Aalto University
School of Electrical
Engineering

Recent Advances in Age and Size-based Scheduling

Samuli Aalto

Aalto University, Finland

ITC 24
4–7 September 2012
Krakow, Poland

Tutorial outline

- Introduction
- Part 1: Fundamental scheduling results
- Part 2: The Gittins index approach revisited
- Part 3: Trade-off between size-based and opportunistic scheduling
- Final remarks

Earlier contributions

- S. Aalto, U. Ayesta and E. Nyberg-Oksanen,
  Two-level processor-sharing scheduling disciplines: Mean delay analysis, in *ACM Sigmetrics/Performance* 2004

- S. Aalto and U. Ayesta,
  Mean delay analysis of multi level processor sharing disciplines, in *IEEE Infocom* 2006

- S. Aalto, U. Ayesta, S. Borst, V. Misra and R. Nunez-Queija,
  Beyond Processor Sharing,
  ACM Sigmetrics Performance Evaluation Review, 2007

- S. Aalto and U. Ayesta,
  On the nonoptimality of the foreground-background discipline for IMRL service times, *Journal of Applied Probability*, 2006

Recent contributions

- S. Aalto, U. Ayesta and R. Righter,
  On the Gittins index in the M/G/1 queue,
  Queueing Systems, 2009

- S. Aalto, U. Ayesta and R. Righter,
  Properties of the Gittins index with application to optimal scheduling,
  Probability in the Engineering and Informational Sciences, 2011

- S. Aalto, A. Penttinen, P. Lassila and P. Osti,
  On the optimal trade-off between SRPT and opportunistic scheduling,
  in *ACM Sigmetrics* 2011

- S. Aalto, A. Penttinen, P. Lassila and P. Osti,
  Optimal size-based opportunistic scheduler for wireless systems,
  Queueing Systems, 2012 (to appear)

Introduction

M/G/1 queue

- Jobs arrive according to a Poisson process
  - IID inter-arrival times
  - exponential inter-arrival time distribution with mean $1/\lambda$
- Jobs are served by a single server
  - IID service times
  - general service time distribution with mean $E[S] = 1/\mu$

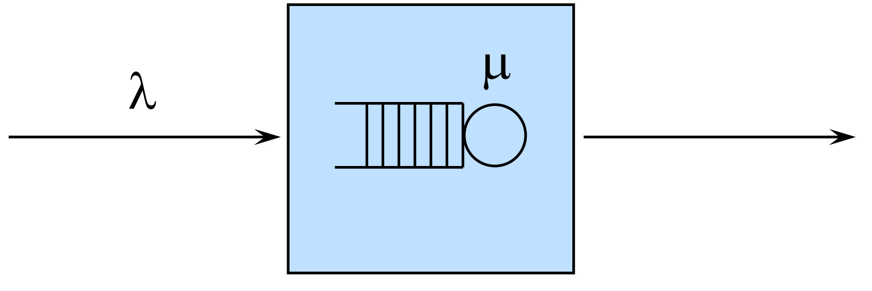

Service discipline

- Service discipline determines the way the service capacity is shared among the jobs in the system

- Service discipline is also known as
  - queueing discipline,
  - scheduling discipline, or
  - scheduling policy

- Service discipline is work-conserving
  if jobs are served whenever the system is non-empty

Some work-conserving disciplines

- First In First Out (FIFO)
  - service in the arrival order (“ordinary queue”)
  - also known as First Come First Served (FCFS)

- Processor Sharing (PS)
  - the service capacity is shared evenly among all jobs (“fair queue”)
  - ideal version of the Round Robin (RR) service discipline

Stability condition

- Any work-conserving discipline is stable if and only if

$$\rho = \frac{\lambda}{\mu} < 1$$

Optimal scheduling problem*

- Service capacity is shared among the jobs so that …
- … the mean delay $E[T]$ is minimized …
- … within the family of disciplines considered

* … in this presentation

Example: M/M/1

- For any work-conserving discipline,

$$E[T]=\frac{E[S]}{1-\rho}=E[S]\left(1+\frac{\rho}{1-\rho}\right)$$

- Conclusion: Any work-conserving discipline is optimal

Example: M/D/1

- For FIFO (by Pollaczek-Khinchin),

$$E[T] = E[S]\left(1 + \frac{\rho}{2(1-\rho)}\right)$$

- For PS (by insensitivity),

$$E[T] = \frac{E[S]}{1-\rho} = E[S]\left(1 + \frac{\rho}{1-\rho}\right) > E[S]\left(1 + \frac{\rho}{2(1-\rho)}\right)$$

- Conclusion: FIFO better than PS

Example: M/G/1

- For FIFO (by Pollaczek-Khinchin),

$$E[T] = E[S] + \frac{\lambda E[S^2]}{2(1-\rho)}$$

- For PS (by insensitivity),

$$E[T] = \frac{E[S]}{1-\rho} = E[S] + \frac{\lambda E[S]^2}{1-\rho}$$

- Conclusion: FIFO better than PS if and only if $C^2[S] \leq 1$

Service time distribution

- Coefficient of variation $C^2[S]$:

$$C^2[S] = \frac{D^2[S]}{E[S]^2} = \frac{E[S^2]}{E[S]^2} - 1$$

- Note that

$$C^2[S] \leq 1 \;\Leftrightarrow\; \frac{E[S^2]}{2} \leq E[S]^2$$

Part 1
Fundamental scheduling results

Outline of Part 1

- Service disciplines
- Service time distributions
- Gittins index approach
- Optimality results
- Summary

Service discipline categories

- Definition: Service discipline is work-conserving
  if jobs are served whenever the system is non-empty
- Definition: Service discipline is non-sharing
  if jobs are served one-by-one
- Definition: Service discipline is non-preemptive
  if jobs are served one-by-one until completion
- Definition: Service discipline is non-anticipating
  if the remaining service times are not utilized
  (while the attained service times may be utilized)

Service disciplines (1)

- First In First Out (FIFO)
  - when the server becomes free,
    the earliest arrived job is taken into service ("ordinary queue")
  - non-preemptive and non-anticipating
  - also known as First Come First Served (FCFS)

- Most Attained Service (MAS)
  - when the server becomes free,
    a job is taken into service in any non-anticipating way
  - non-preemptive and non-anticipating

Service disciplines (2)

- Processor Sharing (PS)
  - the service capacity is shared evenly among all jobs ("fair queue")
  - sharing and non-anticipating
  - ideal version of the Round Robin (RR) service discipline

- Least Attained Service (LAS)
  - the service capacity is shared evenly among the jobs with the least amount of attained service
  - sharing and non-anticipating
  - also known as Foreground Background (FB)

Service disciplines (3)

- Shortest Processing Time (SPT)
  - when the server becomes free,
    the job with the shortest service time is taken into service
  - non-preemptive and anticipating

- Shortest Remaining Processing Time (SRPT)
  - the job with the shortest remaining service time is served
  - non-sharing, preemptive, and anticipating

Service disciplines (4)

- Shortest Expected Processing Time (SEPT)
  - when the server becomes free,
    the job with the shortest expected service time is taken into service
  - non-preemptive and non-anticipating
- Shortest Expected Remaining Processing Time (SERPT)
  - the job with the shortest expected remaining service time is served
  - non-sharing, preemptive, and non-anticipating

Service discipline families

- Non-preemptive non-anticipating disciplines $\Pi^{\text{NPR-NA}}$
  - e.g. FIFO, MAS, SEPT
- Non-preemptive disciplines Π^{NPR}
  - e.g. FIFO, MAS, SEPT + SPT
- Non-anticipating disciplines Π^{NA}
  - e.g. FIFO, MAS, SEPT + PS, LAS, SERPT
- All disciplines Π
  - e.g. all above + SRPT

Outline of Part 1

- Service disciplines
- Service time distributions
- Gittins index approach
- Optimality results
- Summary

Service time distribution

- Hazard rate (HR) function $h(x)$

$$F(x) \triangleq \int_0^x f(y)dy, \qquad h(x) \triangleq \frac{f(x)}{1-F(x)}$$

- Mean residual lifetime (MRL) function $M(x)$

$$M(x) \triangleq E[S-x \mid S>x] = \frac{\int_x^\infty (1-F(y))dy}{1-F(x)}$$

Service time distribution classes (1)

- Definition: Service times are IHR [DHR] if $h(x)$ is increasing [decreasing]
- Definition: Service times are DMRL [IMRL] if $M(x)$ is decreasing [increasing]
- Definition: Service times are NBUE [NWUE] if $M(0) \geq [\leq]\ M(x)$ for any x

Service time distribution classes (2)

- IHR = Increasing Hazard Rate
- DMRL = Decreasing Mean Residual Lifetime
- NBUE = New Better than Used in Expectation

- DHR = Decreasing Hazard Rate
- IMRL = Increasing Mean Residual Lifetime
- NWUE = New Worse than Used in Expectation

