Aalto University
School of Electrical
Engineering

Performance-Energy Trade-off in Queueing Systems with Setup Delay

Samuli Aalto
Aalto University, Finland

Sigmetrics TPC Workshop
12 February 2016
New York, USA

Co-operation with

Misikir Eyob Gebrehiwot, Pasi Lassila
(Aalto University, Finland)

Model

- In this talk, we focus on a single energy-aware server
- M/G/1 queue:
  Jobs arrive according to a Poisson process with rate λ and service times S are generally distributed and i.i.d.

$$\rho = \lambda \, \mathrm{E}[S] < 1$$

- Setup delays D are generally distributed and i.i.d.
- Four energy states with power consumption denoted by P_{state}

BUSY / IDLE / SLEEP / SETUP

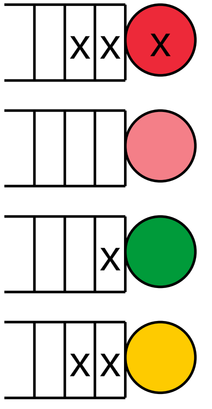

- Assumption:

$$0 \leq P_{\text{sleep}} < P_{\text{idle}} < P_{\text{setup}} \leq P_{\text{busy}}$$

Performance-energy trade-off

- Energy saved by switching the server off when idle
- However, performance impaired if switching the server back on takes some time (setup delay)

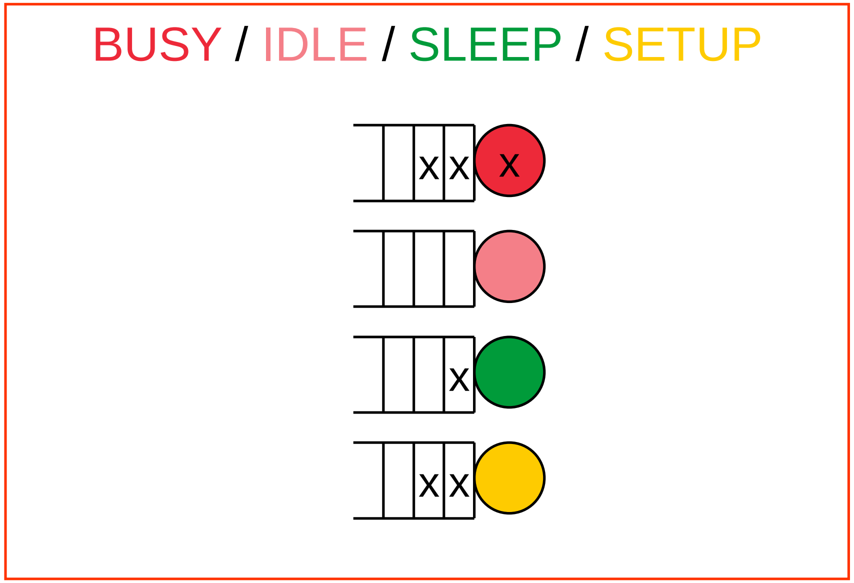

Optimal control problem

- Two possible control actions:
  - When IDLE, how long an interval *I* the server should wait for new jobs before switched off?
  - When SLEEP, how many new jobs *k* the server should wait for before switched on?

BUSY / IDLE / SLEEP / SETUP

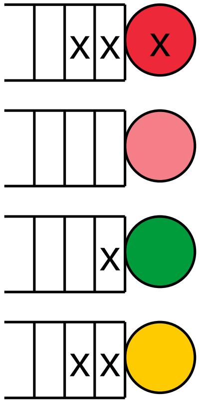

- Objective:
  Choose optimal timer *I* and threshold *k*

- Definition:
  - INSTANTOFF ($I = 0$)
  - NEVEROFF ($I = \infty$)
  - DELAYEDOFF ($0 < I < \infty$)

Performance and energy measures

- Performance:

  $E[T]$ = mean response time per job

  $E[X] = \lambda \cdot E[T]$ = mean number of jobs

- Energy:

  $E[E]$ = mean energy per job

  $E[P] = \lambda \cdot E[E]$ = mean power consumption

Objective functions

- ERWS (Energy-Response-time-Weighted-Sum):

$$w_1 \, \mathrm{E}[T] + w_2 \, \mathrm{E}[P]$$

e.g. Wierman & al. (2009)

