CENTRE NATIONAL
DE LA RECHERCHE
SCIENTIFIQUE

Optimal scheduling in a stochastic single server queue

Urtzi Ayesta (LAAS-CNRS)

9 ième Atelier en Evaluation de Performances,
Aussois, mardi 3 juin, 2008

Why to schedule?

- Simple example. Two jobs available, $S_1=10$ $S_2=1$

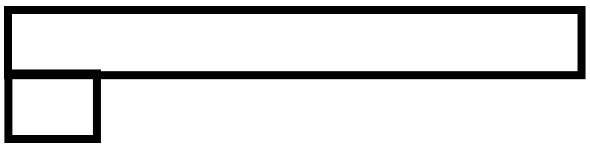

- Policy 1,2: E[T]=(10+11)/2=10,5
- Policy 2,1: E[T]=(1+11)/2=6

- General conclusion: Giving preference to shorter jobs improves the overall performance.
  - The improvement becomes more significant as the variability of the service times increases.

Outline of the talk

- General setting and optimization criteria
- Stability and Little's law
- Single class:
  - Anticipating policies: Optimality of SRPT
  - Non anticipating policies:
    - Optimality of LAS (or FB)
    - Gittins index
- Multi-class systems: Achievable region approach
- Current applications:
  - Scheduling in a network

– A Single Server

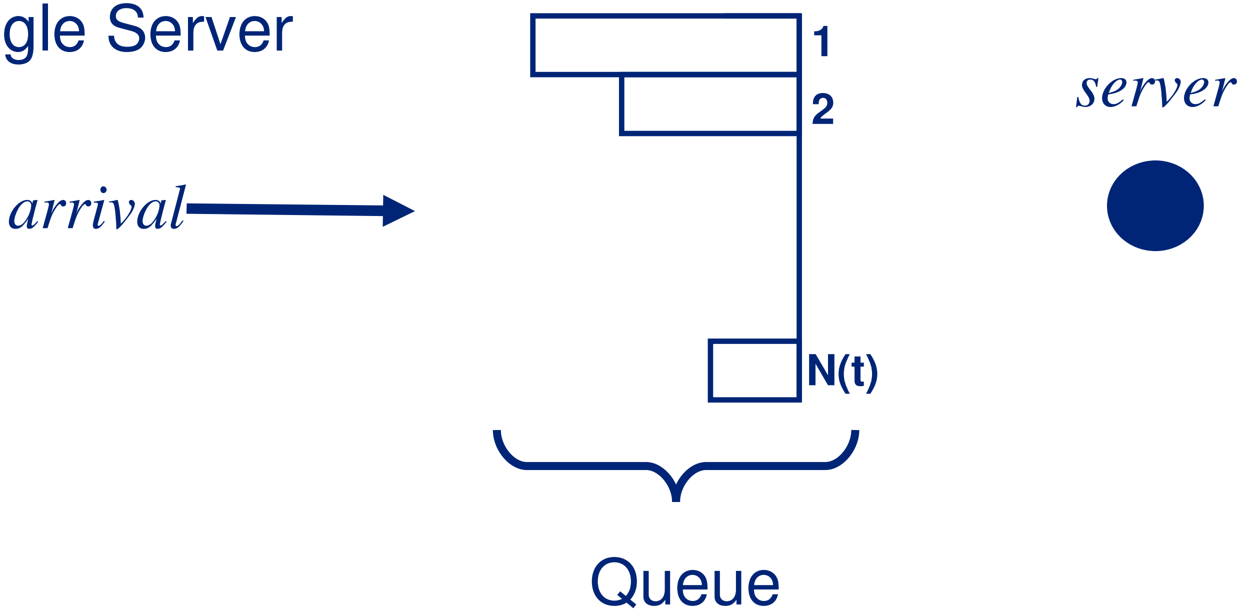

– Let W(t) denote the total work in the system at time *t*

– Let N(t) denote the number of jobs in the system at time *t*.

– For general arrival process and service times, the value of W(t) at arrival epochs can be interpreted as a one dimensional random walk

– For Poisson arrival process and exponential service times, N(t) is a Markovian birth and death process.

Dynamics of the single server

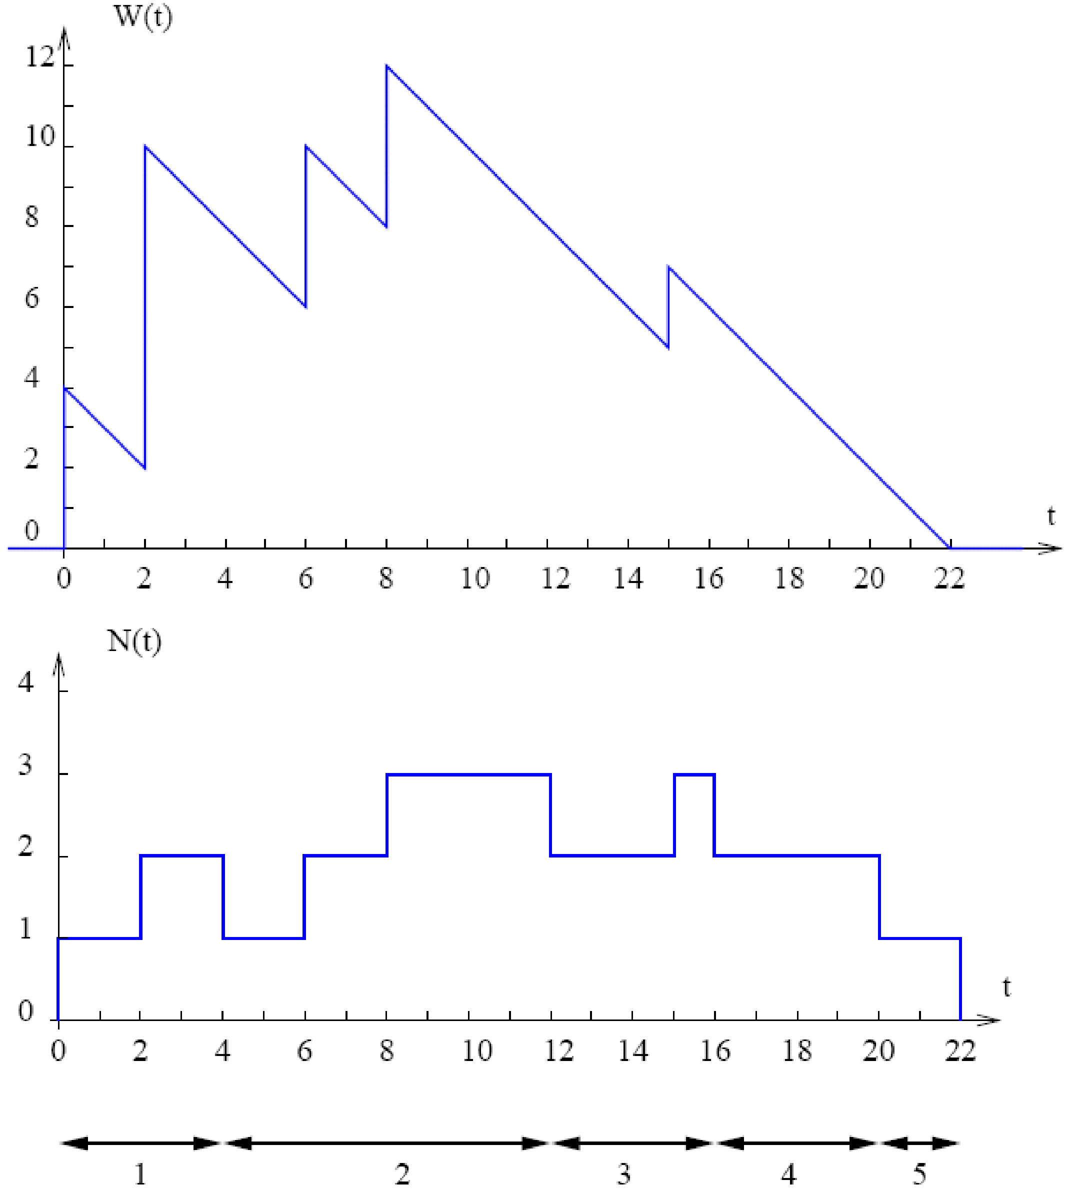

Numéro	Arrivée	Service
1	0	4
2	2	8
3	6	4
4	8	4
5	15	2

- The evolution of W(t) is independent of the scheduling policy
- The evolution of N(t) does depend on the policy