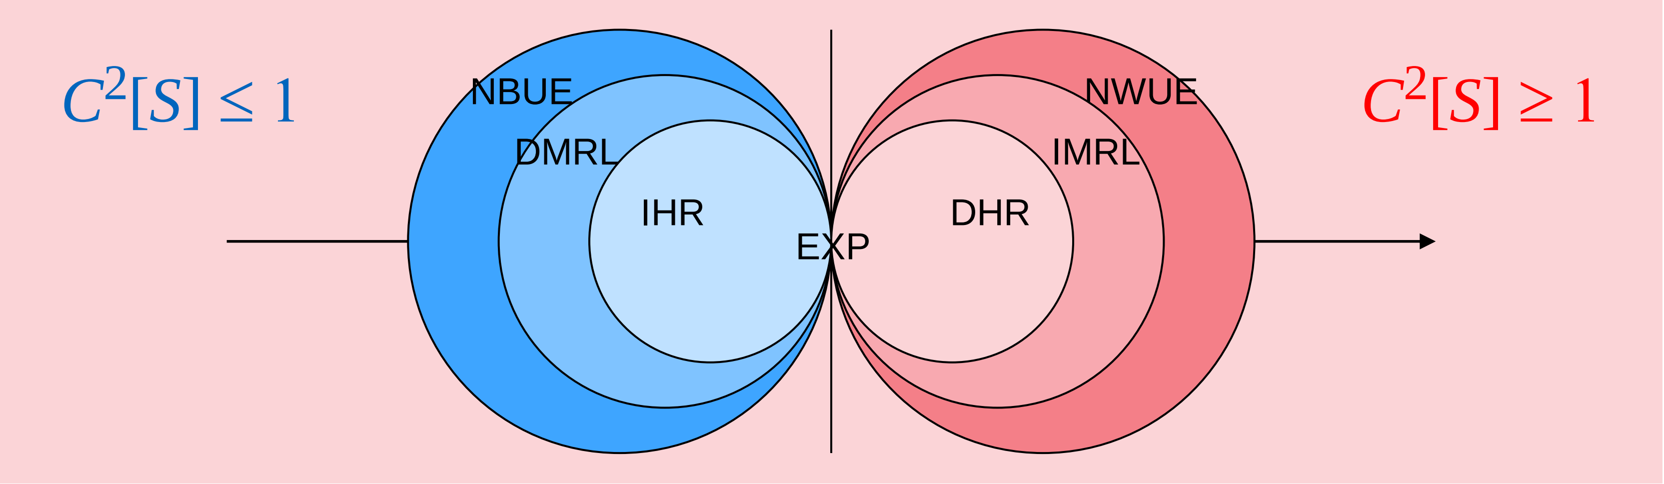

Outline of Part 1

- Service disciplines
- Service time distributions
- Gittins index approach
- Optimality results
- Summary

Hazard rate

- Remaining service time distribution:

$$P\{S - x \le y \mid S > x\} = \frac{F(x+y) - F(x)}{1 - F(x)}$$

- Hazard rate (HR) function $h(x)$:

$$h(x) \triangleq \lim_{\Delta \to 0} \frac{1}{\Delta} P\{S - x \le \Delta \mid S > x\} = \frac{f(x)}{1 - F(x)}$$

Inverse MRL

- Mean residual lifetime (MRL) function $M(x)$:

$$M(x) \triangleq E[S - x \mid S > x] = \frac{\int_x^\infty (1 - F(y))dy}{1 - F(x)}$$

- Inverse MRL function $H(x)$:

$$H(x) \triangleq \frac{1}{E[S - x \mid S > x]} = \frac{1 - F(x)}{\int_x^\infty (1 - F(y))dy}$$

Gittins index (1)

- Consider a job with
  - attained service (age) a
  - served continuously during an interval of length (at most) Δ
- Probability that the service is completed

$$P\{S - a \le \Delta \mid S > a\} = \frac{F(a+\Delta) - F(a)}{1 - F(a)} = \frac{\int_a^{a+\Delta} f(y)dy}{1 - F(a)}$$

- Mean time until the end of service or interval

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

Gittins index (2)

- Efficiency function for age a and service quota Δ:

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

- Limiting values:

$$J(a,0) = h(a), \quad J(a,\infty) = H(a)$$

Gittins index (3)

- Definition: Gittins index $G(\alpha)$ for a job with age α is

$$G(a) \hat{=} \sup_{\Delta \geq 0} J(a, \Delta)$$

- Optimal service quota for a job with age α:

$$\Delta^*(a) \hat{=} \sup\{\Delta \geq 0 \mid J(a, \Delta) = G(a)\}$$

Gittins index discipline

- Gittins index discipline (GI)
  - job i^* with the highest Gittins index $G(a_{i^*})$ is served
  - non-anticipating

- Ordinary M/G/1 queue (with a single job class):

$$G(a_{i*}) \triangleq \max_i G(a_i)$$

- Multiclass M/G/1 queue (with multiple job classes):

$$G_{k_{i*}}(a_{i*}) \triangleq \max_i G_{k_i}(a_i)$$

Optimality of the GI discipline

- Gittins (1989)
- Theorem: For any M/G/1 queue with $\rho < 1$, the GI discipline is optimal among all non-anticipating disciplines,

$$E[T^{\mathrm{GI}}] = \min\{E[T^{\pi}] \mid \pi \in \Pi^{\mathrm{NA}}\}$$

- See also Sevcik (1974), Klimov (1974, 1978)

Outline of Part 1

- Service disciplines
- Service time distributions
- Gittins index approach
- Optimality results
- Summary

Service discipline families

- Non-preemptive non-anticipating disciplines $\Pi^{\text{NPR-NA}}$
  - e.g. FIFO, MAS, SEPT
- Non-preemptive disciplines Π^{NPR}
  - e.g. FIFO, MAS, SEPT + SPT
- Non-anticipating disciplines Π^{NA}
  - e.g. FIFO, MAS, SEPT + PS, LAS, SERPT
- All disciplines Π
  - e.g. all above + SRPT

Optimality of the SEPT discipline

- Cox and Smith (1961)
- Theorem: For any M/G/1 queue with $\rho < 1$,
  the SEPT discipline is optimal
  among all non-preemptive non-anticipating disciplines,

$$E[T^{\mathrm{SEPT}}] = \min\{E[T^{\pi}] \mid \pi \in \Pi^{\mathrm{NPR\text{-}NA}}\}$$

- Special case of the optimality of the $c\mu$-rule (with $c \equiv 1$)

Interpretation by the GI approach

- For the ordinary M/G/1 queue, the result is trivial.
- Consider the multi-class M/G/1 queue. Due to the restriction to the non-preemptive disciplines, the Gittins index is only considered for $\alpha = 0$ and $\Delta = \infty$:

$$G_k(0) = J_k(0,\infty) = H_k(0) = 1/E[S_k]$$

- Thus,

$$G_{k_{i*}}(0) = \max_i G_{k_i}(0) = 1/\min_i E[S_{k_i}]$$

- Conclusion: SEPT = GI discipline

Optimality of the SPT discipline

- Cox and Smith (1961)
- Theorem: For any M/G/1 queue with $\rho < 1$,
  the SPT discipline is optimal
  among all non-preemptive disciplines,

$$E[T^{\mathrm{SPT}}] = \min\{E[T^{\pi}] \mid \pi \in \Pi^{\mathrm{NPR}}\}$$

Interpretation by the GI approach

- Define the class of the job based on its known service requirement s
- Due to the restriction to the non-preemptive disciplines, the Gittins index is only considered for $\alpha = 0$ and $\Delta = \infty$:

$$G_s(0) = J_s(0,\infty) = H_s(0) = 1/s$$

- Thus,

$$G_{s_{i*}}(0) = \max_i G_{s_i}(0) = 1/\min_i s_i$$

- Conclusion: SPT = GI discipline

Service time distribution classes

- IHR = Increasing Hazard Rate
- DMRL = Decreasing Mean Residual Lifetime
- NBUE = New Better than Used in Expectation

- DHR = Decreasing Hazard Rate
- IMRL = Increasing Mean Residual Lifetime
- NWUE = New Worse than Used in Expectation

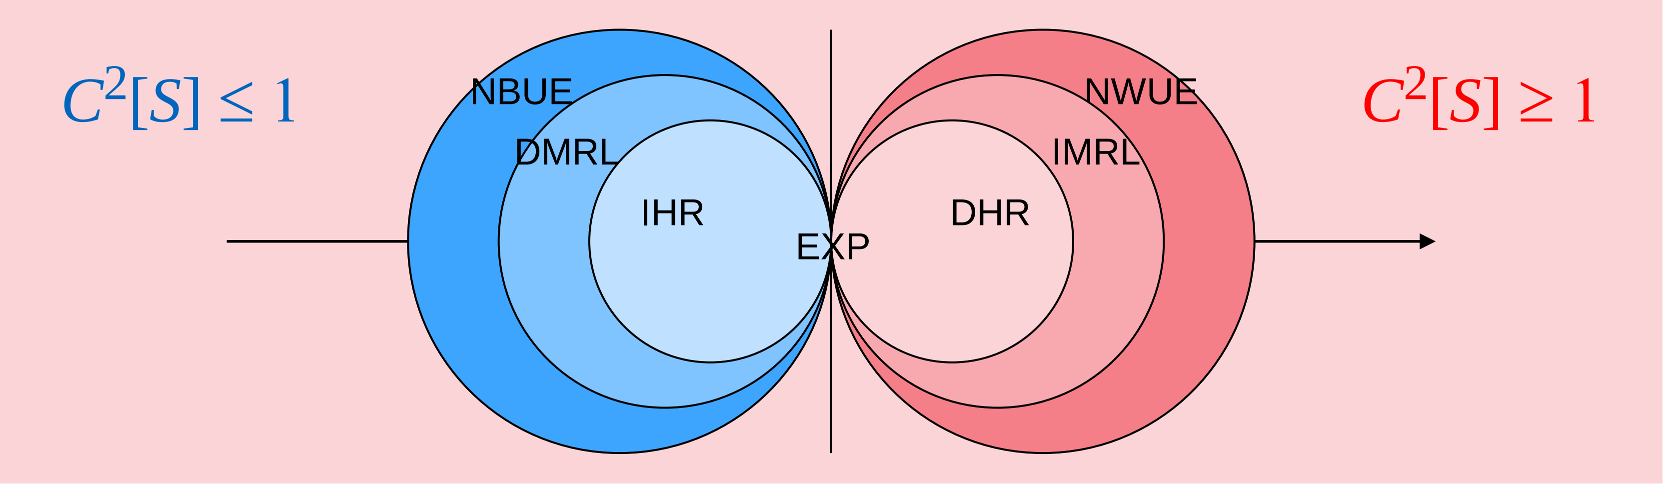

Optimality of the MAS discipline

- Righter, Shanthikumar and Yamazaki (1990)
- Theorem: For the ordinary M/G/1 queue with NBUE service times and $\rho < 1$, any MAS discipline (e.g. FIFO) is optimal among all non-anticipating disciplines,

$$\text{NBUE} \Rightarrow E[T^{\text{MAS}}] = \min\{E[T^{\pi}] \mid \pi \in \Pi^{\text{NA}}\}$$