- General form:

$$w_1 \, \mathrm{E}[T]^{\alpha_1} \, \mathrm{E}[P]^{\beta_1} + w_2 \, \mathrm{E}[T]^{\alpha_2} \, \mathrm{E}[P]^{\beta_2}$$

Maccio & Down (2013)

- ERP (Energy-Response-time-Product):

$$\mathrm{E}[T] \, \mathrm{E}[P]$$

e.g. Gandhi & al. (2010)

- Generalized ERP:

$$w \, \mathrm{E}[T]^{\alpha} \, \mathrm{E}[P]^{\beta}, \quad \alpha, \beta \geq 0$$

Gebrehiwot & al. (2014)

Energy-aware M/G/1-FIFO queue

Optimal *I*-control for the FIFO discipline

- Theorem 1:
  Maccio & Down (2013)
  For ERWS
  with exponential
  service times and setup delays,
  the optimal policy is either
  NEVEROFF or INSTANTOFF

- Theorem 2:
  Gebrehiwot et al. (2014)
  For ERWS and generaliz. ERP
  with generally distributed
  service times and setup delays,
  the optimal policy is either
  NEVEROFF or INSTANTOFF

- Observation:
  NEVEROFF is better
  if P_{idle} is sufficiently small
  compared to P_{setup}

- Counter-example:
  Gebrehiwot et al. (2014)
  For the general form objective
  the result is not necessarily true

Proof of Theorem 2

Theorem 2:
Gebrehiwot et al. (2014)
For ERWS and generalized ERP with generally distributed service times and setup delays, the optimal policy is either NEVEROFF or INSTANTOFF

- Analysis (1)
  - Consider working cycles starting whenever the server is switched off
  - Let C denote the length of one cycle

$$\mathrm{E}[C] = \frac{\mathrm{E}[I] + \frac{k}{\lambda} + \mathrm{E}[D]}{1-\rho}$$

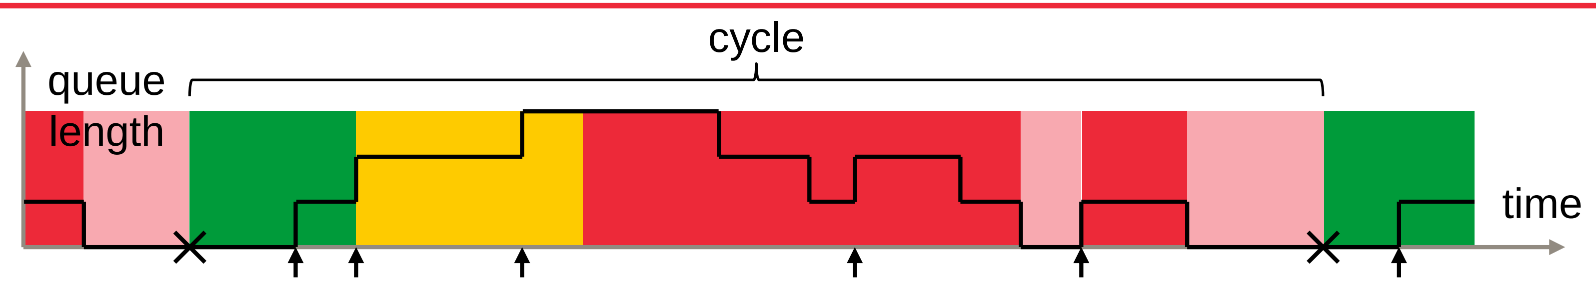

Proof of Theorem 2

Theorem 2:
Gebrehiwot et al. (2014)
For ERWS and generalized ERP with generally distributed service times and setup delays, the optimal policy is either NEVEROFF or INSTANTOFF

- Analysis (2)

$$\mathrm{E}[T] = \mathrm{E}[T_{\text{M/G/1-FIFO}}] + \frac{1}{\mathrm{E}[I] + \frac{k}{\lambda} + \mathrm{E}[D]} \left(\frac{k(k-1)}{2\lambda^2} + \frac{k}{\lambda}\mathrm{E}[D] + \frac{1}{2}\mathrm{E}[D^2] \right)$$

$$\mathrm{E}[P] = \mathrm{E}[P_{\text{M/G/1}}] + \frac{1-\rho}{\mathrm{E}[I] + \frac{k}{\lambda} + \mathrm{E}[D]} \left(\frac{k}{\lambda}(P_{\text{sleep}} - P_{\text{idle}}) + \mathrm{E}[D](P_{\text{setup}} - P_{\text{idle}}) \right)$$