Number of jobs for various policies

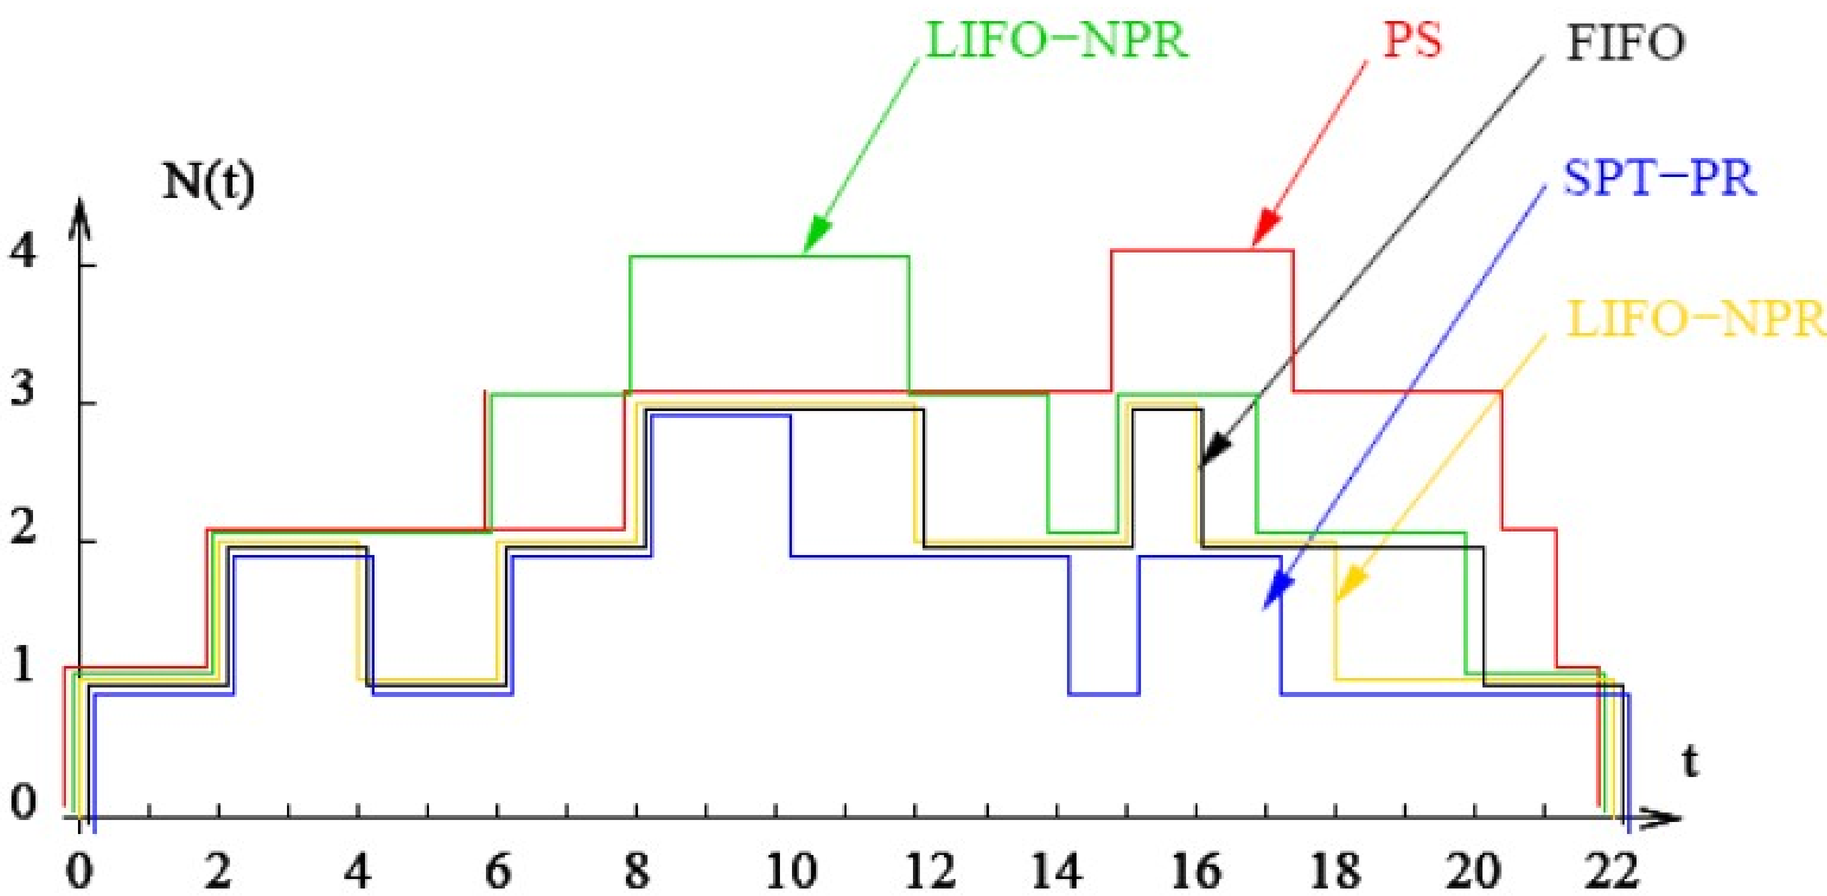

Optimization criteria

- Determine the scheduling discipline that minimizes some performance criterion:
  - Single class system:
    - Sample-path: Minimize $N(t)$ for all t
    - Stochastic ordering: Minimize $P[N(t)>k]$ for all t and all k
    - Mean: Minimize E[N], where N denotes the number of jobs in steady-state
  - Multi-class: Minimize $\Sigma_i c_i E[N_i]$
- Ergodic stochastic process: E[N] et E[T] (sojourn time) are proportional (Little's law [Little61])

