac{c\, r^{\mathrm{g}}}{y}, & r = r^{\mathrm{g}} \text{ ("good" channel)} \\ 0, & r = r^{\mathrm{b}} \text{ ("bad" channel)} \end{cases}$$

Outline

- Introduction
- Whittle index approach
- Our contribution
- Numerical illustrations
- Summary

Scenario 1: Homogeneous users

- 1 class
- Poisson job arrivals
- Pareto job sizes
- 2 channel states

- **PF** = Proportional Fair scheduler
- **PI** = Potential Improv. Ayesta et al. (2010)
- **ASPI** = Attained Service dependent PI Taboada et al. (2014)
- **SW** = Size-aware Whittle index policy

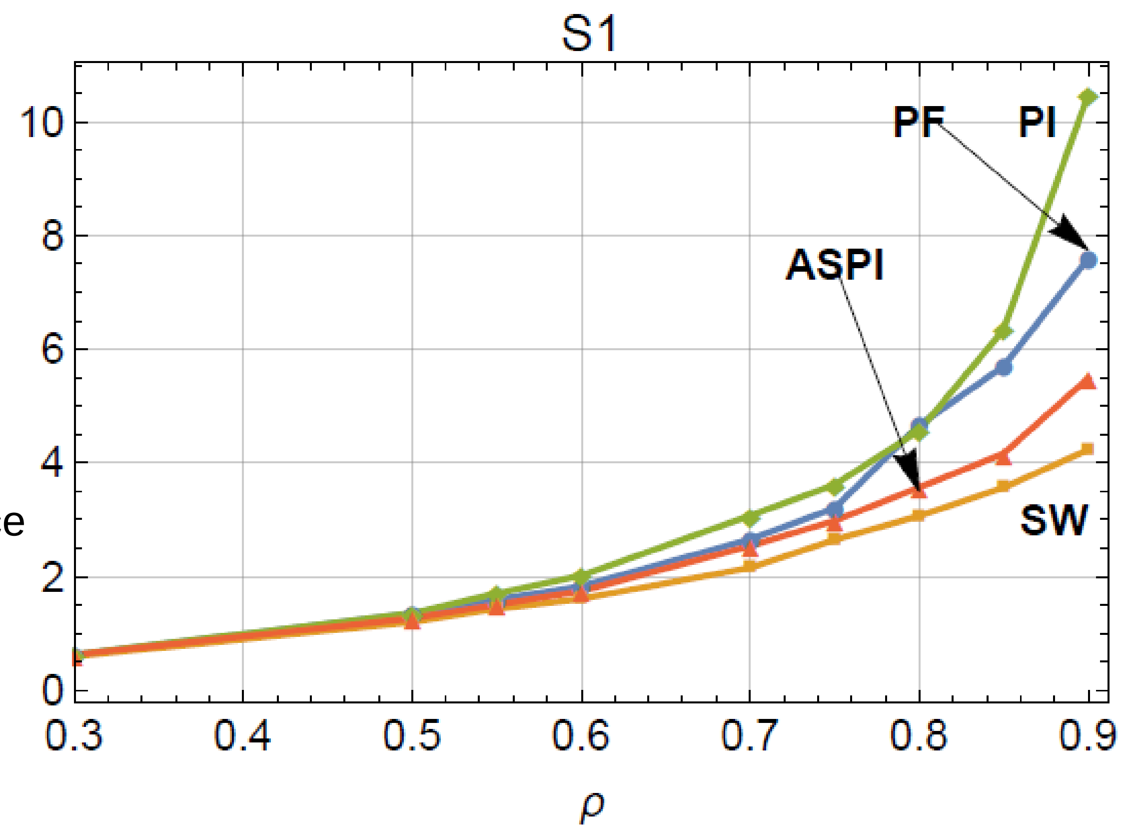

Scenario 1: Homogeneous users

- 1 class
- Poisson job arrivals
- Pareto job sizes
- 2 channel states

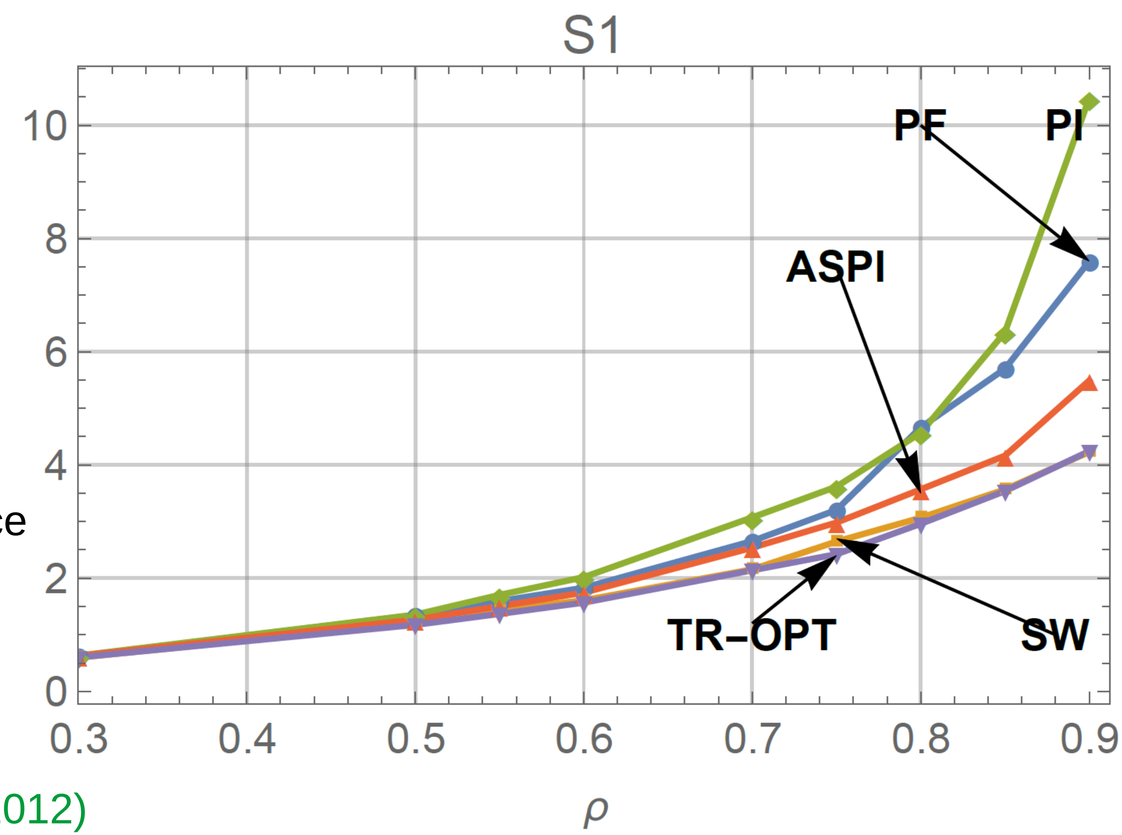

- **PF** = Proportional Fair scheduler
- **PI** = Potential Improv. Ayesta et al. (2010)
- **ASPI** = Attained Service dependent PI Taboada et al. (2014)
- **SW** = Size-aware Whittle index policy
- **TR-OPT**: Aalto et al. (2012)

Scenario 2c: Heterogeneous users

- 2 classes with different channels
- Poisson job arrivals
- Exp. job sizes
- 2 channel states