– distribution of I is irrelevant when the mean value $\mathrm{E}[I]$ is fixed

Proof of Theorem 2

Theorem 2:
Gebrehiwot et al. (2014)
For ERWS and generalized ERP with generally distributed service times and setup delays, the optimal policy is either NEVEROFF or INSTANTOFF

- Optimization (1)

$$\mathrm{E}[T] = a_1 + \frac{b_1}{\mathrm{E}[I]+c}$$

$$\mathrm{E}[P] = a_2 + \frac{b_2}{\mathrm{E}[I]+c}$$

  - where $a_1, a_2, b_1, c > 0$, but b_2 may be negative since $P_{sleep} < P_{idle}$

$$b_2 = (1-\rho)\left(\frac{k}{\lambda}(P_{\text{sleep}} - P_{\text{idle}}) + \mathrm{E}[D](P_{\text{setup}} - P_{\text{idle}})\right)$$

  - NEVEROFF is always optimal if $b_2 > 0$

Proof of Theorem 2

Theorem 2:
Gebrehiwot et al. (2014)
For ERWS and generalized ERP with generally distributed service times and setup delays, the optimal policy is either NEVEROFF or INSTANTOFF

- Optimization (2)
  - objective function for ERWS is clearly monotonic w.r.t. E[*I*]

$$w_1\,\mathrm{E}[T] + w_2\,\mathrm{E}[P] = w_1 a_1 + w_2 a_2 + \frac{w_1 b_1 + w_2 b_2}{\mathrm{E}[I]+c}$$

  - objective function for generalized ERP may be nonmonotonic w.r.t. E[*I*] but it does not have any local minima in (0, ∞)

$$w\,\mathrm{E}[T]^{\alpha}\,\mathrm{E}[P]^{\beta} = w\left(a_1 + \frac{b_1}{\mathrm{E}[I]+c}\right)^{\alpha}\left(a_2 + \frac{b_2}{\mathrm{E}[I]+c}\right)^{\beta}$$

  - Q.E.D.

Generalizations of Theorem 2

Theorem 2:
Gebrehiwot et al. (2014)
For ERWS and generalized ERP with generally distributed service times and setup delays, the optimal policy is either NEVEROFF or INSTANTOFF

- Multiple sleep states
  - Randomized policies (”choose sleep state randomly”)
    Gandhi et al. (2010) for the exponential case
    Gebrehiwot et al. (2014) for general distributions
  - Sequential policies (”sleep deeper and deeper”)
    Gebrehiwot et al. (2015) for general distributions
- Idling timer resetting options
  - Timer I reset only when it expires
    Maccio & Down (2013) for the exponential case
    Gebrehiwot et al. (2014) for general distributions
  - Timer I reset every time an idle period starts
    Gebrehiwot et al. (2014) for general distributions

Counter-example

Counter-example:
Gebrehiwot et al. (2014)
For the general form objective
the result is not necessarily true

- Arrival rate
  $\lambda = 0.1$
- Exponential service times and deterministic setup delays with
  $E[S] = D = 1$
- Power consumptions
  $P_{busy} = P_{setup} = 1,$
  $P_{idle} = 0.6, P_{sleep} = 0$
- Following general form objective function considered:

$$E[T] + 3E[P]^3$$

- Optimal $E[I] = 20.723$

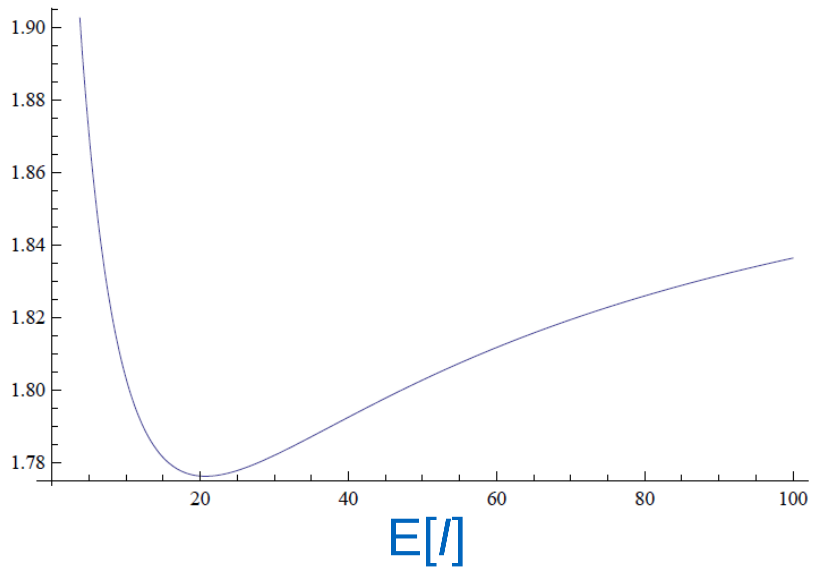

Optimal *k*-control for the FIFO discipline

- Theorem 3:
  Gandhi et al. (2010)
  For ERP
  with exponential service times and deterministic setup delays, the optimal threshold of INSTANTOFF is $k^* = 1$

- Counter-example:
  Gebrehiwot et al. (2014)
  For more variable service time distributions, the result is not necessarily true

- Theorem 4:
  Gebrehiwot et al. (2015)
  For ERP
  with generally distributed service times for which

$$C[S] = \sqrt{E[S^2]/E[S]^2} \leq \sqrt{3}$$

  and deterministic setup delays, the optimal threshold of INSTANTOFF is $k^* = 1$

Energy-aware M/G/1-PS queue

Optimal *I*-control for the PS discipline

- Theorem 5:
  Gebrehiwot et al. (2016)
  For ERWS and ERP
  with generally distributed
  service times and setup delays,
  the optimal policy is either
  NEVEROFF or INSTANTOFF

Proof of Theorem 5

Theorem 5:
Gebrehiwot et al. (2016)
For ERWS and ERP with generally distributed service times and setup delays, the optimal policy is either NEVEROFF or INSTANTOFF

- Analysis (1)
  - Consider again working cycles starting whenever the server is switched off
  - Let C denote the length of one cycle

$$E[C] = \frac{E[I] + \frac{k}{\lambda} + E[D]}{1-\rho}$$

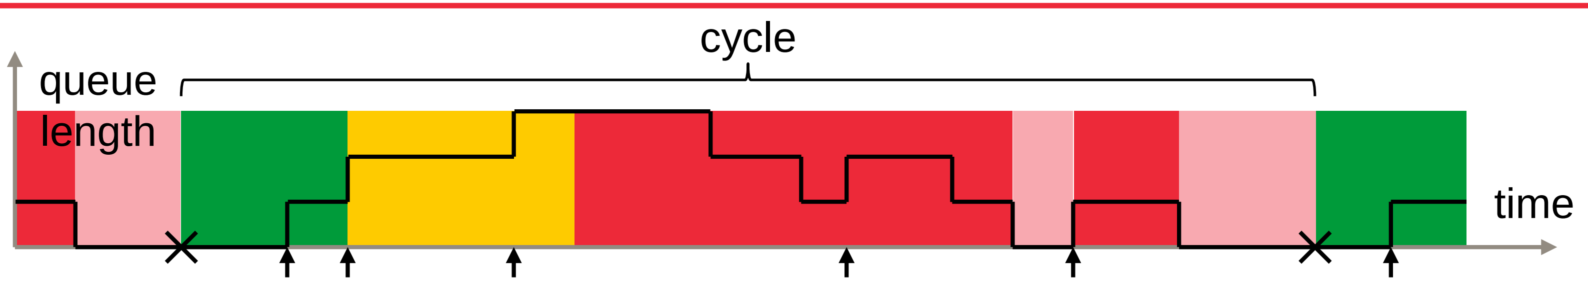

Aalto University
School of Electrical Engineering

Proof of Theorem 5

Theorem 5:
Gebrehiwot et al. (2016)
For ERWS and ERP with generally distributed service times and setup delays, the optimal policy is either NEVEROFF or INSTANTOFF

- Analysis (2)
  - Response time consists of waiting time (from arrival to the beginning of service) and residence time (the rest): $T = W + R$
  - Consider a test job with service time s

$$E[W] = \mathrm{E}[W \mid S = s] = \frac{1-\rho}{\mathrm{E}[I] + \frac{k}{\lambda} + \mathrm{E}[D]} \left(\frac{k(k-1)}{2\lambda^2} + \frac{k}{\lambda}\mathrm{E}[D] + \frac{1}{2}\mathrm{E}[D^2] \right)$$