Classification of Scheduling disciplines

- Knowledge:
  - Anticipating
  - Non-anticipating

- Preemption:
  - Preemptive
  - Non-preemptive

Stability

- The total workload in the system is independent of the scheduling discipline

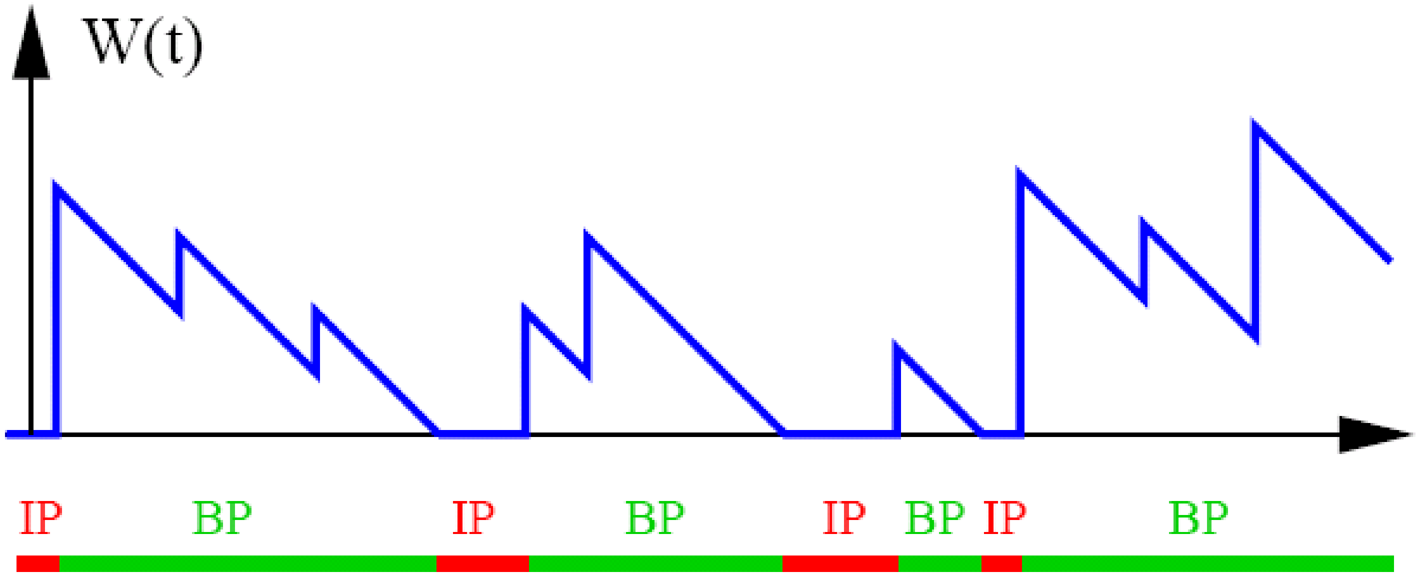

- N(t)=0 if and only if W(t)=0
- The stability condition is independent of the scheduling discipline
- Let $\rho = E[S]/E[\alpha] < 1$ denote the load in the system, where $E[\alpha]$ is the mean interarrival time, and $E[S]$ the mean service time

- Theorem: The stochastic process is stable if and only if $\rho < 1$ [Lindley52,KW55,Loynes62]

Little's law

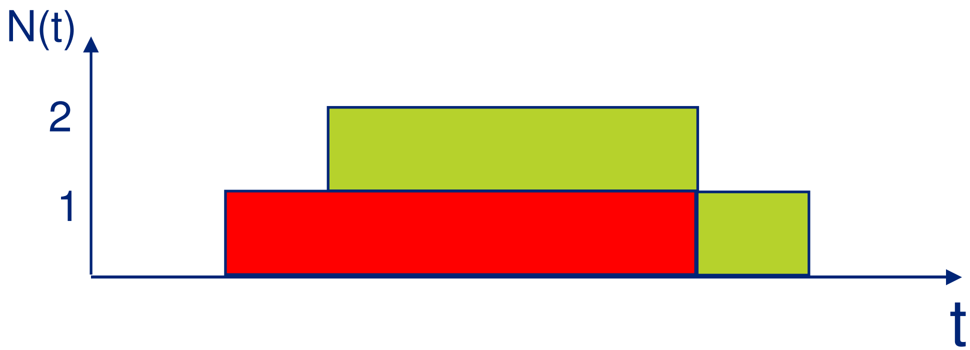

$$\int_0^t N(s)ds = T_1 + T_2$$

– In general we have $\frac{1}{t}\int_0^t N(s)ds \approx \frac{1}{t}\sum_{i=1}^{A(t)} T_i = \frac{A(t)}{t}\frac{1}{A(t)}\sum_{i=1}^{A(t)} T_i$

– As t $\longrightarrow \infty$, A(t) / t = λ, where λ is the mean arrival rate

– As t $\longrightarrow \infty$, $\frac{1}{A(t)}\sum^{A(t)} T_i = E[T]$

– Note that $\lim \frac{1}{t}\int_0^t N(s)ds = \mathrm{E[N]}$ so we get $\lambda E[T] = E[N]$

Anticipating disciplines: Optimality of SRPT

- SRPT : Shortest remaining processing time first
- Theorem [Schrage68,Smith78]: SRPT minimizes the number of jobs in the system sample-path wise, that is, for all *t*

$$N^{SRPT}(t) \leq N^{\pi}(t)$$

  for any other admissible scheduling discipline π .
- Pure sample-path result: Arrival process and service time requirements can be arbitrary (any correlation structure is allowed)
- [Sketch of the proof] Order the jobs such that

$$R_1(t) \geq R_2(t) \geq .... \geq R_{N(t)}(t)$$

- By definition of SRPT, for all *t*

$$R_1^{SRPT}(t) \geq R_1^{\pi}(t)$$

Anticipating: Optimality of SRPT (cont.)

– It can be shown that for any *j* and for all *t*

$$\sum_{i=1}^{j} R_i^{SRPT}(t) \geq \sum_{i=1}^{j} R_i^{\pi}(t)$$

But this together with the fact that

$$W^{SRPT}(t) = \sum_{i=1}^{N^{SRPT}(t)} R_i^{SRPT}(t) = \sum_{i=1}^{N^{\pi}(t)} R_i^{\pi}(t) = W^{\pi}(t)$$

implies that $N^{SRPT}(t) \leq N^{\pi}(t)$

Non-anticipating disciplines

- The size is not known, but we know the *attained service* of jobs.
  - The most appropriate scheduling discipline depends on the service time distribution characteristics

- Let F(x) denote the service time distribution, that is, $F(x)=P[S\leq x]$. And let f(x) denote the density function, that is, $f(x)=dF(x)/dx$.

- Hazard rate of a distribution function:
  $h(x)dx = P[x< S \leq x+dx \mid S > x]$

$$h(x) = \frac{f(x)}{1-F(x)}$$

Non-anticipating discipline (cont.)

Many human related random variables have a decreasing hazard rate: duration of chat/voice conversations, size of downloading files etc.

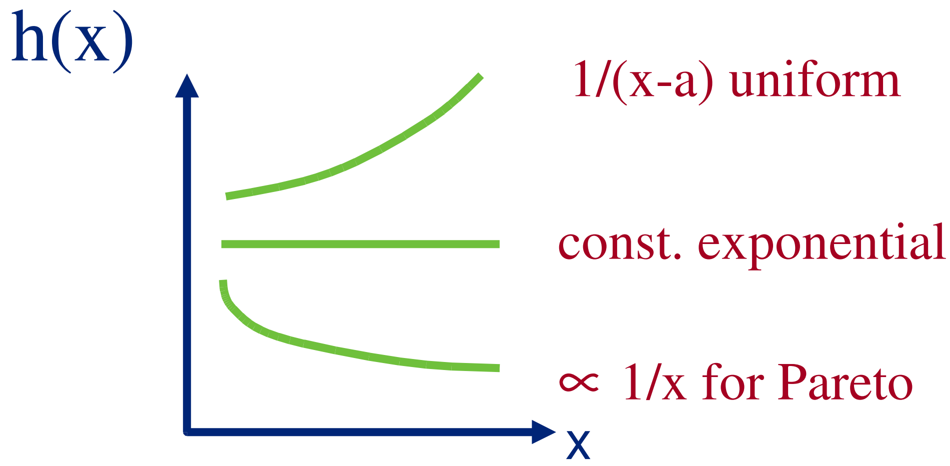

If the hazard rate is increasing, then the optimal policy is FCFS [RS89]

If the hazard rate is decreasing, then the optimal policy is FB (LAS)

Non-anticipating discipline (cont.)

