Discriminatory processor sharing revisited

Konstantin Avrachenkov (INRIA)
Urtzi Ayesta (CWI)
Patrick Brown (France Telecom R&D)
R. Núñez Queija (CWI and TU/e)

The DPS model

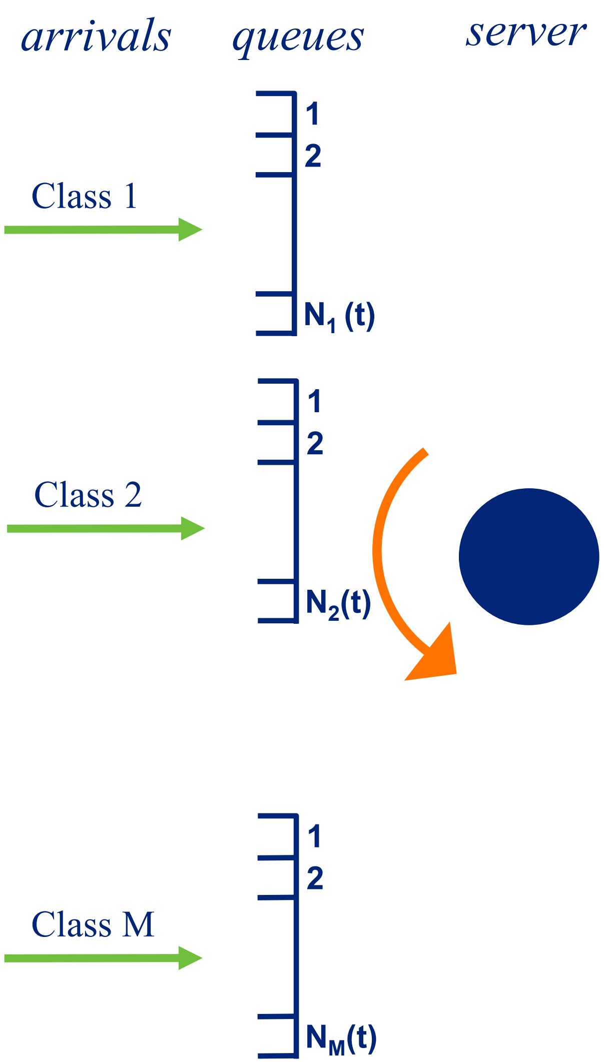

- Class $k \in \{1,2,\ldots,M\}$
  - Poisson arrival, λ_k
  - Service time requirements, $F_k(x) = P_k(X_k \leq x)$
  - Load $\rho_k = \lambda_k E[X_k]$
- stability, $\rho = \sum_{j=1}^{M} \rho_j < 1$
- vector of weights $\{g_j > 0; j=1,\ldots,M\}$
- service rate, $\dfrac{g_k}{\sum_{j=1}^{M} g_j N_j(t)}$

The DPS model: Applications

- Flow level models for the Internet:
  - Heterogeneous characteristics, round-trip delays ...
  - Service differentiation: application based, different ADSL users…

- Performance of weighted round robin: Resource sharing in a computer CPU.

- ...

Outline of the talk

- Background: Review of known results
- Expected response time: Finiteness
- Conservation law for DPS
- Expected conditional response time
  - Ordering according to weights
  - Closed form expression for the asymptotic bias
- Comparison between DPS and PS

Background

- References:
  - Kleinrock [1967]
  - O'Donovan [1974]
  - Fayolle, Mitrani and Iasnogorodski [1980]
  - Rege and Sengupta [1994,1996]
  - Borst, Van Ooteghem and Zwart [2003]
  - Altman, Jimenez and Kofman [2004]
  - Application to networks in Bonald and Massoulie [2001]; Bu and Towsley[2001]; Guo and Matta [2002];

Expected response time: Finiteness

➔ Let $T_k(t)$ be the expected conditional response time of a class-*k* job with required service time *t.*

Theorem: Let $E[X_k]<\infty$, $\forall k$ and $\rho<1$. Then if $g_k>0$,

$$E[T_k] = \int_0^\infty T_k(x)\,dF_k(x) < \infty$$

Conservation law for DPS

- Let $U_j(t)$ be the unfinished work of class-*j* jobs at time *t.*
- Let U(t) be the total unfinished work, $U(t) = \sum_{j=1}^{M} U_j(t)$
- A scheduling discipline is work-conserving if *U(t)* decreases at a rate 1 whenever *U(t)*>0.