  - Let $r(s) = \mathrm{E}[R \mid S = s]$. Then $r'(s)$ is the unique solution of

$$r'(s) = 1 + 2\lambda E[W]\overline{F}(s) + \lambda \int_0^{\infty} r'(t)\overline{F}(s+t)dt + \lambda \int_0^{s} r'(t)\overline{F}(s-t)dt$$

Theorem 5:
Gebrehiwot et al. (2016)
For ERWS and ERP with generally distributed service times and setup delays, the optimal policy is either NEVEROFF or INSTANTOFF

Proof of Theorem 5

- Analysis (3)
  - Let $g(s) = (r'(s) - 1/(1-\rho))/(2\lambda E[W])$. Then $g(s)$ is the uniq. sol. of

$$g(s) = \bar{F}(s) + \lambda \int_0^\infty g(t)\bar{F}(s+t)dt + \lambda \int_0^s g(t)\bar{F}(s-t)dt$$

where

$$\bar{F}(s) = P\{S > s\}$$

  - Observation: $g(s)$ depends on the arrival rate λ and the service time distribution $F(s)$ but not at all on $E[I]$

Proof of Theorem 5

Theorem 5:
Gebrehiwot et al. (2016)
For ERWS and ERP with generally distributed service times and setup delays, the optimal policy is either NEVEROFF or INSTANTOFF

- Analysis (4)

$$\mathrm{E}[T] = \mathrm{E}[T_{\mathrm{M/G/1\text{-}PS}}] + \frac{1-\rho}{\mathrm{E}[I] + \frac{k}{\lambda} + \mathrm{E}[D]} \left(\frac{k(k-1)}{2\lambda^2} + \frac{k}{\lambda}\mathrm{E}[D] + \frac{1}{2}\mathrm{E}[D^2] \right) \left(1 + 2\lambda \int_0^{\infty} g(y)\bar{F}(y)dy \right)$$

$$\mathrm{E}[P] = \mathrm{E}[P_{\mathrm{M/G/1}}] + \frac{1-\rho}{\mathrm{E}[I] + \frac{k}{\lambda} + \mathrm{E}[D]} \left(\frac{k}{\lambda}(P_{\mathrm{sleep}} - P_{\mathrm{idle}}) + \mathrm{E}[D](P_{\mathrm{setup}} - P_{\mathrm{idle}}) \right)$$

  - distribution of *I* is irrelevant when the mean value E[*I*] is fixed

Proof of Theorem 5

Theorem 5:
Gebrehiwot et al. (2016)
For ERWS and ERP with generally distributed service times and setup delays, the optimal policy is either NEVEROFF or INSTANTOFF

- Optimization (1)

$$E[T] = a_1 + \frac{b_1}{E[I]+c}$$

$$E[P] = a_2 + \frac{b_2}{E[I]+c}$$

  - where $a_1, a_2, b_1, c > 0$, but b_2 may be negative since $P_{sleep} < P_{idle}$

$$b_2 = (1-\rho)\left(\frac{k}{\lambda}(P_{sleep} - P_{idle}) + E[D](P_{setup} - P_{idle})\right)$$

  - NEVEROFF is always optimal if $b_2 > 0$

Theorem 5:
Gebrehiwot et al. (2016)
For ERWS and ERP with generally distributed service times and setup delays, the optimal policy is either NEVEROFF or INSTANTOFF

Proof of Theorem 5

- Optimization (2)
  - objective function for ERWS is clearly monotonic w.r.t. E[*I*]

$$w_1 \, \mathrm{E}[T] + w_2 \, \mathrm{E}[P] = w_1 a_1 + w_2 a_2 + \frac{w_1 b_1 + w_2 b_2}{\mathrm{E}[I] + c}$$

  - objective function for ERP may be nonmonotonic w.r.t. E[*I*] but it does not have any local minima in $(0, \infty)$

$$\mathrm{E}[T]\,\mathrm{E}[P] = \left(a_1 + \frac{b_1}{\mathrm{E}[I] + c}\right)\left(a_2 + \frac{b_2}{\mathrm{E}[I] + c}\right)$$

  - Q.E.D.