- We denote by FB (Foreground-Background) or LAS (Least Attained Service) the policy that gives service to the job with the least attained service

- Theorem [RS89,Yashkov87] If the hazard rate distribution is decreasing (DFR), then LAS is stochastically optimal, that is $N^{LAS}(t) \underset{st.}{\leq} N^{\pi}(t)$, so for all *t* and all *k* we have $P[N^{LAS}(t) > k] \leq P[N^{\pi}(t) > k]$

Intuition behind the result:

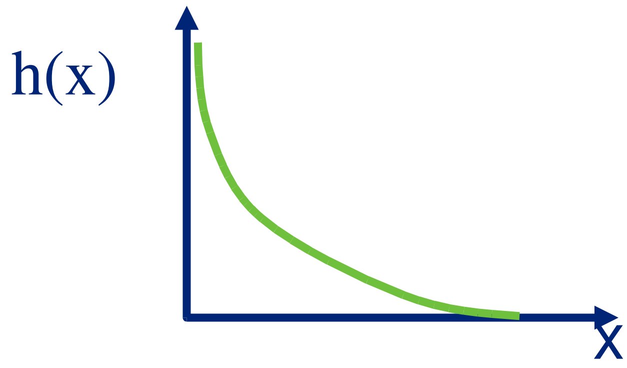

Proof: By Induction on time horizon and an interchange argument

Induction: Let T be the time horizon. Stochastic optimality of LAS holds for T=0. Assume it holds for $T=n\Delta$, and prove it holds for $T'=(n+1)\Delta$.

Contradiction: Suppose that for horizon T', π is optimal but does not follow LAS

- So at time 0, task j is processed under π but there exists a task i such that $s_j > s_i$, so $h(s_j) < h(s_i)$ (h is decreasing)
- After time Δ, the horizon is $T=n\Delta$, so π does LAS.
- Let τ be the first time that π schedules task i
- By the DFR assumption we have $h(s_j + \Delta) \le h(s_j) < h(s_i)$.
- Thus j will not be scheduled again before time τ

Interchange: Construct π' by scheduling task i at time 0, and task j at time τ

Let D' be the number of jobs completed during the first Δ time units and $\hat{D}$ be the number of jobs completed after time Δ

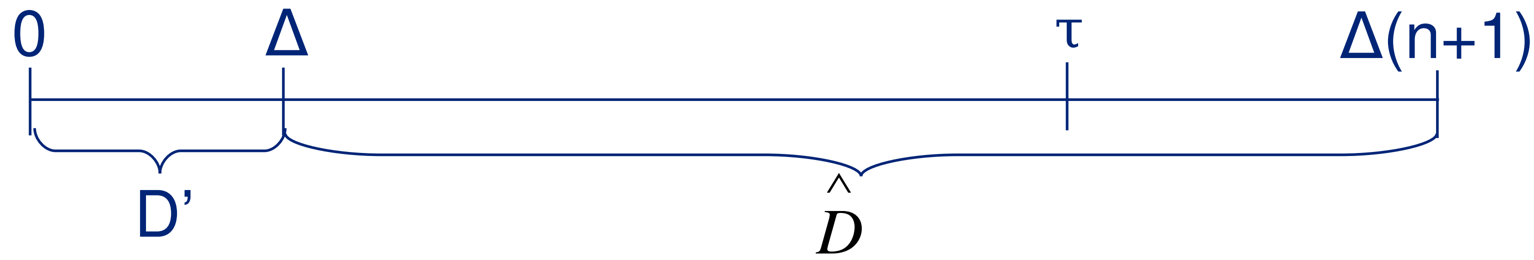

$$P_{\pi}\left(D(T')\geq k\right)=P_{\pi}\left(D(T')\geq k \mid T'\geq \tau\right)P_{\pi}\left(T'\geq \tau\right)+P_{\pi}\left(\hat{D}\geq k, T'<\tau\right)$$
$$+P_{\pi}\left(\hat{D}=k-1, T'<\tau\right)P_{\pi}\left(D'=1\right)$$

After time τ it is irrelevant: $P_{\pi}\left(D(T')\geq k \mid T'\geq \tau\right)=P_{\pi'}\left(D(T')\geq k \mid T'\geq \tau\right)$

If T'< τ, $\hat{D}$ under π and π' are the same, so:

$$P_{\pi}\left(\hat{D}\geq k, T'<\tau\right)=P_{\pi'}\left(\hat{D}\geq k, T'<\tau\right) \text{ and } P_{\pi}\left(\hat{D}=k-1, T'<\tau\right)=P_{\pi'}\left(\hat{D}=k-1, T'<\tau\right)$$

For all *k*, we end up with

$$P_{\pi'}\left(D(T')\geq k\right)-P_{\pi}\left(D(T')\geq k\right)=P_{\pi}\left(\hat{D}=k-1, T'<\tau\right)\left(P_{\pi'}\left(D'=1\right)-P_{\pi}\left(D'=1\right)\right)$$
$$=P_{\pi}\left(\hat{D}=k-1, T'<\tau\right)\left(h(s_i)-h(s_j)\right)>0$$

Starvation ?

Size-based disciplines (like SRPT and LAS) reduce the total number of jobs in the system

- Preference is given to small jobs
  - Performance of small jobs improves
- To what extent is the performance of large jobs degraded?
  - Does starvation of large jobs occur?

Compare sized-based policies with a fair policy (PS)

Fair Policy: Processor-Sharing model

- Processor-Sharing (PS):
  All present jobs in the system get a fair share of service.
  If there are N jobs, each job gets served at rate 1/N.

- An acceptable model for (i) data networks at high load (ii) web servers and (iii) CPU

- Well-studied [Kleinrock, Cohen, Kelly, Boxma, Robert …]

$$\frac{E[T \mid S = x]}{x} = \frac{1}{1-\rho}$$

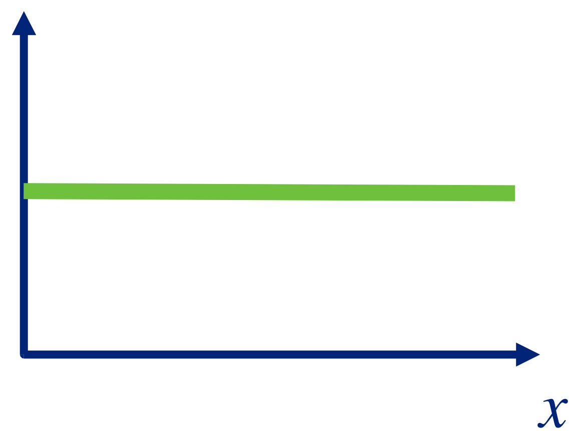

Comparison of size-based and PS

- Theorem [NQ01, HSW02] For FB and SRPT,

$$\lim_{x \to \infty} E[T \mid S = x]/x = \frac{1}{1-\rho}$$

- Theorem [BHB01] If ρ <1/2, then for any distribution and for all $x \geq 0$

$$E[T^{SRPT} \mid S = x] \leq E[T^{PS} \mid S = x]$$

- Theorem [B06] If $E[S^2]=\infty$, then for all $x \geq 0$

$$E[T^{LAS} \mid S = x] \leq E[T^{PS} \mid S = x]$$

Despite these results, size-based scheduling is rarely implemented.

