Unifying conservation law for single server queues[a]

Urtzi Ayesta

LAAS-CNRS

urtzi@laas.fr

EURANDOM, 20 May 2008

[a]Published in JAP, 44(4), 1078-1087, 2007

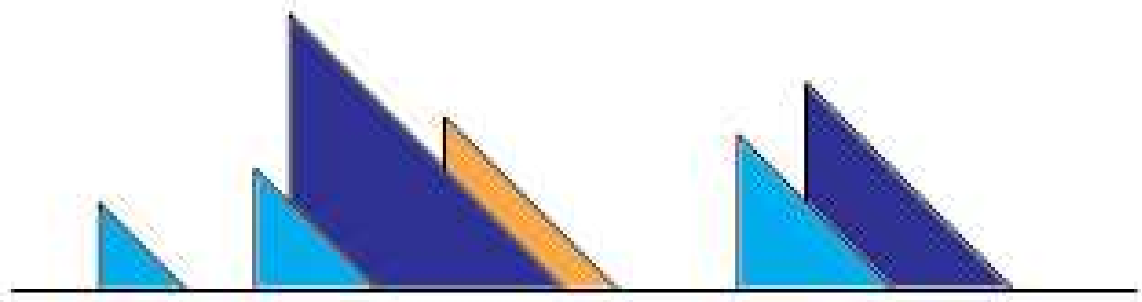

Figure 1: Cumulative burden for FCFS

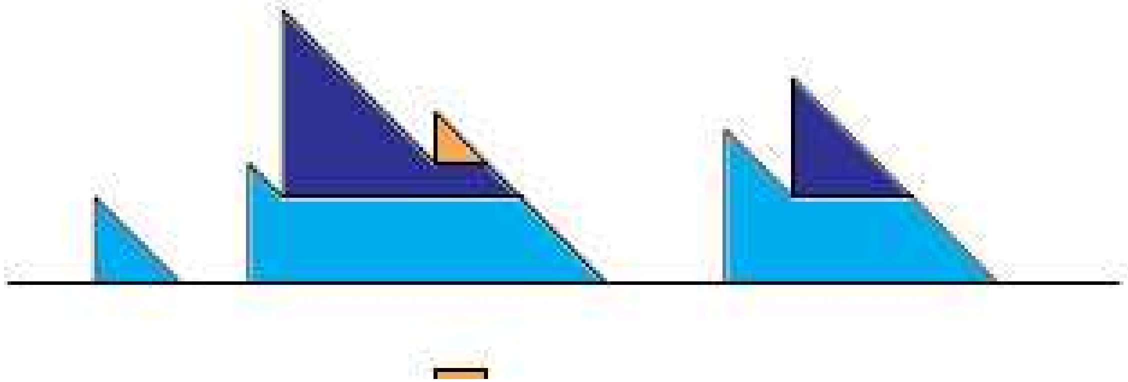

Figure 2: Cumulative burden for LCFS

Outline of the talk

- General Conservation Law

- Conservation laws for:
  - Single-class queue
  - Multi-class queue. Achievable Region approach

- Unifying conservation law

- Mean delay reduction of anticipating disciplines

Classification of Scheduling Disciplines

- Knowledge:
  - Non-anticipating
  - Anticipating

- Preemption:
  - Non-preemptive
  - Preemptive

General Conservation Law

- Consider a single-server queue with M classes.

- Let $U_j^\pi(t)$ be the unfinished work at time t of class-j jobs under policy π and let

$$U^\pi(t) = \sum_{j=1}^{M} U_j^\pi(t)$$

denote the total unfinished work in the system.

- $U^\pi(t)$ has vertical jumps at arrival epochs equal in size to the corresponding service requirements.

Definition: Work-conserving discipline

We say that the scheduling discipline is work-conserving if $\frac{dU^\pi(t)}{dt} = -1$ whenever $U^\pi(t) > 0$.

General Conservation Law: A sample path argument shows that $\sum_{j=1}^{M} U_j^\pi(t) = U^\pi(t) = U(t)$ for all work-conserving discipline $\pi \in \Pi$.