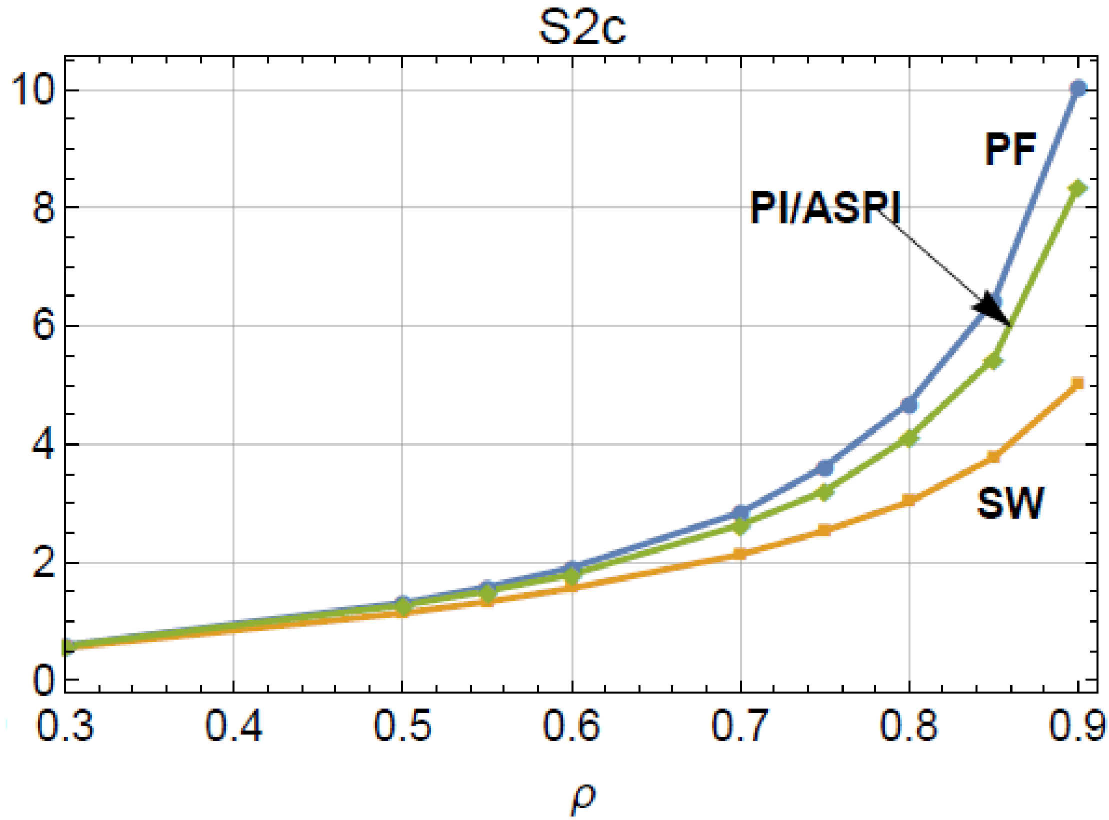

Scenario 3: Multiple channel states

- 2 classes with different channels
- Poisson job arrivals
- Exp. job sizes
- 5/3 channel states