Interpretation by the GI approach

- Aalto, Ayesta and Righter (2009)
- Lemma: For NBUE service times, $J(0,\Delta) \leq J(0,\infty)$ for all Δ.
- Lemma implies that

$$G(0) = \sup_{\Delta \geq 0} J(0,\Delta) = J(0,\infty) = H(0)$$

- On the other hand, due to the NBUE property,

$$G(a) \geq H(a) \geq H(0) = G(0)$$

- Conclusion: MAS = GI discipline

Service time distribution classes

- IHR = Increasing Hazard Rate
- DMRL = Decreasing Mean Residual Lifetime
- NBUE = New Better than Used in Expectation

- DHR = Decreasing Hazard Rate
- IMRL = Increasing Mean Residual Lifetime
- NWUE = New Worse than Used in Expectation

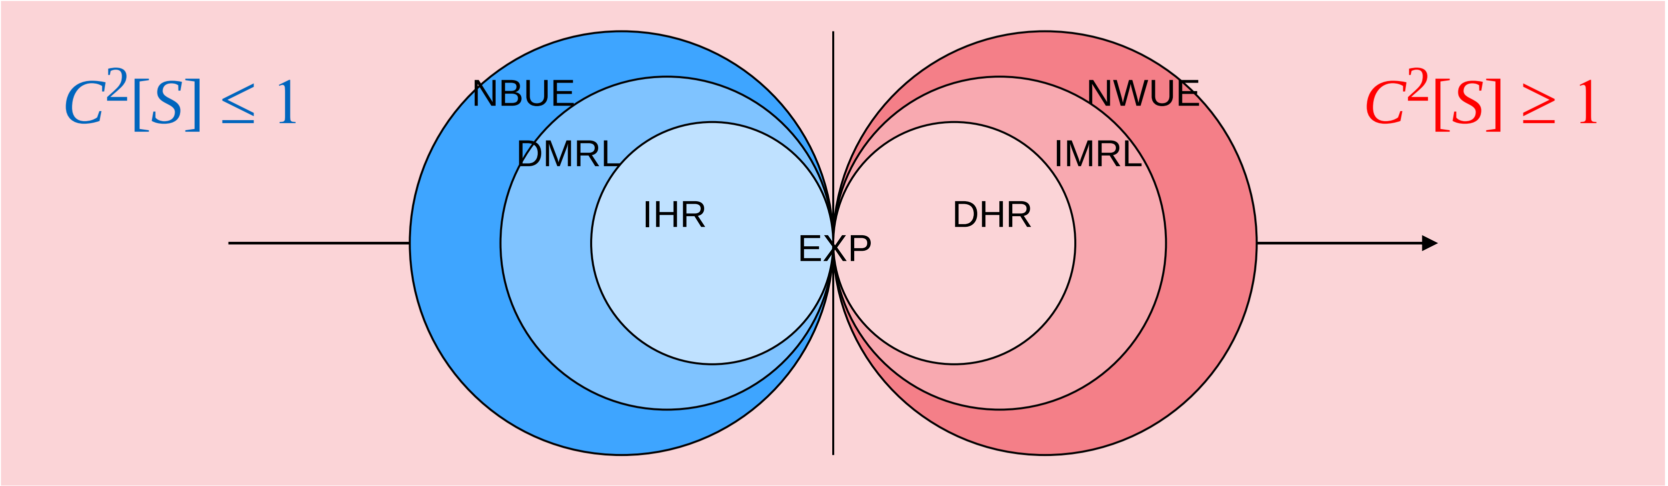

Optimality of the LAS discipline

- Yashkov (1987); Righter and Shanthikumar (1989)
- Theorem: For the ordinary M/G/1 queue with DHR service times and $\rho < 1$, the LAS discipline is optimal among all non-anticipating disciplines,

$$\text{DHR} \Rightarrow E[T^{\text{LAS}}] = \min\{E[T^{\pi}] \mid \pi \in \Pi^{\text{NA}}\}$$

- See also Aalto and Ayesta (2006)

Interpretation by the GI approach

- Aalto, Ayesta and Righter (2009)
- Lemma: For DHR service times,
  $J(a,\Delta)$ is decreasing (with respect to Δ) for all a, Δ.
- Lemma implies that

$$G(a) = \sup_{\Delta \geq 0} J(a,\Delta) = J(a,0) = h(a)$$

- On the other hand, due to the DHR property,

$$G(a_{i*}) = \max_i G(a_i) = \max_i h(a_i) = h(\min_i a_i)$$

- Conclusion: LAS = GI discipline

Optimality of the SRPT discipline

- Schrage (1968); Smith (1978)
- Theorem: For any M/G/1 queue with $\rho < 1$, the SRPT discipline is optimal among all disciplines,

$$E[T^{\mathrm{SRPT}}] = \min\{E[T^{\pi}] \mid \pi \in \Pi\}$$

Interpretation by the GI approach

- Define the class of the job based on its known service requirement s
- The Gittins index is now given by

$$G_s(a) = \sup_{\Delta \geq 0} J_s(a, \Delta) = J_s(a, s-a) = 1/(s-a)$$

- Thus,

$$G_{s_{i*}}(a_{i*}) = \max_i G_{s_i}(a_i) = 1/\min_i (s_i - a_i)$$

- Conclusion: SRPT = GI discipline

Outline of Part 1

- Service disciplines
- Service time distributions
- Gittins index approach
- Optimality results
- Summary

Summary of Part 1

Mean delay minimization in M/G/1

NPR-NA:

SEPT optimal

NPR:

SPT optimal

NA:

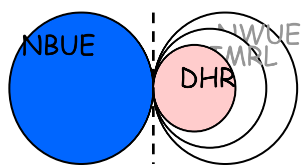

GI optimal

NBUE ⇒ MAS [SERPT] optimal

DHR ⇒ LAS [SERPT] optimal

All:

SRPT optimal

Part 2
The Gittins index approach revisited

Outline of Part 2

- Introduction
- Gittins index
- Continuity and monotonicity result
- Monotonicity in finite intervals
- Service time distribution classes
- Optimality results
- Summary

Optimal scheduling problem

- Transient system (no arrivals)
  - Given a single-server queue
    with n IID jobs and service time distribution $F(x)$,
    what is the optimal non-anticipating service policy
    so that the mean delay is minimized?
- Dynamic system (Poisson arrivals)
  - Given an M/G/1 queue
    with arrival rate λ and service time distribution $F(x)$,
    what is the optimal non-anticipating service policy
    so that the mean delay is minimized?

Optimality results

- For both problems,
  the optimal anticipating policy is SRPT,
  but it requires exact information about the service times

- For both problems,
  the optimal non-anticipating policy is GI,
  based on the amount of the attained service and
  the service time distribution

Gittins index discipline

- Gittins index discipline (GI)
  - job i^* with the highest Gittins index $G(a_{i^*})$ is served
  - non-anticipating
- Observations:
  - GI is not necessary unique
  - MAS is GI
    if and only if $G(a) \geq G(0)$ for all a
  - LAS is GI
    if and only if $G(a)$ is decreasing for all a

Example

$$n = 3$$

Hazard rate h(x)