- Let $\overline{U}_j^\pi$ denote the time average unfinished work of class j, $j = 1, \ldots, M$.
- The work-conservation property implies that $\overline{U} = \sum_{j=1}^{M} \overline{U}_j^\pi$, is a constant that depends only on the inter-arrival and service time distributions.

Review Conservation Laws

The work-conserving property has led to the development of so-called work-conservation laws.

- Single-class: Integral equation for the conditional sojourn time
  - Non-anticipating discipline [Kleinrock76]
  - Anticipating discipline [O'Donovan 74]
- Multi-class: Linear relation that the expected unconditional sojourn times of the various classes must satisfy
  - Non-anticipating disciplines with exponential service time distributions [CM80]
  - Non-preemptive non-anticipating disciplines with general service time distributions [Kleinrock76]

Important body of literature *cont.*

- Textbooks: Kleinrock (1976), Gelenbe and Mitrani (1980), Heyman and Sobel (1982), Wolff (1989) and Baccelli and Brémaud (2003)

- Conservation law: Boxma (1989), Sigman (1991), Brémaud (1993), Miyazawa (1994)

- Achievable Region approach: Coffman and Mitrani (1980), Federgruen and Groenevelt (1988), Shantikumar and Yao (1992), Dacre, Glazebrook and Niño-Mora (1999), Green and Stidham (2000)

Example: Conservation law for a multi-class queue, non-anticipating discipline and exponential service time distributions

- Memoryless property: $\overline{U}_j^\pi = \frac{\overline{N}_j^\pi}{\mu_j}$
- Little's law: $\overline{N}_j^\pi = \lambda_j \overline{T}_j^\pi$
- $\overline{U}_j^\pi = \rho_j \overline{T}_j^\pi$
- Summing over all the classes: $\sum_{j=1}^M \rho_j \overline{T}_j^\pi = \sum_{j=1}^M \overline{U}_j^\pi = \overline{U}$.
- If Poisson arrivals $\overline{U} = \frac{\sum_{j=1}^M \rho_j / \mu_j}{1-\rho}$

Example: Achievable region for a multi-class queue, non-anticipating discipline and exponential service time distributions:

For $\mathcal{S} \subseteq \mathcal{M}$, let $f(\mathcal{S}) = \frac{\sum_{j\in\mathcal{S}} \rho_j/\mu_j}{1-\sum_{j\in\mathcal{S}} \rho_j}$.

For any policy π: $\sum_{j\in\mathcal{S}} \rho_j \overline{T}_j^{\pi} \geq f(\mathcal{S})$ and $\sum_{j=1}^{M} \rho_j \overline{T}_j^{\pi} = f(\mathcal{M})$