$$\frac{E[T \mid S = x]}{x}$$

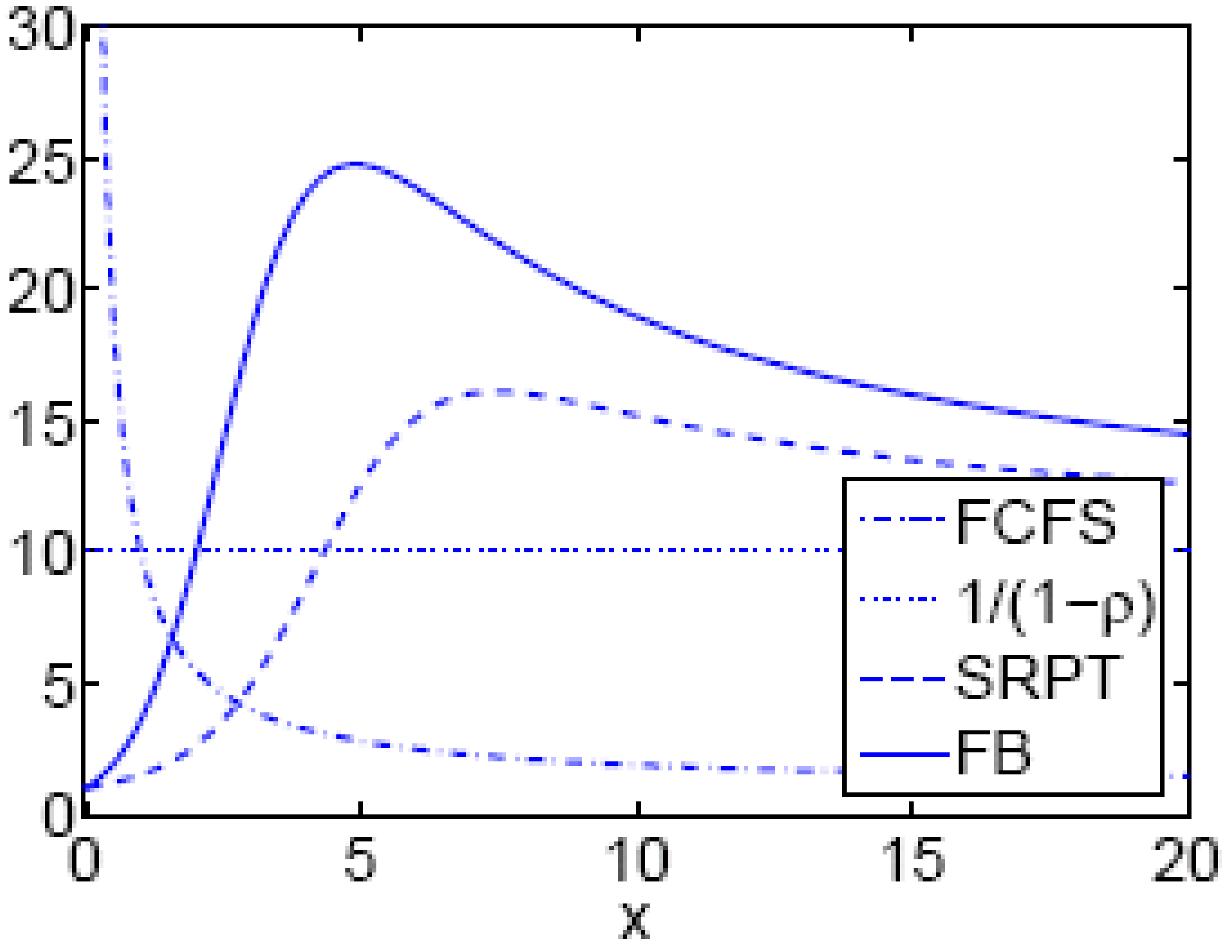

Non-anticipating discipline: What when the Hazard rate is not monotone?

– Theorem [Gittins]: Let the arrival process be Poisson. The Gittins index policy minimizes the mean number of jobs in the system among all non-anticipating scheduling policies, that is

$$E[N^{Gittins}] \leq E[N^{\pi}]$$

for any other admissible scheduling discipline π .

Introduced by Sevcik [1974] for static scheduling.

Optimality with Poisson arrivals due to Gittins [1989].

– Sevcik considered M tasks and showed that the Smallest-Rank policy (equivalent to the Gittins index) minimizes $\sum_{i=1}^{M} E[T_i]$

Gittins index

Job with attained service a has the Gittins index $G(a) = \sup_{\Delta \geq 0} J(a, \Delta)$

with $$J(a,\Delta) = \frac{\int_a^{a+\Delta} f(y)\,dy}{\int_a^{a+\Delta} \overline{F}(y)\,dy} = \frac{\text{reward}}{\text{investment}}$$

- reward: P[a ≤ S ≤ a+Δ | S>a]
- investment: E[min(S - a, Δ) | S>a]

Gittins index policy:

Pick the job with highest index value G(a) and assign him a service quota.

This job will be served until (i) it receives Δ *(a) units of service, (ii) it departs from the system or (iii) a new job with higher Gittins index arrives to the queue

Multi-Class single server

We denote for all *i*, $\boldsymbol{\mu}_i = \dfrac{1}{E[S_i]}$

Theorem:

Consider a single-server multi-class queue with either:

- Non-anticipating discipline and exponential service time distribution
- Non-preemptive discipline and general service time distribution

Then an objective function of the form $\min\left\{\sum_j c_j E[N_j^{\boldsymbol{\pi}}], \boldsymbol{\pi} \in \boldsymbol{\Pi}\right\}$

is optimized by the strict-priority rule $\pi(\varphi)$, where

$$c_{\varphi_1}\mu_{\varphi_1} \geq c_{\varphi_2}\mu_{\varphi_2} \geq \cdots \geq c_{\varphi_M}\mu_{\varphi_M}$$

$\rightarrow$Known as the cμ -rule. Equivalent to the Smith's rule [Smith56]

Achievable region [Coffman&Mitrani1980]

– Seeks solutions to stochastic optimization problems by
  1. characterize the set of all possible performances (the achievable region)
  2. optimize the overall performance objective over the achievable region
  3. identify the corresponding policy

– Used to address a wide variety of control problems:

  multi-class queues (single and multiple server), indexable systems, multi-armed bandits, special classes of queueing networks etc.

  $\rightarrow$ Here we focus on multi-class single server systems

Achievable region: Conservation law

– Let *N*={1,2,…M} denote the set of classes

– Let $W_j^{\pi}(t)$ be the unfinished work at time t of class-j jobs under policy π

and let $$W(t) = \sum_{j=1}^{M} W_j^{\pi}(t)$$ denote the total work in the system

General Conservation law:

A sample-path argument shows that *W*(*t*) is independent of the policy π

Example: Conservation law for non-anticipating discipline and exponential service times

- At every time t $\quad W_j^{\pi}(t) = \sum_{i=1}^{N_j^{\pi}(t)} R_i(t).$
- Memoryless property implies $E(W_j^{\pi}) = E(N_j^{\pi})\dfrac{1}{\mu_j}$
- Summing over all classes: $$E(W) = \sum_{j=1}^{M} \frac{E(N_j^{\pi})}{\mu_j}$$

Conservation law

If Poisson arrivals, then $E(W) = \dfrac{\sum_{j=1}^{M} \rho_j / \mu_j}{1-\rho}$, and the conservation law

becomes $$\sum_{j=1}^{M} \frac{E(N_j^{\pi})}{\mu_j} = \frac{\sum_{j=1}^{M} \rho_j / \mu_j}{1-\rho}$$

Example (cont.): Achievable region

For any $S \subseteq M$ let $f(S) = \dfrac{\sum_{j \in S} \rho_j / \mu_j}{1 - \rho_S}$

Conservation law gives: $\displaystyle \sum_{j=1}^{M} \frac{E(N_j^{\pi})}{\mu_j} = \frac{\sum_{j=1}^{M} \rho_j / \mu_j}{1 - \rho}$

So for any admissible policy π : $\displaystyle \sum_{j \in M} \frac{E(N_j^{\pi})}{\mu_j} = f(M)$ and $\displaystyle \sum_{j \in S} \frac{E(N_j^{\pi})}{\mu_j} \geq f(S), \forall S$