$$F(x) = \int_0^x f(y)dy, \qquad h(x) = \frac{f(x)}{1 - F(x)}$$

Example 1
Constant hazard rate

$$h(x) = 1$$

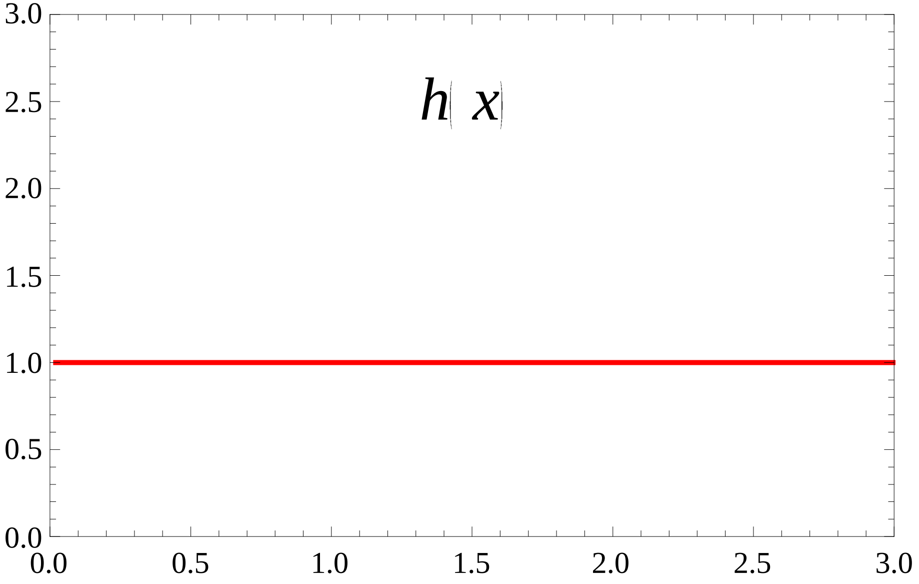

Example 1
Constant hazard rate

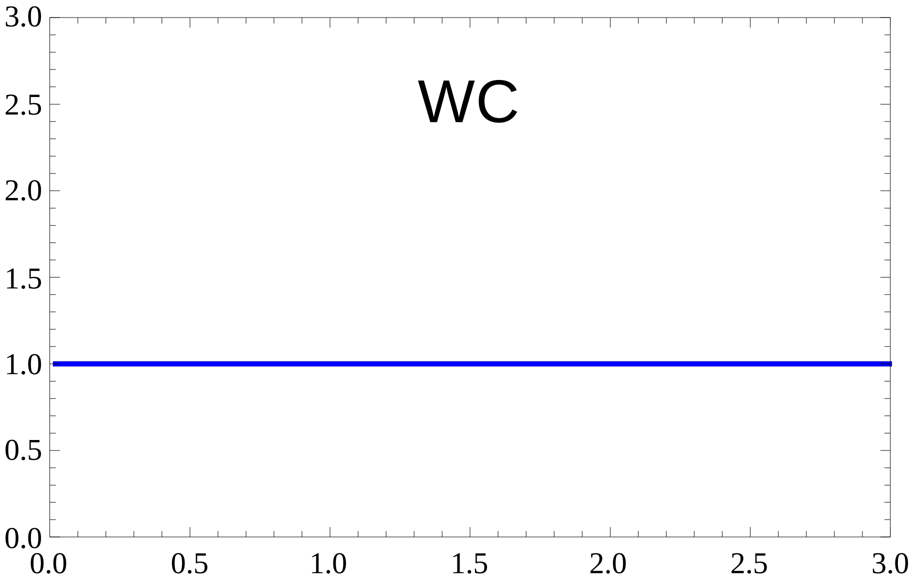

Example 2
Increasing hazard rate

$$h(x) = \begin{cases} 1, & x < 1 \\ 2, & x \geq 1 \end{cases}$$

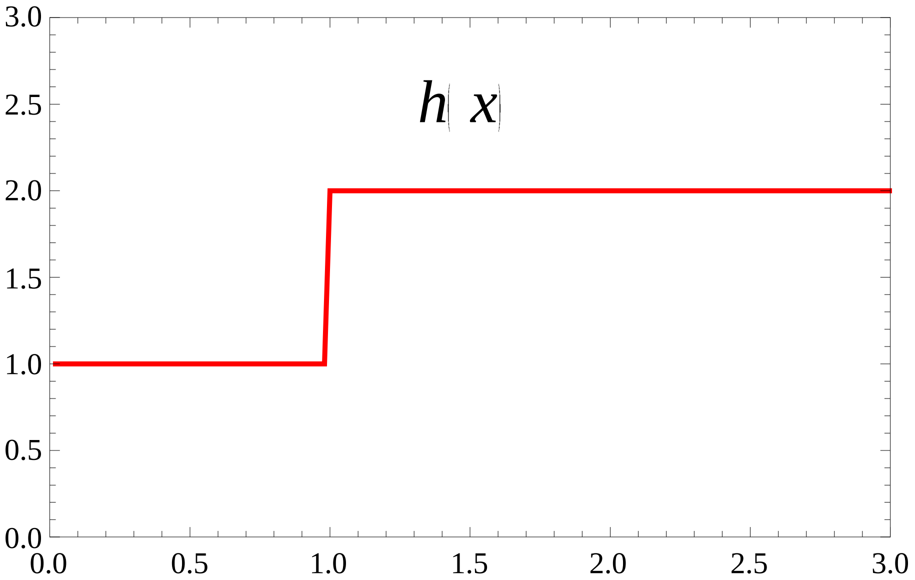

Example 2
Increasing hazard rate

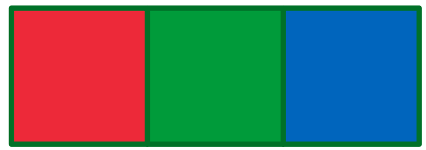

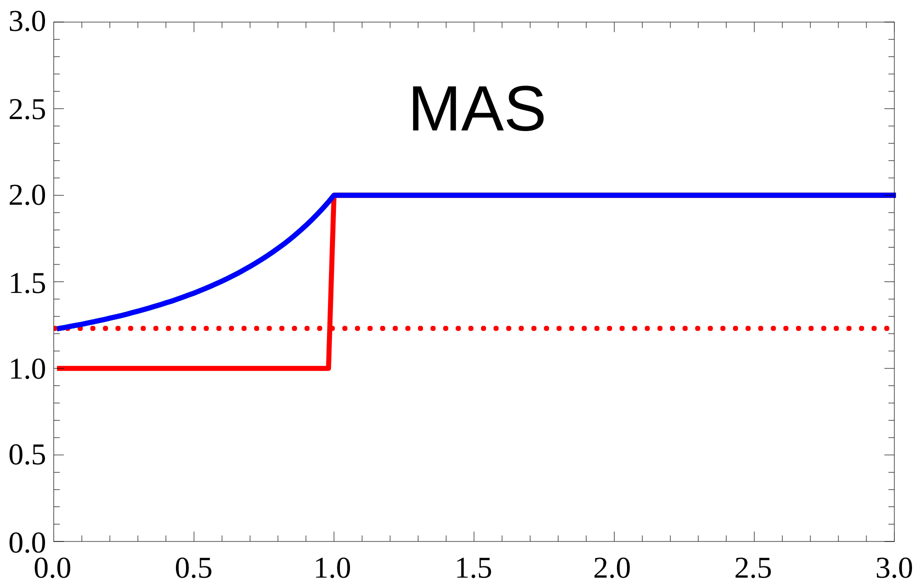

Example 3
Decreasing hazard rate

$$h(x) = \begin{cases} 2, & x < 1 \\ 1, & x \geq 1 \end{cases}$$

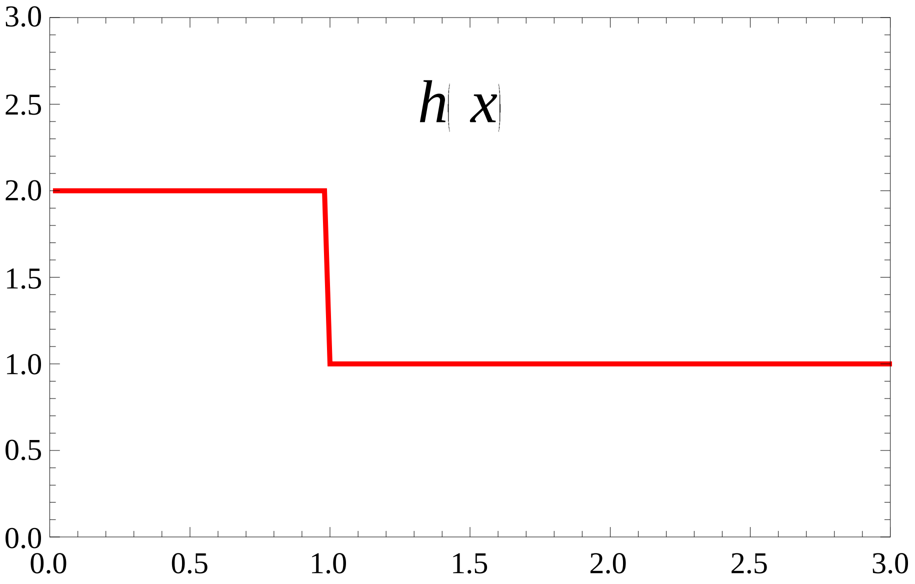

Example 3
Decreasing hazard rate

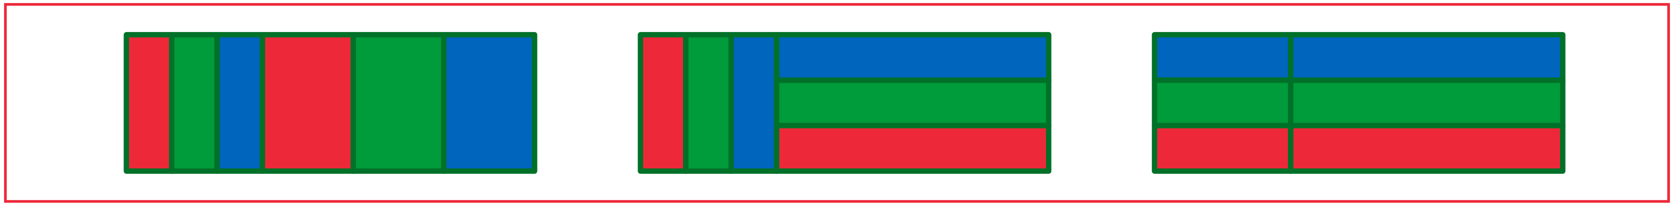

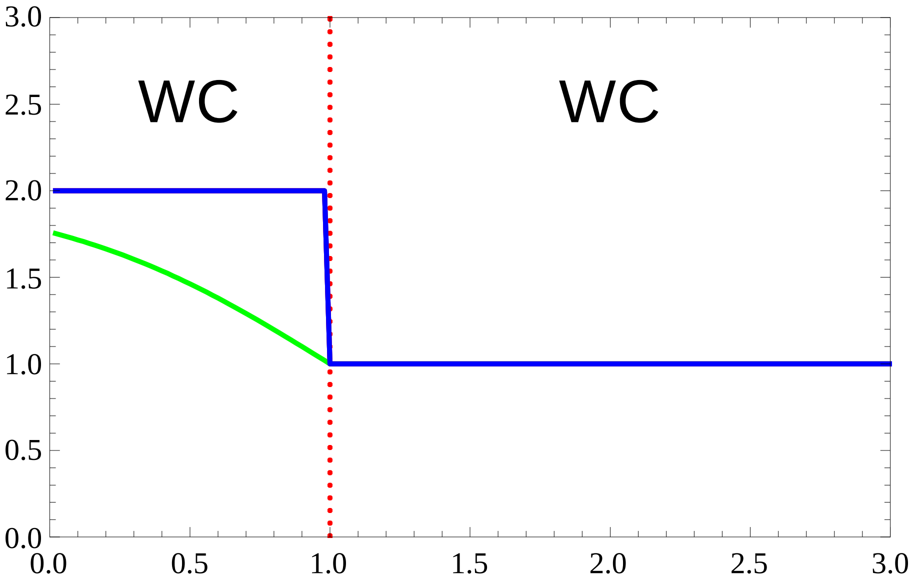

Example 4
Increasing-decreasing hazard rate

$$h(x) = \begin{cases} 1, & x < 1,\ x > 2 \\ 2, & 1 \le x < 2 \end{cases}$$

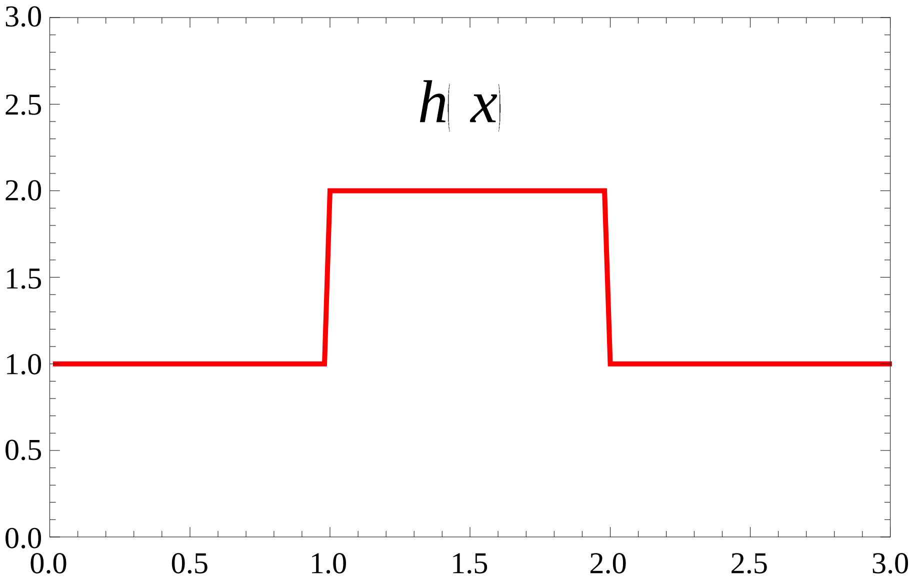

Example 4
Increasing-decreasing hazard rate

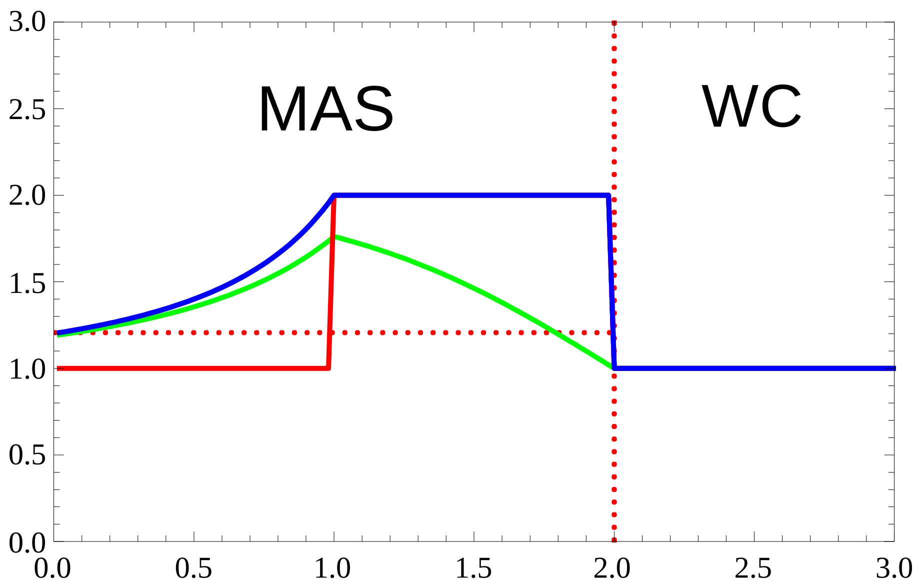

Example 5
Decreasing-increasing hazard rate

$$h(x) = \begin{cases} 2, & x < 1,\ x > 2 \\ 1, & 1 \leq x < 2 \end{cases}$$

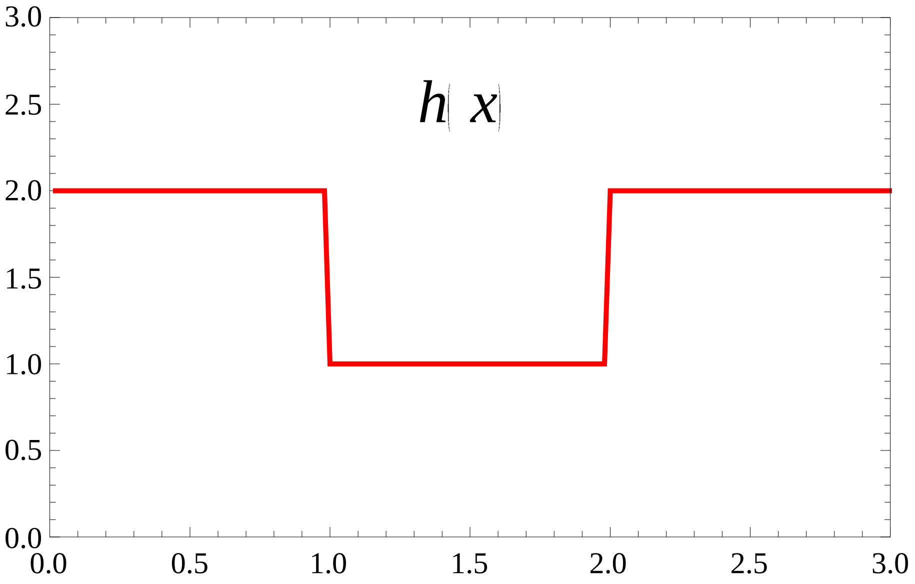

Example 5
Decreasing-increasing hazard rate

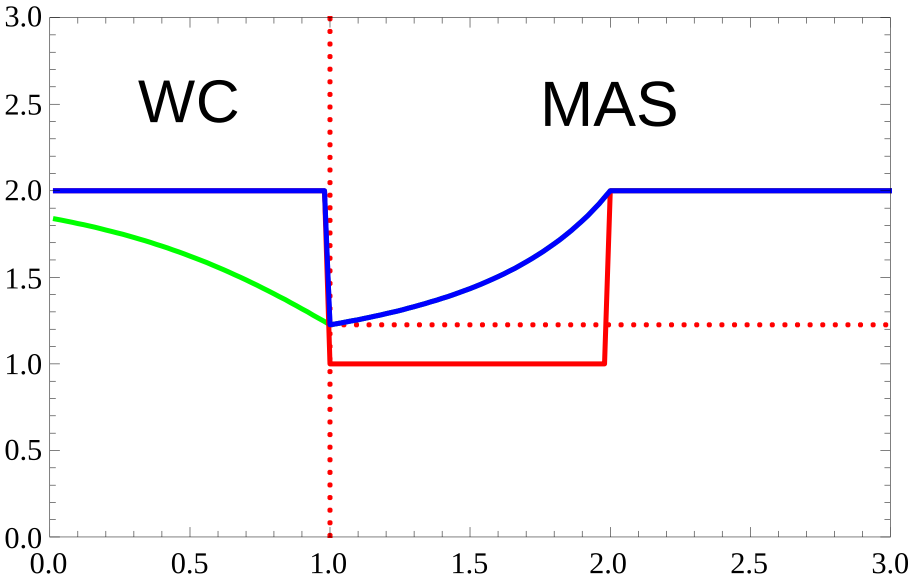

Outline of Part 2

- Introduction
- Gittins index