Achievable region: Polyhedra of dimension $M-1$. $M!$ vertices

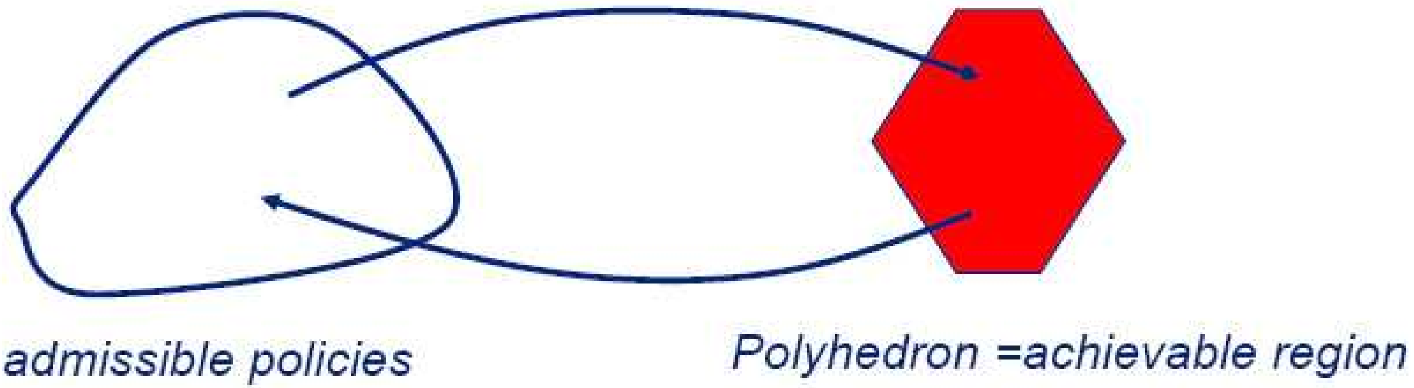

Example: 3 classes

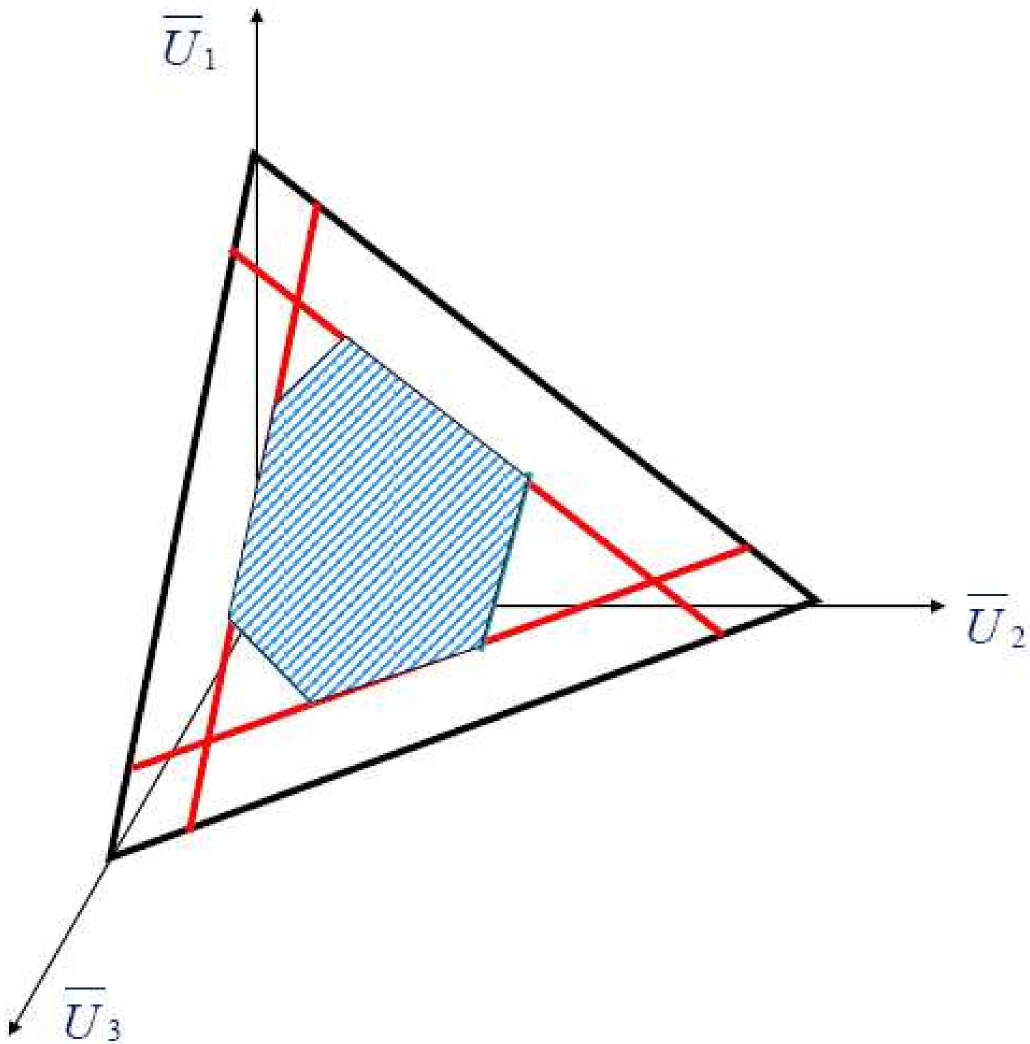

Notation and assumptions

- $GI/GI/1$ queue under work-conserving discipline
- Let λ be the mean arrival rate
  - With probability p_j an arrival is a class-j job
- Let $F_j(\cdot)$ denote the service time distribution of class j and $\overline{F}_j(\cdot) = 1 - F_j(\cdot)$ the complementary distribution.
- Stable regime, i.e., $\rho = \lambda \sum_{j=1}^{M} p_j E[X_j] < 1$
- $E[X_j^2] < \infty$, $j = 1, \ldots, M$. This ensures that the expected unfinished work at arrival epochs, $\overline{V}$, and random epochs, $\overline{U}$, is finite.

Notation and assumptions *cont.*

- $\overline{U}$ is independent of the scheduling discipline, hence $\overline{U} = \overline{U}^{FCFS}$. In the case of Poisson arrivals, by the Pollaczek-Khinchin formula we get $\overline{U} = \overline{U}^{FCFS} = \frac{\sum_{j=1}^{M} \lambda_j E[X_j^2]}{2(1-\rho)}$.
- Let $T_j^{\pi}(x)$ be the expected conditional response time of a class-j job with service time x.