  $\Rightarrow$ *U(t)* is independent of the discipline.

- Proposition: Assume $\mathrm{E}\left[X_j^2\right] < \infty$, $\forall j$ and $\rho<1$, then

$$\sum_{j=1}^{M} \lambda_j \int_0^{\infty} T_j(x) \overline{F}_j(x) dx = \mathrm{E}[U]$$

where $\mathrm{E}[U] = \dfrac{\sum_{j=1}^{M} \lambda_j \mathrm{E}\left[X_j^2\right]}{2(1-\rho)}$

Conditional response time: Ordering according to weights

Let $\tau_k(x)$ be the response time of a class-*k* job that requires *x* units of service.

Theorem: If $g_k > g_l$, then $\tau_k(x) \leq_{st} \tau_l(x)$, that is,

$$P(\tau_k(x) > y) \leq P(\tau_l(x) > y), \text{ for all } y \geq 0.$$

In particular

$$T_k(x) = \mathrm{E}[\tau_k(x)] \leq \mathrm{E}[\tau_l(x)] = T_l(x)$$

Conditional response time: Asymptotic bias

Theorem: Let $\mathrm{E}\left[X_k^2\right] < \infty \quad \forall k$, then,

$$\lim_{x \to \infty} \left(T_k(x) - \frac{x}{1-\rho} \right) = \frac{\sum_{j=1}^{M} \lambda_j \left(1 - g_k / g_j\right) \mathrm{E}\left[X_j^2\right]}{2(1-\rho)^2}$$

- Bias is independent of class' own second moment
- Bias of class with largest (smallest) weight is always negative (positive)

- Possible to choose the weights so that all but one class have negative/positive biases.

Conditional response time: Asymptotic bias (cont.)

- Example two classes $\mathrm{E}\left[X_1^2\right] >> \mathrm{E}\left[X_2^2\right]$

$$\lim_{x\to\infty}\left(T_1(x)-\frac{x}{1-\rho}\right)=\frac{\lambda_2(1-g_1/g_2)\mathrm{E}\left[X_2^2\right]}{2(1-\rho)^2} \qquad \lim_{x\to\infty}\left(T_2(x)-\frac{x}{1-\rho}\right)=\frac{\lambda_1(1-g_2/g_1)\mathrm{E}\left[X_1^2\right]}{2(1-\rho)^2}$$

- Example two classes
  - $g_2 > g_1$: very good for large class-2 jobs, while large class-1 jobs "close" to PS
  - $g_1 > g_2$: very bad for large class-2 jobs, while large class-1 jobs "close" to PS

Mean delay comparison of DPS with PS

- Exponential service time distributions, with means μ_j^{-1} , j=1,…, M
- Condition 1: There exists a j* such that $\begin{cases} T_k^{DPS}(x) \geq T^{PS}(x), & \forall k \geq j^* \\ T_k^{DPS}(x) \leq T^{PS}(x), & \forall k < j^* \end{cases}$
- **Theorem:** Assume Condition 1 is satisfied, and let

$$g_1 \geq g_2 \geq \ldots \geq g_M \quad \text{and} \quad \mu_1 \geq \mu_2 \geq \ldots \geq \mu_M$$

$$\text{then} \qquad \mathrm{E}\left[T^{DPS}\right] \leq \mathrm{E}\left[T^{PS}\right]$$

where $\mathrm{E}\left[T^{DPS}\right] = \sum_{j=1}^{M} \frac{\lambda_j}{\lambda} \mathrm{E}\left[T_j^{DPS}\right]$ and $\lambda = \sum_{j=1}^{M} \lambda_j$

Numerical examples: $1/\mu_1=5000, 1/\mu_2=20, 1/\mu_3=2$

$\mathrm{E}\left[T^{DPS}\right]= 2085$ $\mathrm{E}\left[T^{PS}\right]= 56.8$ $\mathrm{E}\left[T^{DPS}\right]= 38.3$