- Continuity and monotonicity result
- Monotonicity in finite intervals
- Service time distribution classes
- Optimality results
- Summary

Hazard rate

- Remaining service time distribution:

$$P\{S-x\le y\,|\,S>x\}=\frac{F(x+y)-F(x)}{1-F(x)}$$

- Hazard rate (HR) function $h(x)$:

$$h(x)\triangleq \lim_{\Delta\to 0}\frac{1}{\Delta}P\{S-x\le \Delta\,|\,S>x\}=\frac{f(x)}{1-F(x)}$$

Inverse MRL

- Mean residual lifetime (MRL) function $M(x)$:

$$M(x) \triangleq E[S - x \mid S > x] = \frac{\int_x^\infty (1-F(y))dy}{1-F(x)}$$

- Inverse MRL function $H(x)$:

$$H(x) \triangleq \frac{1}{E[S - x \mid S > x]} = \frac{1-F(x)}{\int_x^\infty (1-F(y))dy}$$

Efficiency function

- Efficiency function for age a and service quota Δ:

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

- Limiting values:

$$J(a,0) = h(a), \quad J(a,\infty) = H(a)$$

Gittins index

- Definition: Gittins index $G(\alpha)$ for a job with age α is

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

- Optimal service quota for a job with age α:

$$\Delta^*(a) \triangleq \sup\{\Delta \geq 0 \mid J(a, \Delta) = G(a)\}$$

Example 6
Oscillating hazard rate

$$h(x) = \frac{1 + \sin x}{5}$$

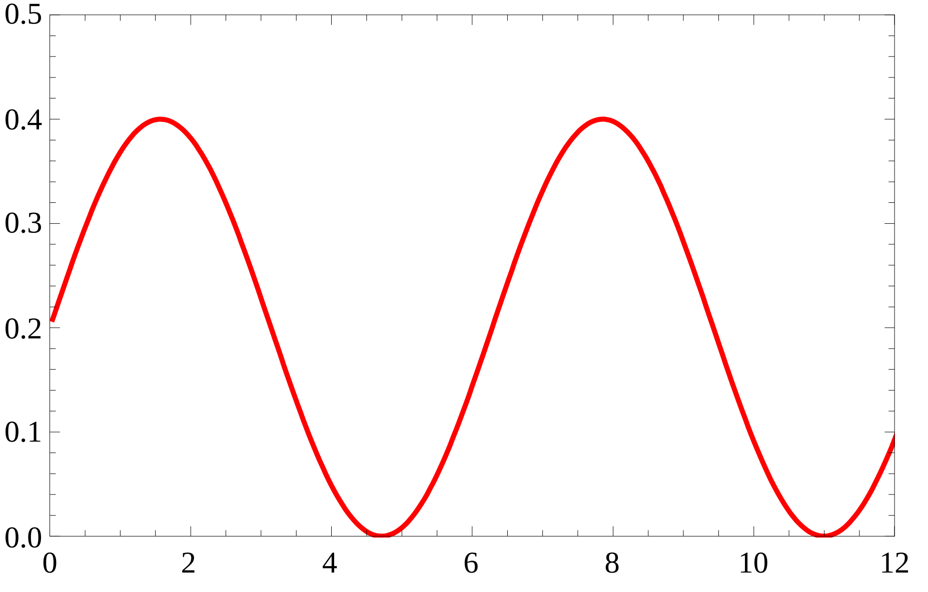

Example 6
Oscillating hazard rate

Gittins index G(x)
inverse MRL H(x)
hazard rate h(x)