- Let $T_j^{\pi}(u;x)$ denote the expected conditional time that a class-j job with total service time x spends in the system in order to obtain u units of service, $u \leq x$.
  - In particular $T_j^{\pi}(x;x) = T_j^{\pi}(x)$
  - For non-anticipating disciplines, $T_j^{\pi}(u;x) = T_j^{\pi}(u)$, for all $u \leq x$.

Unifying Conservation Law

Theorem: Consider a $GI/GI/1$ multi-class queue under a work-conserving scheduling discipline. The expected conditional response times of the various classes satisfy

$$\sum_{j=1}^{M} \lambda_j \int_{x=0}^{\infty} \overline{F}_j(x) \left(T_j^{\pi}(x) + \int_{u=0}^{x} \frac{\partial T_j^{\pi}(u;x)}{\partial x} \mathrm{d}u \right) \mathrm{d}x = \overline{U}.$$

If in addition we assume that class-j jobs, $j = 1, \ldots, M$, arrive according to a Poisson process, then

$$\overline{U} = \frac{\sum_{j=1}^{M} \lambda_j E[X_j^2]}{2(1-\rho)}.$$

Sketh of Proof:

- Since $\rho < 1$, the busy period has a finite length *with probability* 1. Regenerative process and thus stationary and ergodic.