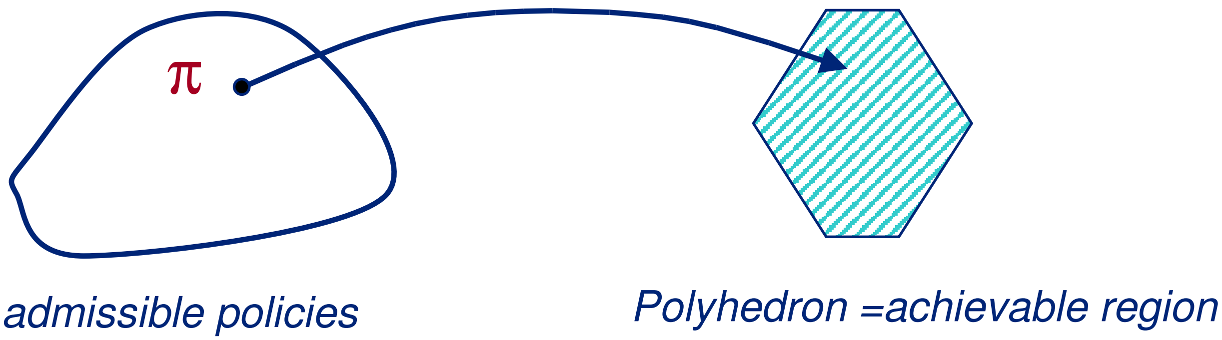

Example: 3 classes

Corners of achievable region correspond to priority strategies,

e.g. policy 1>2>3

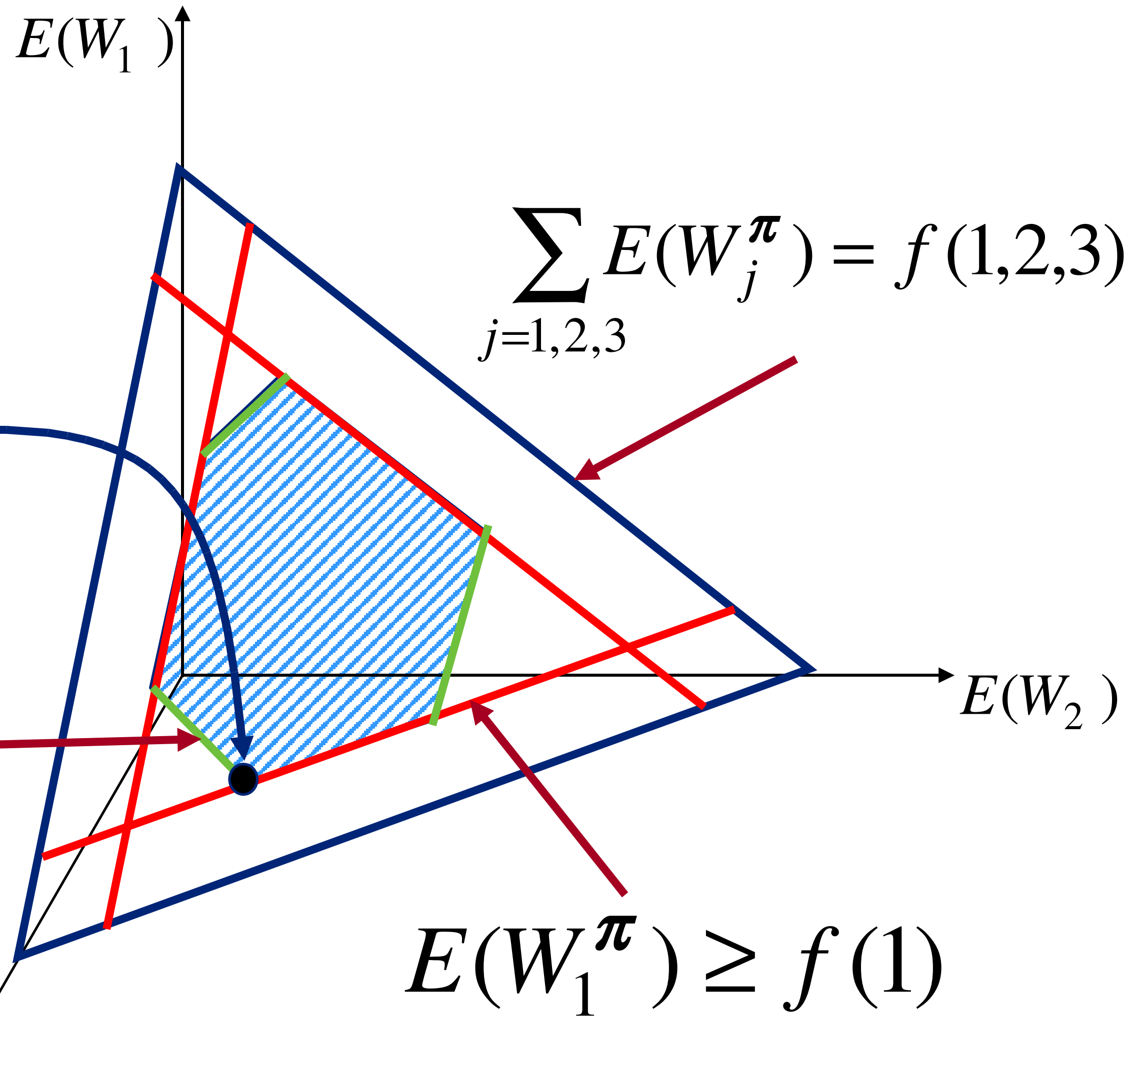

$$\sum_{j=1,2} E(W_j^{\pi}) \geq f(1,2)$$

In general: Achievable region is a polyhedra of dimension M-1 with M! vertices

Example: 3 classes

Optimality of cμ -rule:

The minimum of the function

$$\sum_j c_j E[N_j^{\pi}] = \sum_j \mu_j c_j E[W_j^{\pi}]$$

will be achieved in one of the vertices (linear programming argument) of the achievable region and corresponds to the priority policy:

$$c_{\varphi_1}\mu_{\varphi_1} \geq c_{\varphi_2}\mu_{\varphi_2} \geq c_{\varphi_3}\mu_{\varphi_3}$$

$$E(W_j^{\pi}) = E(N_j^{\pi})\frac{1}{\mu_j}$$

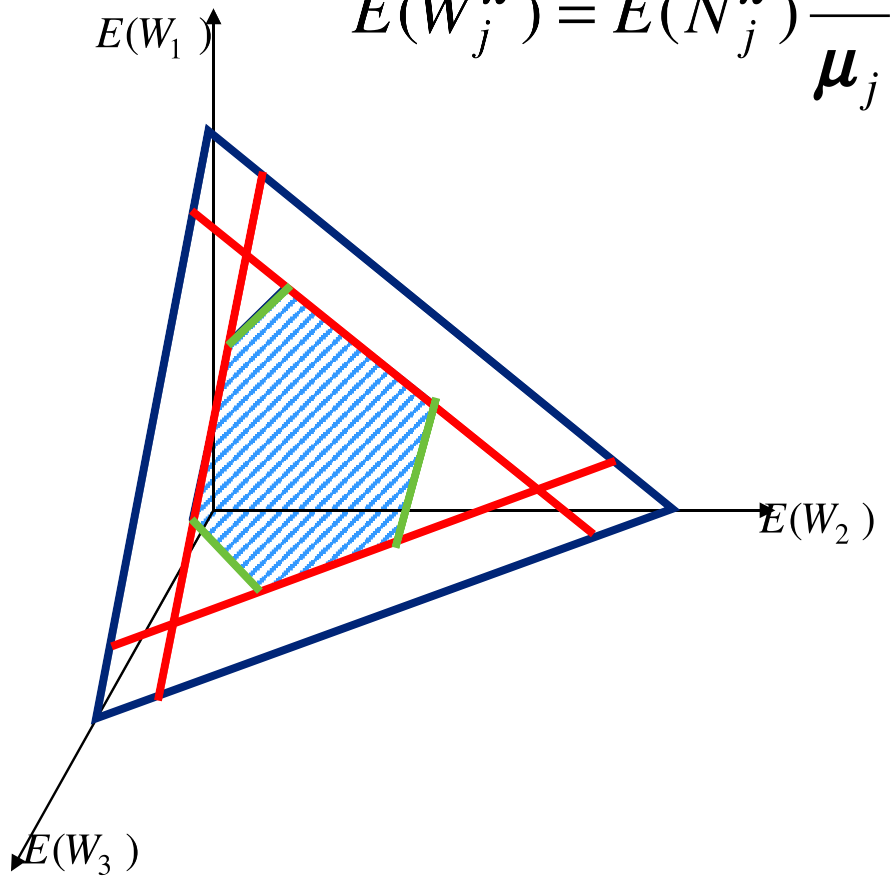

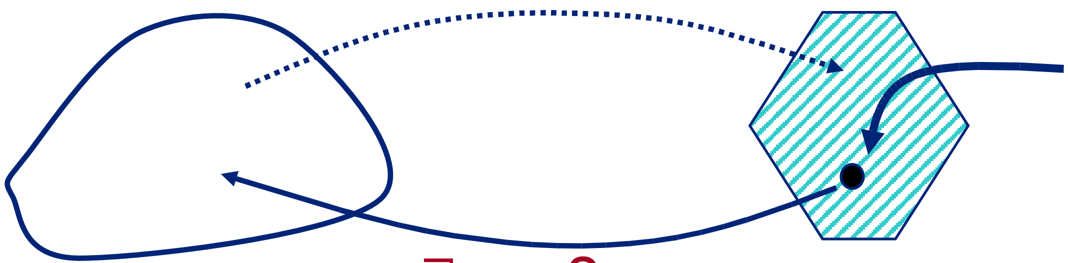

Result: for every (W_1, W_2, W_3) in the achievable region, there exists a π such that $(E(W_1^{\pi}), E(W_2^{\pi}), E(W_3^{\pi})) = (W_1, W_2, W_3)$