Generalizations of Theorem 5

Theorem 5:
Gebrehiwot et al. (2016)
For ERWS and ERP with generally distributed service times and setup delays, the optimal policy is either NEVEROFF or INSTANTOFF

- Batch arrivals
  - Batch Poisson arrivals
    with *k* referring to the number batches arrived
    Gebrehiwot et al. (2016)
- Multiple sleep states
  - Randomized policies (”choose sleep state randomly”)
    Gebrehiwot et al. (2016)
- Idling timer resetting options
  - Timer *I* reset only when it expires
    Gebrehiwot et al. (2016)
  - Timer *I* reset every time an idle period starts
    Gebrehiwot et al. (2016)
- Generalized ERP

Numerical illustrations

- INSTANTOFF vs NEVEROFF
- Deterministic setup delays with $D = 10$ s
- Power consumptions $P_{busy} = P_{setup} = 200$ W, $P_{idle} = 120$ W, $P_{sleep} = 15$ W
- Different loads and service time distributions considered with fixed mean $E[S] = 1$ s
- Observation:
  Deterministic service times give always the worst results!

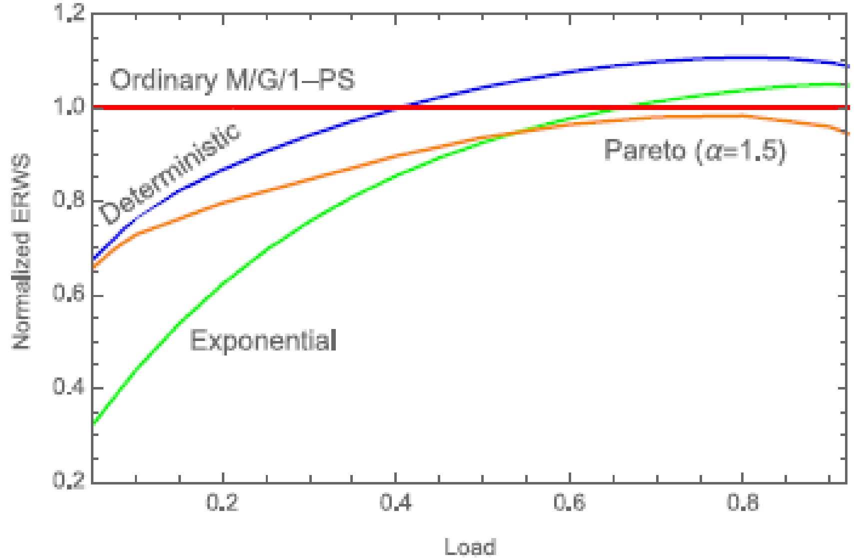

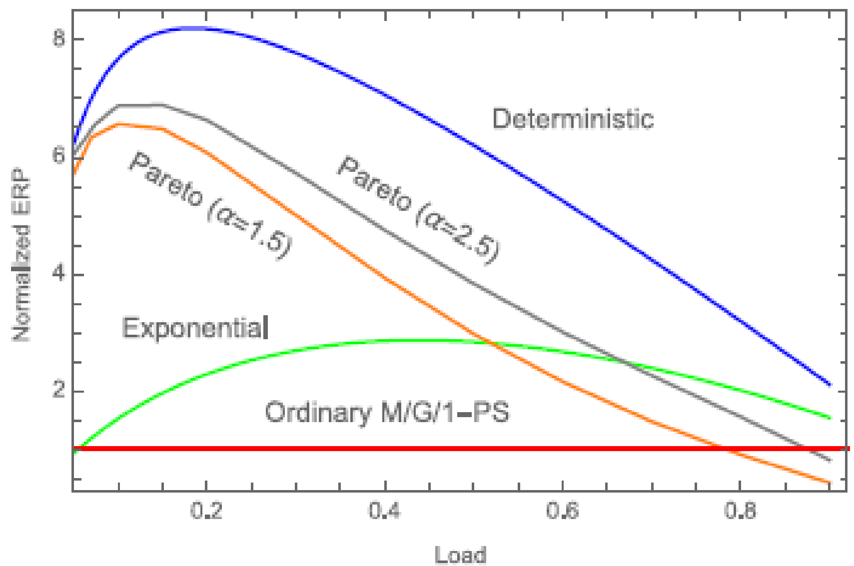

Worst case service time distribution for batch PS queues

Batch PS queueing models

- PS queue with multiple jobs in the system when the service starts
  - Energy-aware PS queue with setup delay
  - PS queue with multiple vacations
  - Batch arrival PS queue
  - Two-level MLPS queue (2nd level)
- Note:
  - In the ordinary PS queue there is only one job in the system when the service starts (after an idle period)