- Let $W_j^{i,\pi}$ denote the cumulative burden of the i-th class-j job

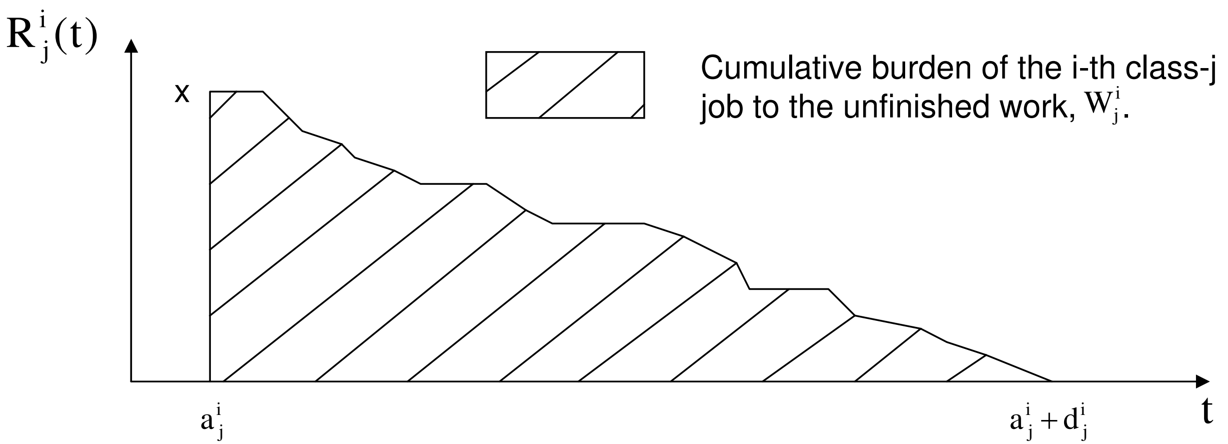

- Applying the Palm inversion formula to the unfinished work, we obtain $\overline{U}_j^{\pi} = \lambda_j E[W_j^{\pi}]$

- Summing over all classes $\sum_{j=1}^{M} \lambda_j E[W_j^{\pi}] = \sum_{j=1}^{M} \overline{U}_j^{\pi} = \overline{U}$

- Explanation for $\overline{U}_j^{\pi} = \lambda_j E[W_j^{\pi}]$: Let s be a time epoch such that $U_j^{\pi}(s) = 0$, then $\int_0^s U_j^{\pi}(t)dt = \sum_{i=1}^{A_j(s)} W_j^{i,\pi}$. Dividing by s and taking the limit $s \to \infty$:

$$\overline{U}_j^{\pi} = \lim_{s\to\infty} \frac{1}{s} \int_0^s U_j^{\pi}(t)dt = \lim_{s\to\infty} \frac{A_j(s)}{s} \frac{1}{A_j(s)} \sum_{i=1}^{A_j(s)} W_j^{i,\pi} = \lambda E[W_j^{\pi}].$$

Non-anticipating discipline:

- $T_j^{\pi}(u;x) = T_j^{\pi}(u)$ and
- Hence $\frac{\partial T_j^{\pi}(u;x)}{\partial x} = 0$, for all $0 \leq u \leq x$

Then we need to show that

$$\sum_{j=1}^{M} \lambda_j \int_{x=0}^{\infty} \overline{F}_j(x) T_j^{\pi}(x) \mathrm{d}x = \overline{U}.$$