Proof based on the convexity of the achievable region:

– A point in the interior can be expressed as a convex combination of the vertices of the achievable region

$$\vec{W} = \sum_{j \in vertex} \alpha_j \vec{W}^{\varphi_j}, \text{ with } \sum_{j \in vertex} \alpha_j = 1$$

– Take for π the mixing strategy: At the beginning of the busy period, use priority policy φ_j with probability α_j

TABLE 1

Indexable Problems and their Performance Regions

System	*Criterion*	*Indexability*	*Performance region*
Batch of jobs	LC[a]	Smith (1956): D[b]	Queyranne (1993): D, P[c]
		Rothkopf (1966b)	This paper: P
	DC[d]	Rothkopf (1966a): D	This paper P
		Gittins & Jones (1974)	
Batch of jobs with out-tree	LC	Horn (1972): D	This paper: EP[e]
		Meilijson & Weiss (1977)	
prec. constraints	DC	Glazebrook (1976)	This paper: EP
Multiclass $M/G/1$	LC	Cox & Smith (1961)	Coffman & Mitrani (1980): P
			Gelenbe & Mitrani (1980): P
	DC	Harrison (1975a, 1975b)	This paper: EP
Multiclass[f] $M/G/c$	LC	Federgruen & Groenevelt (1988b)	Federgruen & Groenevelt (1988b)
		Shanthikumar & Yao (1992)	Shanthikumar & Yao (1992): P
Multiclass $G/M/c$	LC	Federguen & Groenevelt (1988a)	Federgruen & Groenevelt (1988a)
		Shanthikumar & Yao (1992)	Shanthikumar & Yao (1992): P
Mutliclass Jackson network[g]	LC	Ross & Yao (1989)	Ross & Yao (1989): P
Multiclass $M/G/1$	LC	Klimov (1974)	Tsoucas (1991): EP
with feedback	DC	Tcha & Pliska (1977)	This paper: EP
Multiarmed bandits	DC	Gittins & Jones (1974)	This paper: EP
Branching bandits	LC	Meilijson & Weiss (1977)	This paper: EP
	DC	Weiss (1988)	This paper: EP

Source: Bertsimas, D. and Niño-Mora, J. (1996). Conservation laws, extended polymatroids and multiarmed bandit problems: a polyhedral approach to indexable systems. Mathematics of Operations Research, vol. 21 no. 2, 257-306.

Bandwidth-Sharing networks

– Traffic flows (web pages, email, music and movies…)
– One flow may consist of many packets

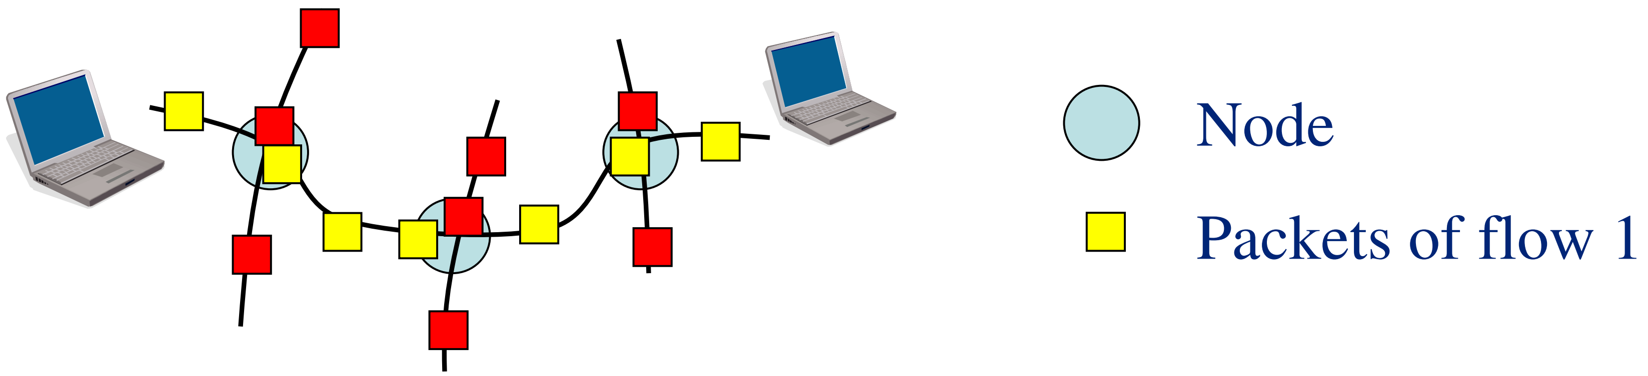

– A flow is served **simultaneously** at the same service rate **at all nodes**
– What is an efficient way to schedule these flows?

Stability is scheduling dependent

- **Class *i*** is stable iff P(N_i=0) > 0
- **Network** is stable if all classes are stable

Consider a linear network

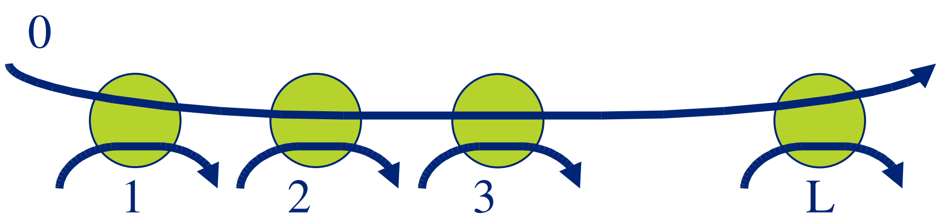

→**Necessary** condition **for stability** of network: $\boldsymbol{\rho_0+\rho_i< 1}$ **for all *i***

Stability is scheduling dependent (cont.)

– Prioritize all classes 1,…,*L*

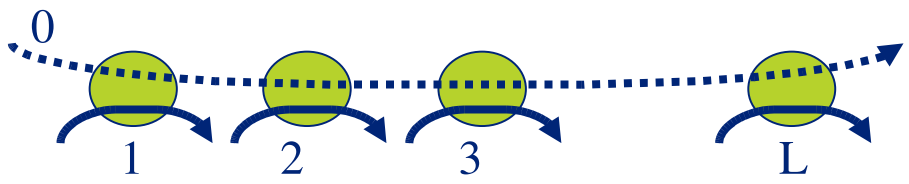

- Class 0 is served only if all classes 1,…,*L* are empty
- Stable iff $\rho_0 < \mathrm{P}\left(N_1 = 0,\ldots,N_L = 0\right) = \prod_{i=1}^{L}\left(1-\rho_i\right)$
- More stringent stability condition compared to $\boldsymbol{\rho_0+\rho_i < 1}$ **for all *i***

– Theorem [VBN05]: In a linear network, size-based scheduling (like SRPT and LAS) may lead to instability at arbitrarily low loads.