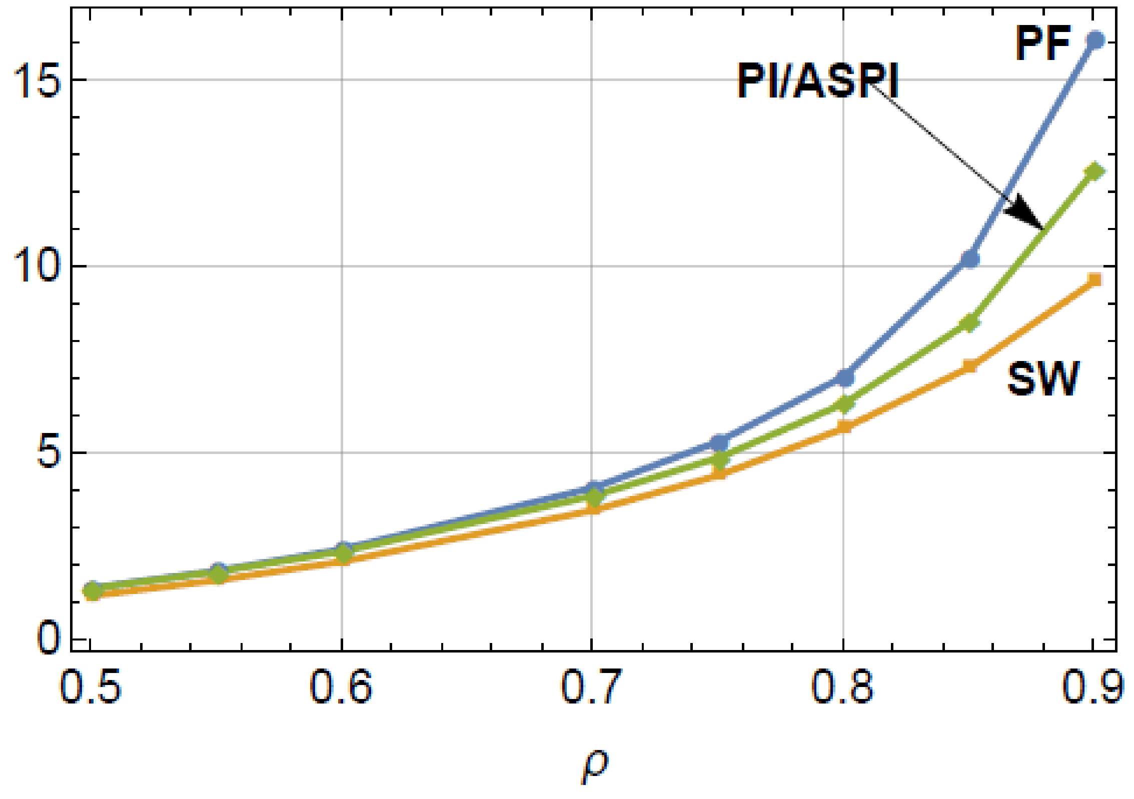

Outline

- Introduction
- Whittle index approach
- Our contribution
- Numerical illustrations
- Summary

Summary

- We considered the size-aware opportunistic scheduling problem for elastic downlink data traffic with two-state time-varying channels
- By the Whittle index approach and a discrete-time 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 job size for the bad channel state
- Secondary index:
  - inversely proportional to the remaining size for the good channel state
  - zero for the bad channel state

References

- Gittins and Jones (1974)
  A dynamic allocation index for the sequential design of experiments,
  in *Progress in Statistics*, North-Holland
- Gittins (1989, 2nd ed. 2011)
  Multi-armed Bandit Allocation Indices, Wiley
- Sadiq and de Veciana (2010)
  Balancing SRPT prioritization vs opportunistic gain in wireless systems with flow dynamics,
  in *Proc. of ITC 22*
- Aalto, Penttinen, Lassila and Osti (2011)
  On the optimal trade-off between SRPT and opportunistic scheduling,
  in *Proc. of ACM Sigmetrics*
- Aalto, Penttinen, Lassila and Osti (2012)
  Optimal size-based opportunistic scheduler for wireless systems,
  Queueing Systems
- Aalto, Lassila and Osti (2015)
  Whittle index approach to size-aware scheduling with time-varying channels,
  in *Proc. of ACM Sigmetrics*
- Whittle (1988)
  Restless bandits: activity allocation in a changing world,
  Journal of Applied Probability
- Ayesta, Erasquin and Jacko (2010)
  A modeling framework for optimizing the flow-level scheduling with time-varying channels,
  Performance Evaluation
- Jacko (2011)
  Value of information in optimal flow-level scheduling of users with Markovian time-varying channels,
  Performance Evaluation
- Cecchi and Jacko (2013)
  Scheduling of users with Markovian time-varying transmission rates,
  in *Proc. of ACM Sigmetrics*
- Taboada, Jacko, Ayesta and Liberal (2014)
  Opportunistic scheduling of flows with general size distribution in wireless time-varying channels,
  in *Proc. of ITC 26*
- Taboada, Liberal and Jacko (2014)
  An opportunistic and non-anticipating size-aware scheduling proposal for mean holding cost minimization in time-varying channels,
  Performance Evaluation

The End