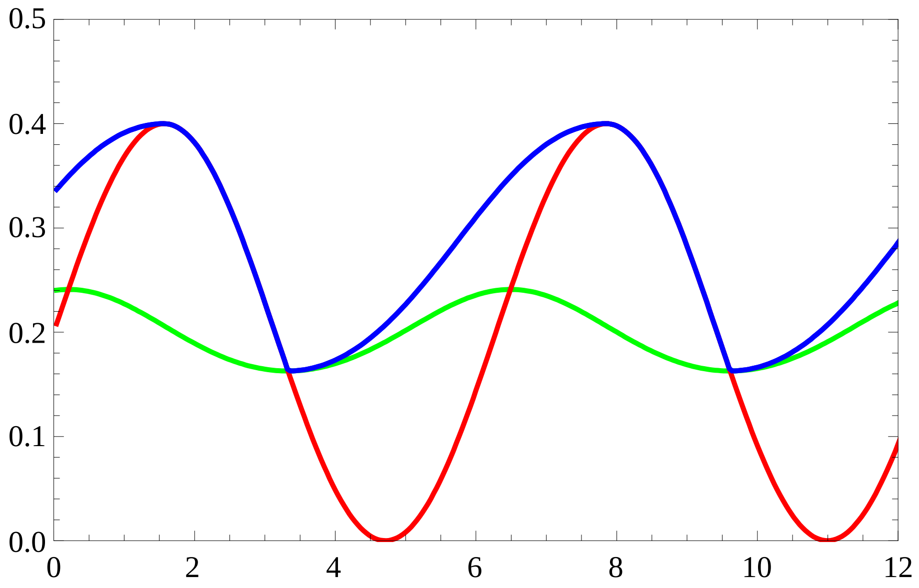

Example 6
Oscillating hazard rate

Gittins index G(x)
inverse MRL H(x)
hazard rate h(x)

NOTE!

h(x) continuous

$\Rightarrow$

G(x) continuous

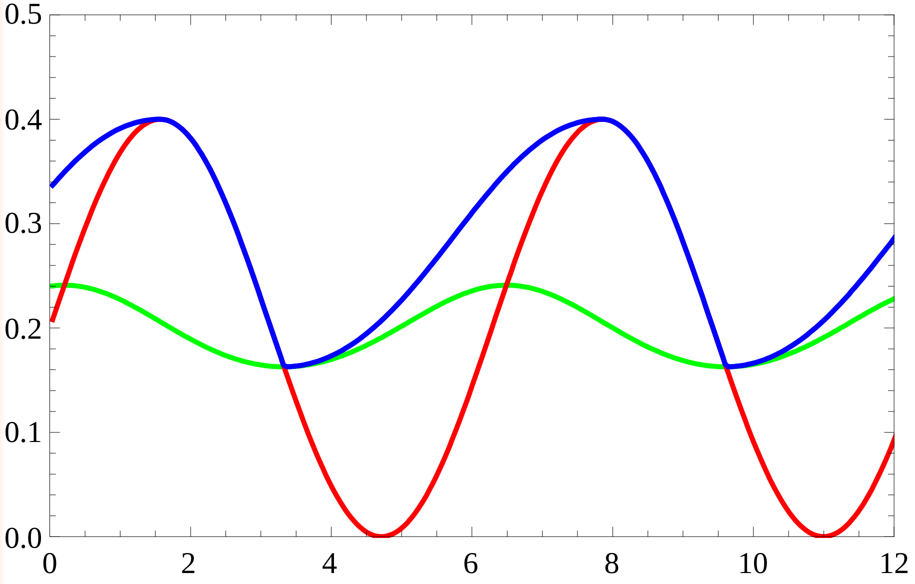

Example 6
Oscillating hazard rate

Gittins index G(x)
inverse MRL H(x)
hazard rate h(x)

NOTE!

G(x) decreasing

$\Rightarrow$

h(x) decreasing
and
G(x) = h(x)

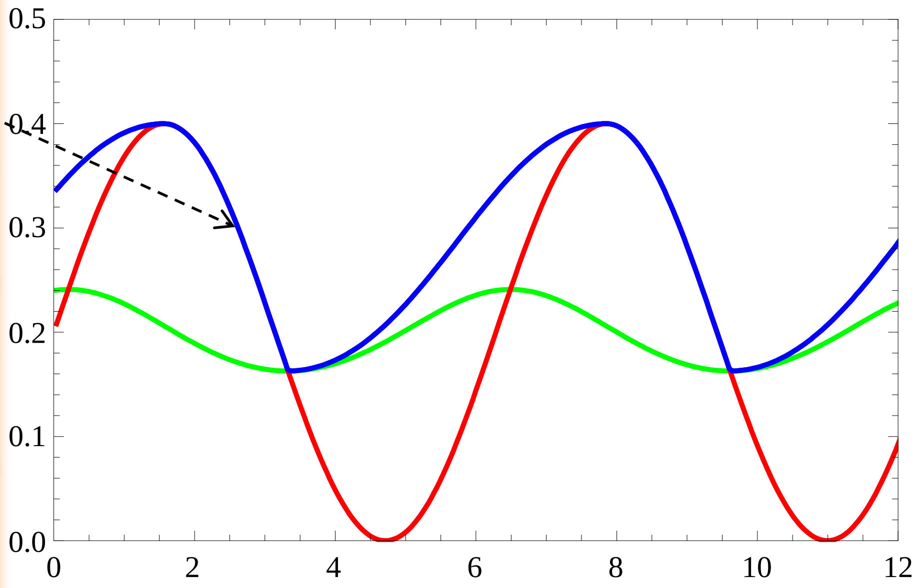

Example 6
Oscillating hazard rate

Gittins index G(x)
inverse MRL H(x)
hazard rate h(x)

NOTE!

h(x) increasing
or
H(x) increasing
$\Rightarrow$
G(x) increasing

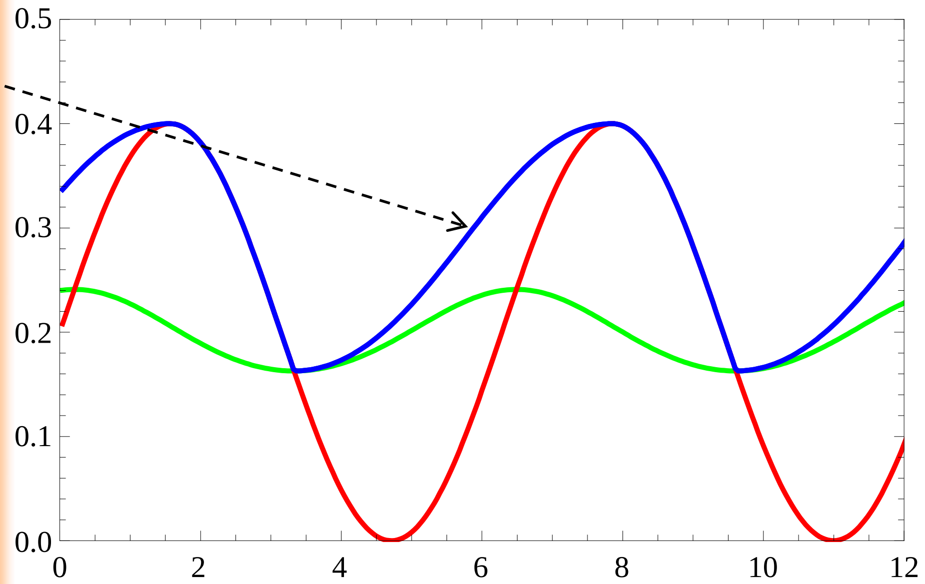

Example 6
Oscillating hazard rate

Gittins index G(x) (rescaled)

optimal service quota $\Delta^*(x)$

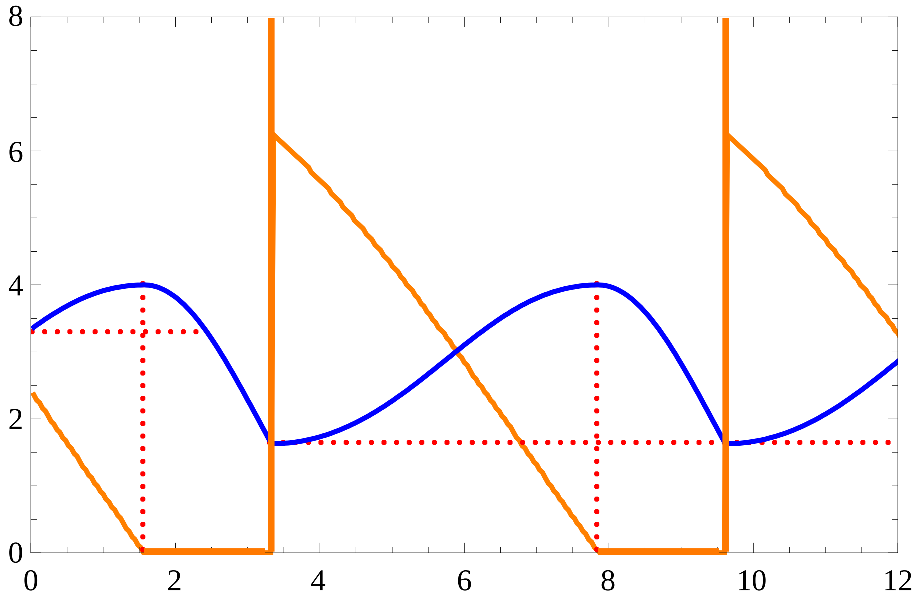

Example 6
Oscillating hazard rate

Gittins index G(x) (rescaled)

optimal service quota $\Delta^*(x)$

NOTE!