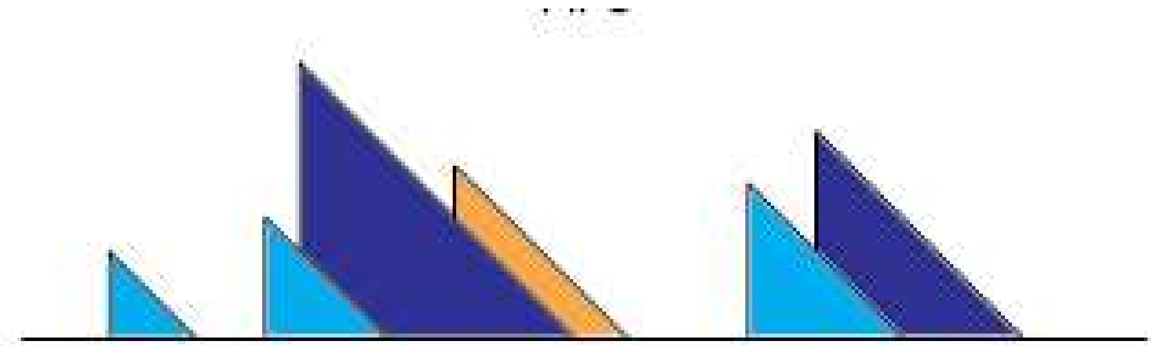

Figure 3: Cumulative burden for FCFS

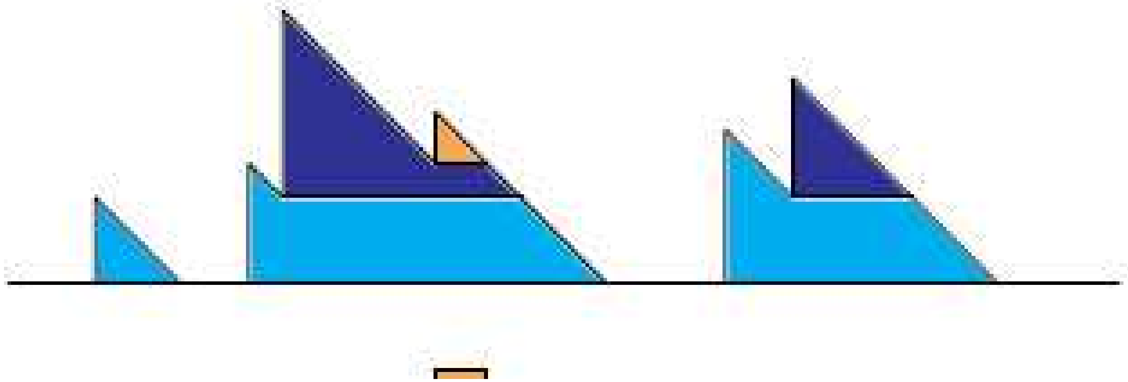

Figure 4: Cumulative burden for LCFS

Expression for $E[W_j^\pi]$

- Let $\tau_j^{i,\pi}(u;x)$ denote the amount of time that the i-th job, which has size x, needs to obtain $u \le x$ units of service.
- Since the policy is non-anticipating: $\tau_j^{i,\pi}(u) := \tau_j^{i,\pi}(u;x)$.
- $E[\tau_j^{i,\pi}(u)] = T_j^\pi(u)$.

$$\begin{aligned} W_j^{i,\pi}(x) &= \int_{t=0}^{\tau_j^{i,\pi}(x)} R_j^{i,\pi}(a^i+t)\mathrm{d}t = \int_{u=0}^{x} \tau_j^{i,\pi}(x-u;x)\mathrm{d}u \\ &= \int_{u=0}^{x} \tau_j^{i,\pi}(u)\mathrm{d}u = \int_{u=0}^{\infty} \tau_j^{i,\pi}(u)1_{\{u\le x\}}\mathrm{d}u \end{aligned}$$

- Unconditioning on x and taking expectation:

$$E[W_j^\pi] = \int_{u=0}^{\infty} T_j^\pi(u)\overline{F}_j(u)\mathrm{d}u.$$

Particular cases: Single-class

- Non-anticipating [Kleinrock,1976]: $\overline{U} = \lambda \int_{x=0}^{\infty} \overline{F}(x) T^{\pi}(x) \mathrm{d}x$.
- Anticipating [O'Donovan,1974]:

$$\overline{U} = \lambda \int_{r=0}^{\infty} r \int_{x=r}^{\infty} \left(-\partial_r T^{\pi}(x-r;x)\right) \mathrm{d}F(x),$$

- Non-preemptive anticipating [Baccelli Brémaud, 2003]. It holds that $T(u;x) = u + \overline{V}^{\pi}(x)$, $\forall 0 \le u \le x$, where $\overline{V}(x)$ denotes the expected waiting time in the queue. Substituting we get

$$\overline{U} = \frac{1}{2}\lambda E[X^2] + \int_{x=0}^{\infty} x \overline{V}^{\pi}(x) \mathrm{d}F(x)$$

Multi-class

Non-anticipating discipline and exponential service time distributions [Coffman Mitrani, 1980]

- Exponential assumption: $\overline{F}_j(x)dx = E[X_j]\mathrm{d}F_j(x)$.
- Plugging into the general equation

$$\begin{aligned}\overline{U} &= \sum_{j=1}^{M} \lambda_j \int_{x=0}^{\infty} T_j^{\pi}(x)\overline{F}_j(x)\mathrm{d}x \\ &= \sum_{j=1}^{M} \lambda_j E[X_j] \int_{x=0}^{\infty} T_j^{\pi}(x)\mathrm{d}F_j(x) = \sum_{j=1}^{M} \rho_j \overline{T}_j^{\pi},\end{aligned}$$

Multi-class

Non-preemptive non-anticipating discipline and general service time distributions [Schrage,1970]