Expected conditional response time

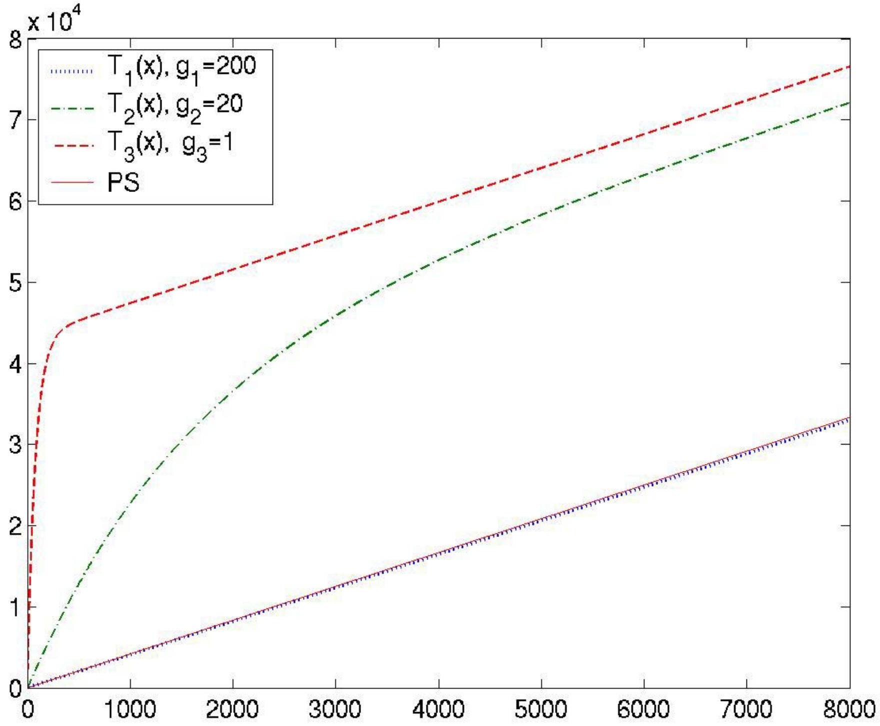

job size

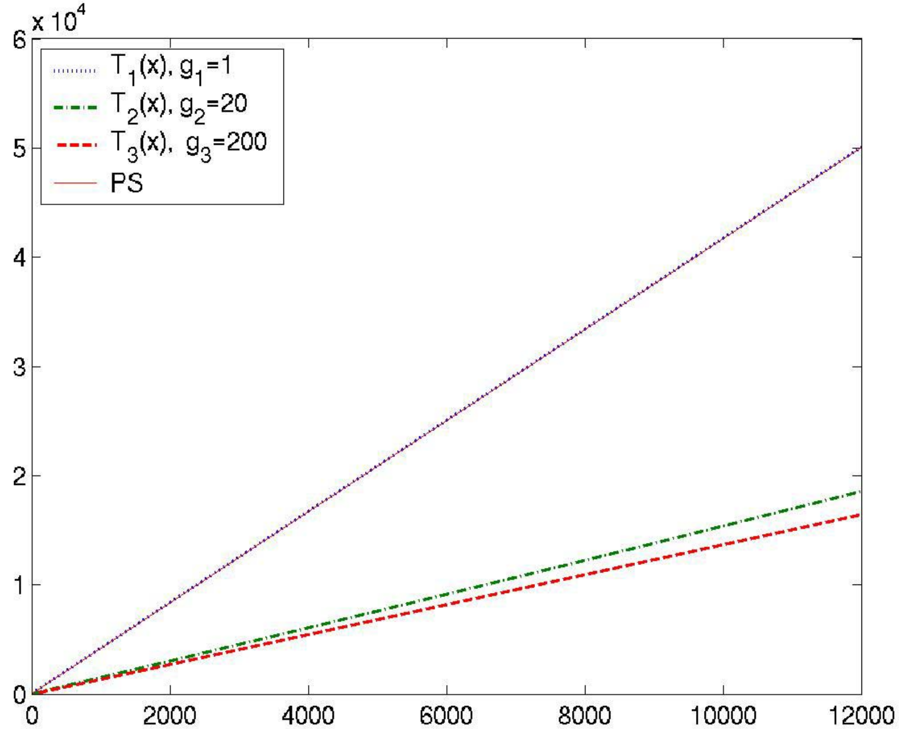

job size

Conclusion

- New results for the DPS model: Finiteness of the mean unconditional response time, Conservation law for DPS, stochastic ordering of conditional response times, exact asymptotics for the conditional response time and sufficient condition in order DPS outperforms PS.

- Several open questions:
  - Does the asymptote provide an upper/lower bound ?
  - Does the bias remain finite with infinite second moment ?
  - For general service time distributions, how must the weights be chosen in order DPS outperforms PS ?
  - ...