Here $\Delta^*(x) = \infty$

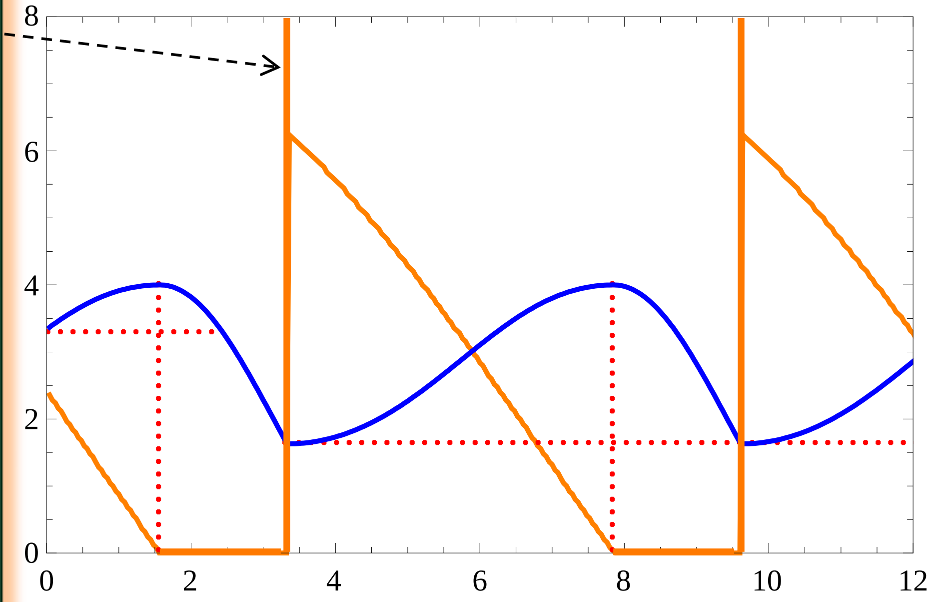

Outline of Part 2

- Introduction
- Gittins index
- Continuity and monotonicity result
- Monotonicity in finite intervals
- Service time distribution classes
- Optimality results
- Summary

Continuity result

- Property:

$$f(x) \text{ is continuous for all } x$$
$$\Leftrightarrow h(x) \text{ is continuous for all } x$$
$$\Leftrightarrow J(x,d) \text{ is continuous for all } x,d$$

- Proposition:

$$h(x) \text{ is continuous for all } x$$
$$\Rightarrow G(x) \text{ is continuous for all } x$$

Monotonicity result 1

- Proposition:

$$h(x)\text{ strictly decreasing for all }x\in(a,b)$$
$$\Rightarrow$$
$$G(x)\text{ strictly decreasing for all }x\in(a,c),$$
$$G(x)\text{ increasing for all }x\in(c,b)$$

Example 6
Oscillating hazard rate

Gittins index G(x)
inverse MRL H(x)
hazard rate h(x)

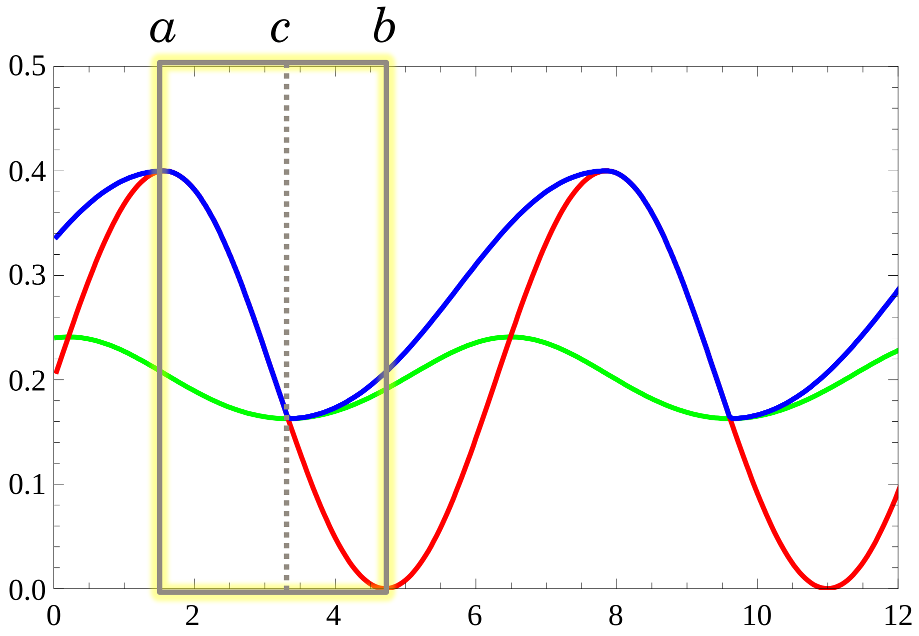

Monotonicity result 2

- Proposition:

$$h(x) \text{ increasing for all } x \in (a,b)$$

$$\Rightarrow$$

$$G(x) \text{ increasing for all } x \in (a,b)$$

Example 6
Oscillating hazard rate

Gittins index G(x)
inverse MRL H(x)
hazard rate h(x)

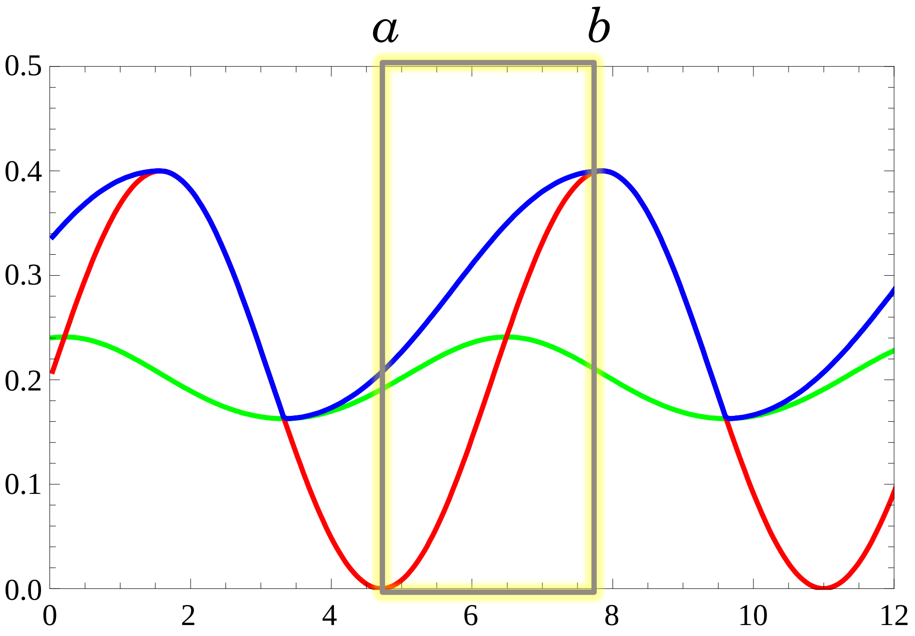

Continuity and monotonicity result

- Summary:

$$h(x) \text{ is continuous and piecewise monotonic for all } x$$

$$\Rightarrow G(x) \text{ is continuous and piecewise monotonic for all } x$$

Example 6
Oscillating hazard rate

Gittins index G(x)
inverse MRL H(x)
hazard rate h(x)

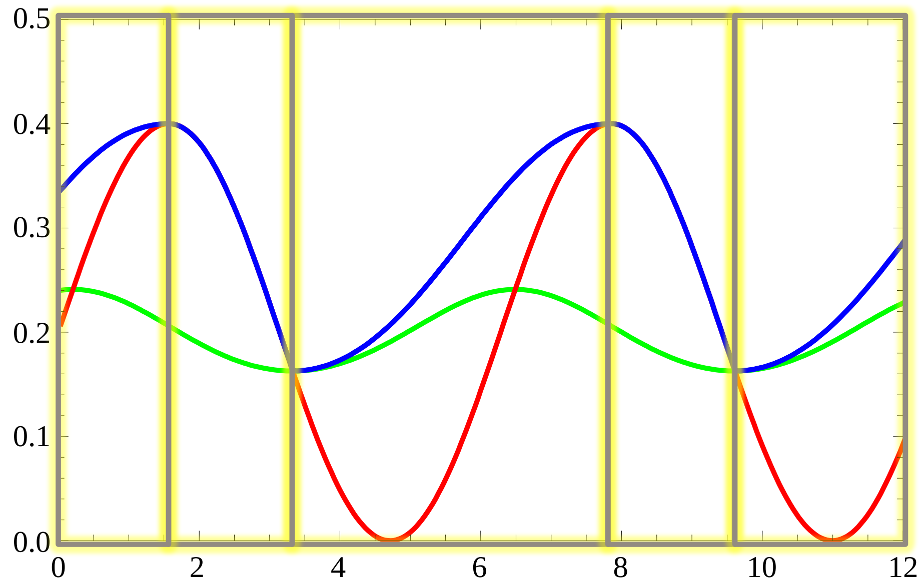

Example 6
Oscillating hazard rate

Gittins index G(x)
inverse MRL H(x)
hazard rate h(x)

NOTE!

Continuity needed here

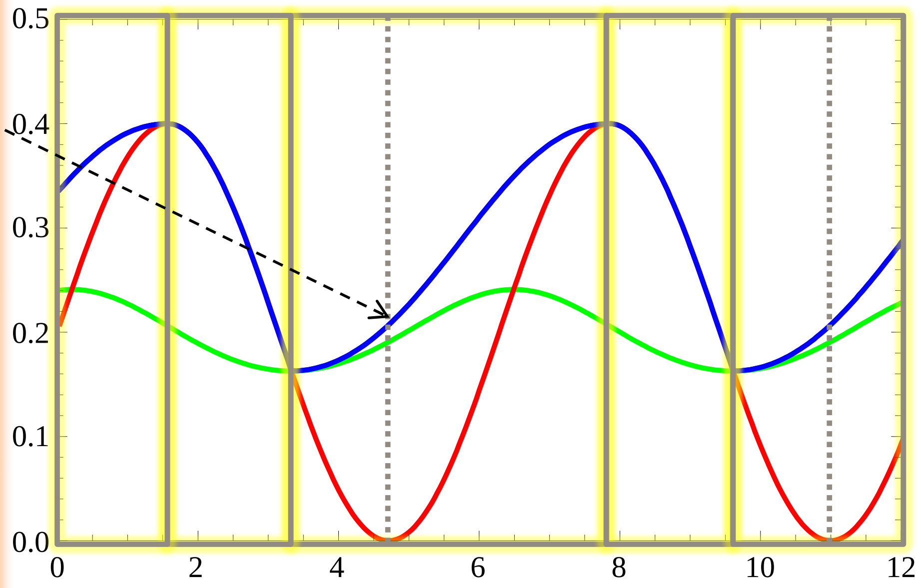

Aalto University
School of Electrical Engineering

Outline of Part 2

- Introduction
- Gittins index
- Continuity and monotonicity result
- Monotonicity in finite intervals
- Service time distribution classes
- Optimality results
- Summary