Intuition

- Deterministic service times
  - $S_1 = 2$
  - $S_2 = 2$
  - $S_3 = 2$
  - Average response time: 6.0
  - All jobs leave at the same time

- Variable service times
  - $S_1 = 1$
  - $S_2 = 2$
  - $S_3 = 3$
  - Average response time: 4.7 < 6.0
  - Shortest job leaves first

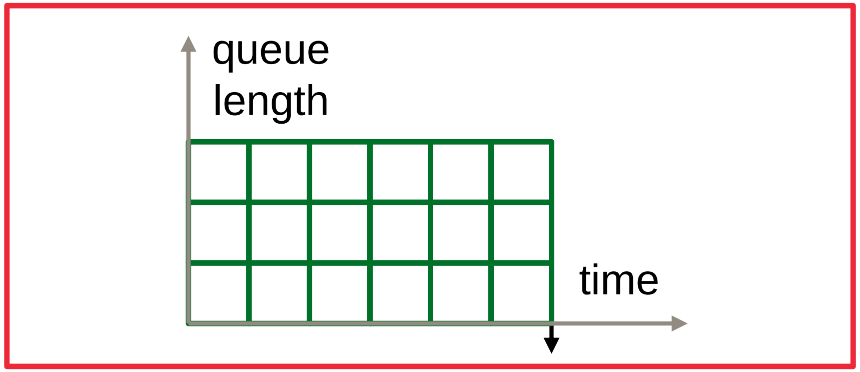

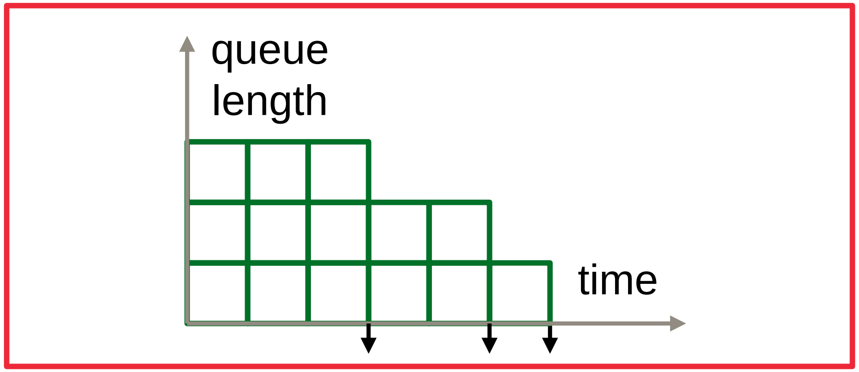

Worst case service time distribution

- Theorem 6:
  (Unpublished)
  For batch PS queues with (batch) Poisson arrivals, deterministic service times give the worst mean response time $E[T]$ among all service time distributions

- Conjecture 7:
  (Unpublished)
  For batch PS queues with (batch) Poisson arrivals, deterministic service times give even the worst conditional mean response time $E[T \mid S = s]$ among all service time distributions

- Proof: Utilizes an upper bound in Avrachenkov et al. (2005)

- Proposal: Anybody interested in cooperation to solve the remaining problems in its proof?

References

- Avrachenkov, Ayesta & Brown (2005)
  Batch arrival processor-sharing with application to multi-level processor-sharing scheduling, *Queueing Systems*
- Wierman, Andrew & Tang (2009)
  Power-aware speed scaling in processor sharing systems,
  in *IEEE INFOCOM*
- Gandhi, Gupta, Harchol-Balter & Kozuch (2010)
  Optimality analysis of energy-performance trade-off for server farm management,
  Performance Evaluation
- Maccio & Down (2013)
  On optimal policies for energy-aware servers,
  in *IEEE MASCOTS*
- Gebrehiwot, Aalto & Lassila (2014)
  Optimal sleep-state control of energy-aware M/G/1 queues,
  in *ValueTools*
- Gebrehiwot, Aalto & Lassila (2015)
  Optimal energy-aware control policies for FIFO servers,
  submitted
- Gebrehiwot, Aalto & Lassila (2016)
  Energy-performance trade-off for processor sharing queues with setup delay,
  Operations Research Letters

The End