α-fair bandwidth-sharing policies

Denote by $s_i(t)$ be the rate given to class *i* flows

α –fair policy: chooses $s_i(t)$ that solves
$$\max \sum_{i=0}^{L} N_i(t)^{\alpha} \frac{s_i(t)^{\alpha}}{1-\alpha}$$
$$s.t. \quad s_0(t)+s_i(t) \le 1, \ i=1,\ldots,L$$

- **α** = 0: Maximizes throughput: $\max \sum_{i=0}^{L} s_i(t)$
- **α** ⟶ 1: Proportional fairness
- **α** = 2: TCP
- **α** ⟶ **∞**: Max-min fairness

Stability of α-fair allocation

The process $\left(\vec{N}(t)\right)_{t\geq 0} = \left(N_1(t),\ldots,N_L(t)\right)_{t\geq 0}$ is Markovian

with transition rates

$$\begin{cases} \left(\vec{N}(t)\right) \to \left(\vec{N}(t)\right) + \vec{e}_i : & \lambda_i \\ \left(\vec{N}(t)\right) \to \left(\vec{N}(t)\right) - \vec{e}_i : & \mu_i s_i(t) \end{cases}$$

Theorem [BM01]: The process $\left(\vec{N}(t)\right)_{t\geq 0}$ corresponding to the **α** –fair policy is stable under the necessary conditions:

$\boldsymbol{\rho_0 + \rho_i < 1}$ **for all *i*.**

→ The result that **α** –fair policies are stable is not restricted to linear network, but holds for general network topology!

Conclusions

Other objective functions:

- Single-Class:
  - Beyond the first moment: Reducing the variance can be very beneficial from the quality of service point of view [Wierman]
  - Measure of Fairness, how to characterize it? [Harchol-Balter, Levy,.]
- Multi-class: Objective functions that are not a linear combination of the mean number of jobs (convex etc.). [Glazebrook, Nino-Mora, Whittle …]

In practice policies like SRPT and LAS are not implemented:

- Implementation of SRPT and LAS may suffer from overheads.
- Simpler implementation of size-based scheduling: 2 classes (small jobs and large jobs). Interest in performance evaluation [Aalto, Ayesta] and implementation [Biersack et al.]

Scheduling in Wireless systems, the capacity of the system is time-varying [Borst, Bonald, Proutiere, Tse …]

- The design of simple and near optimal policies is an important and challenging problem

Bibliography

Stability

- J.Kiefer and J.Wolfowitz, On the Theory of queues with many servers, Transactions of the American Mathematical Society, 78, 1—18, 1955
- R. Loynes. The stability of a queue with non-independent inter-arrival and service times. *Proc. Cambridge Phil. Soc.*, 58:497-520, 1962.
- D.V. Lindley. The theory of queues with a single server. *Proc. Cambridge Phil. Soc.*, 48:277-289,1952.

Size based scheduling:

- L.E. Schrage and L.W. Miller, The Queue {M/G/1} with the Shortest Remaining Processing Time Discipline, Operations Research, 14, 670—684, 1966.
- D.R. Smith. A new proof of the optimality of the shortest remaining processing time discipline. *Operations Research*, 26:197-199, 1978.
- R. Righter and J.G. Shanthikumar, Scheduling multiclass single server queueing systems to stochastically maximize the number of successful departures, Probability in the Engineering and Informational Sciences, 3, 323-334, 1989
- S.F. Yashkov. Processor sharing queues: Some progress in analysis. *Queueing Systems*, 2:1-17, 1987.

Stochastic comparison

- W.A. Massey. Stochastic orderings for markov processes on partially ordered spaces. *Mathematics of Operations Research*, 12:350{367, 1987.
- F.J. Lopez and G. Sanz. Markovian couplings staying in arbitrary subsets of the state space. *Journal Applied Probability*, 39:197{212, 2002.

$c\mu$ -rule

- W.E. Smith. Various optimizers for single stage production. *Naval Res. Logist. Quart.*, 3:59-66, 1956.

Interchange arguments

- C. Buyukkoc, P. Varaya, and J. Walrand. The $c\mu$ rule revisited. *Adv. Appl. Prob.*, 17:237-238,1985.
- Z. Liu, P. Nain, and D. Towsley. Sample path methods in the control of queues. *Queueing Systems*, 21,1995.
- P. Nain, P. Tsoucas, and J. Walrand. Interchange arguments in stochastic scheduling. *Journal of Applied Probability*, 26(4):815-826, 1989.

Gittins index

- J.C. Gittins. *Multi-armed Bandit Allocation Indices*. Wiley, Chichester, 1989.
- K. Sevcik. Scheduling for minimum total loss using service time distributions. *Journal of the ACM*, 21:66{75, 1974
- I. Meilijson and G. Weiss. Multiple feedback at a single server station. *Stochastic Processes and Applications*, 5:195-205, 1977.

Asymptotic/Starvation

- M. Harchol-Balter and K. Sigman and A. Wierman, Asymptotic convergence of Scheduling policies with respect to slowdown, Performance Evaluation, 49, 241—256, 2002
- N. Bansal and M. Harchol-Balter. Analysis of SRPT scheduling: investigating unfairness. *ACM SIGMETRICS Perform. Eval. Rev.*, 29(1):279-290, 2001.
- R. Nunez-Queija, Processor-Sharing Models for Integrated-Services Networks. PhD Thesis, Eindhoven University of Technology, January 2000, ISBN 90-646-4667-8
- P. Brown. Comparing FB and PS scheduling policies. *SIGMETRICS Perform. Eval. Rev.*,34(3):18-20, 2006.

Scheduling in wireless

– S.C. Borst. User-level performance of channel-aware scheduling algorithms in wireless data networks. In *Proceedings of IEEE INFOCOM*, September 2003.

Achievable region and conservation laws

– E.G. Coffman and I. Mitrani. A characterization of waiting time performance realizable by single-server queues. *Operations Research*, 3(28):810--821, 1980.

– E. Gelenbe and I. Mitrani. *Analysis and Synthesis of Computer Systems*. London: Academic Press, 1980.

– J. Nino-Mora. Restless bandits, partial conservation laws and indexability. *Adv. Appl. Prob.*, 33(1):76-98, 2001.

– U. Ayesta. A unifying conservation law for single server queues. *Journal of Applied Probability*, 44(4):1078{1087, 2007

Stability in networks

– J.G. Dai. On positive harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *Annals of Applied Probability*, 5:49-77, 1995.

– T. Bonald and L. Massoulié. Impact of fairness on Internet performance. In *Proceedings of ACM SIGMETRICS/Performance*, pages 82-91, 2001.

– I.M. Verloop, S.C. Borst, and R. Nunez-Queija. Stability of size-based scheduling disciplines in resource-sharing networks. *Performance Evaluation*, 62:247-262, 2005.

Textbooks

– D.P. Heyman and M.J. Sobel. *Stochastic Models in operations research, Volume I: Stochastic Processes and operating characteristics*. McGraw-Hill, 1982.
– L. Kleinrock. *Queueing Systems, vol. 2*. John Wiley and Sons, 1976.
– R.W. Wolff. *Stochastic Modeling and the theory of Queues*. Prentice-Hall, 1989.
– M. El-Taha and S. Stidham. *Sample-path analysis of queueing systems*. Kluwer Academic Publishers, 1999.
– F. Baccelli and P. Bremaud, Elements of Queuing Theory: Palm Martingale Calculus and Stochastic Recurrences, Springer 2003

CENTRE NATIONAL
DE LA RECHERCHE
SCIENTIFIQUE

Urtzi Ayesta (LAAS-CNRS)

Questions, comments?

urtzi@laas.fr

Copy of the slides may be found at

www.laas.fr/~urtzi/aur/2008-06-2_aussois.pdf

9 ième Atelier en Evaluation de Performances,

Aussois, mardi 3 juin, 2008