- $T_j(x) = x + \overline{V}_j^{\pi}, \forall x \geq 0$, where $\overline{V}_j^{\pi}$ denotes the expected waiting time in the queue for a class-j job
- Substituting:

$$
\begin{aligned}
\overline{U} &= \sum_{j=1}^{M} \lambda_j \int_{x=0}^{\infty} T_j^{\pi}(x) \overline{F}_j(x) \mathrm{d}x \\
&= \frac{1}{2} \sum_{j=1}^{M} \lambda_j E[X_j^2] + \sum_{j=1}^{M} \lambda_j \overline{V}_j^{\pi} \int_{x=0}^{\infty} \overline{F}_j(x) \mathrm{d}x \\
&= \frac{1}{2} \sum_{j=1}^{M} \lambda_j E[X_j^2] + \sum_{j=1}^{M} \rho_j \overline{V}_j^{\pi}.
\end{aligned}
$$

Application to DPS

- If there are $N_k(t)$, $k = 1, \ldots, M$, class-k jobs in the queue, a class-k job gets served with rate

$$\frac{g_k}{\sum_{j=1}^{M} g_j N_j(t)}$$

Using the conservation law $\overline{U} = \sum_{j=1}^{M} \lambda_j \int_{x=0}^{\infty} T_j^{\pi}(x) \overline{F}_j(x) \mathrm{d}x$, it can be shown that:

Theorem: In a DPS queue with Poisson arrivals it holds:

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

Performance of Anticipating Disciplines

Proposition: Assume exponentially distributed service times. Let π_1 be a non-anticipating discipline and let π_2 be an anticipating discipline such that for all $0 \leq u \leq x$,

$$\frac{\partial T^{\pi_2}(u,x)}{\partial x} \geq 0.$$

Then:

$$\overline{T}^{\pi_1} \geq \overline{T}^{\pi_2}.$$

Sketch of the proof:

- Since π_1 is non-anticipating it follows:

$$\overline{U} = \lambda \int_{x=0}^{\infty} \overline{F}(x) T^{\pi_1}(x) \mathrm{d}x = \lambda E[X] \int_{x=0}^{\infty} T^{\pi_1}(x) \mathrm{d}F(x) = \rho \overline{T}^{\pi_1}.$$

Sketch of the proof (cont.):

- In the case of π_2, the conservation law can be written as

$$\overline{U} = \rho \overline{T}^{\pi_2} + \lambda \int_{x=0}^{\infty} \overline{F}(x) \int_{u=0}^{x} \frac{\partial T^{\pi_2}(u;x)}{\partial x} \mathrm{d}u \mathrm{d}x.$$

- Taking the difference we obtain

$$\rho \left(\overline{T}^{\pi_1} - \overline{T}^{\pi_2} \right) = \lambda \int_{x=0}^{\infty} \overline{F}(x) \int_{u=0}^{x} \frac{\partial T^{\pi_2}(u;x)}{\partial x} \mathrm{d}u \mathrm{d}x.$$

Condition $\frac{\partial T^{\pi_2}(u,x)}{\partial x} \geq 0$ holds $\pi = \{SRPT, SPT\}$.
Plausible to hold for other size-based policies as SMART, FSP etc.

Conclusions

- Multi-class generalization of the Swiss-army formula [Brémaud,93] or $H = \lambda G$ [Brumelle,71]:

$$\sum_{j=1}^{M} H_j = \sum_{j=1}^{M} \lambda_j G_j$$

- Achievable region for anticipating disciplines:

$$\lambda \int_{x=a}^{b} \overline{F}(x) \left(T(x) + \int_{u=0}^{x} \frac{\partial T(u;x)}{\partial x} \mathrm{d}u \right) \mathrm{d}x \geq \overline{U}(a \leq X \leq b).$$

and

$$\lambda \int_{x=a}^{b} \overline{F}(x) \left(T(x) + \int_{u=0}^{x} \frac{\partial T(u;x)}{\partial x} \mathrm{d}u \right) \mathrm{d}x = \overline{U}(0, \infty).$$