Monotonicity in finite intervals 1

- Proposition:

$$G(x) \text{ is strictly increasing for all } x \in (a,b)$$

$$\Leftrightarrow$$

$$G(x) > h(x) \text{ for all } x \in (a,b)$$

Example 6
Oscillating hazard rate

Gittins index G(x)
inverse MRL H(x)
hazard rate h(x)

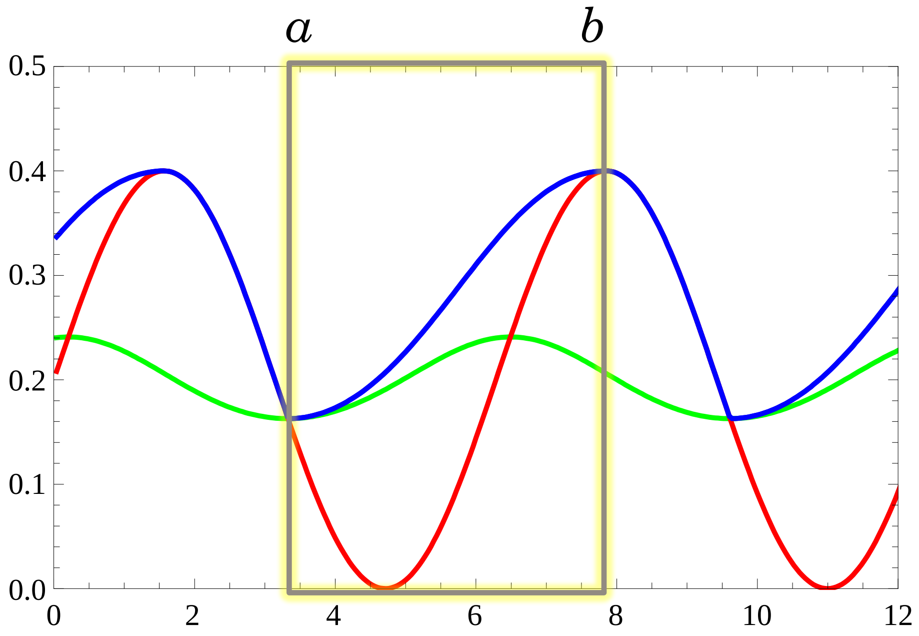

Monotonicity in finite intervals 2

- Proposition:

$$G(x) \text{ is increasing for all } x \in (a,b)$$

$$\Leftrightarrow$$

$$\Delta^{*}(x) > 0 \ \text{ for all } x \in (a,b)$$

Example 6
Oscillating hazard rate

Gittins index G(x) (rescaled)
optimal service quota Δ*(x)

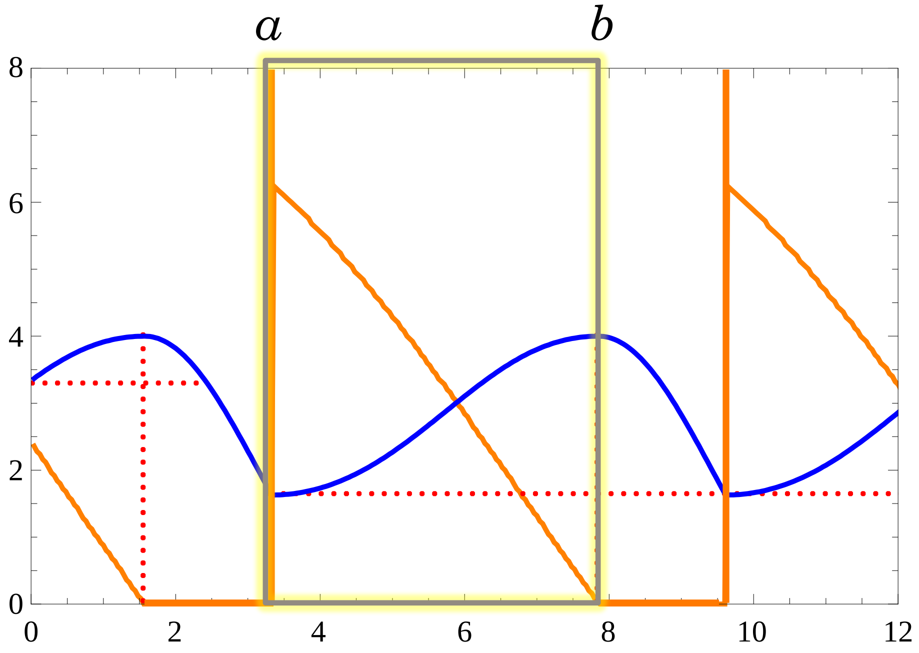

Monotonicity in finite intervals 3

- Proposition:

$$G(x) \text{ is constant for all } x \in (a,b)$$

$$\Leftrightarrow$$

$$G(x) = h(x) \text{ and } \Delta^*(x) > 0 \text{ for all } x \in (a,b)$$

Example 3
Decreasing hazard rate

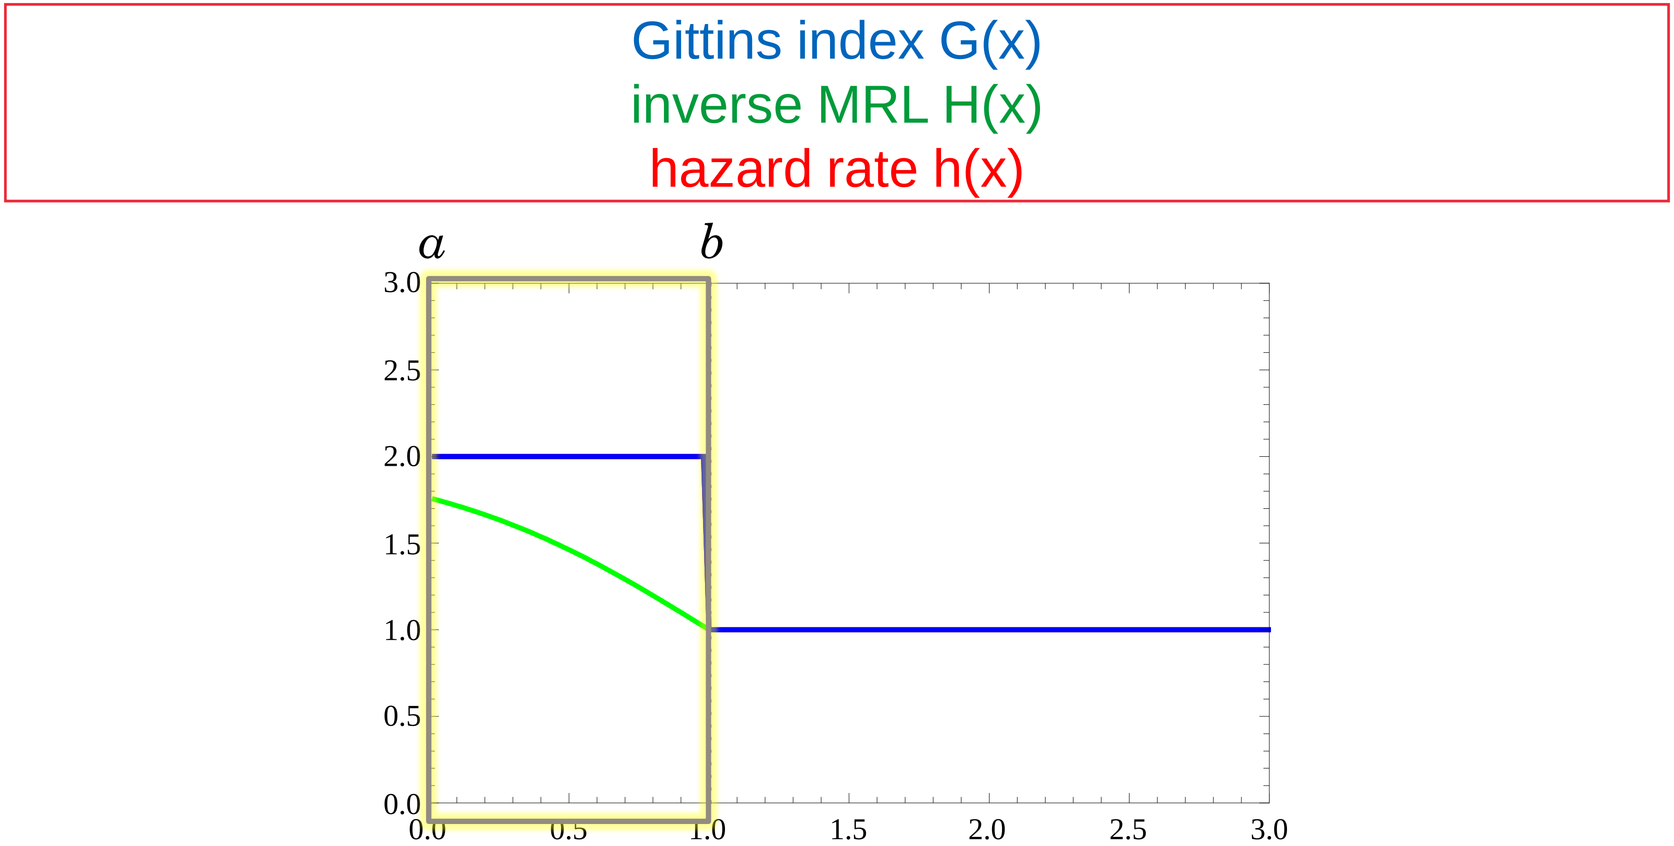

Monotonicity in finite intervals 4

- Proposition:

$$G(x) \text{ is decreasing for all } x \in (a,b)$$

$$\Leftrightarrow$$

$$G(x) = h(x) \text{ for all } x \in (a,b)$$

Example 6
Oscillating hazard rate

Gittins index G(x)
inverse MRL H(x)
hazard rate h(x)

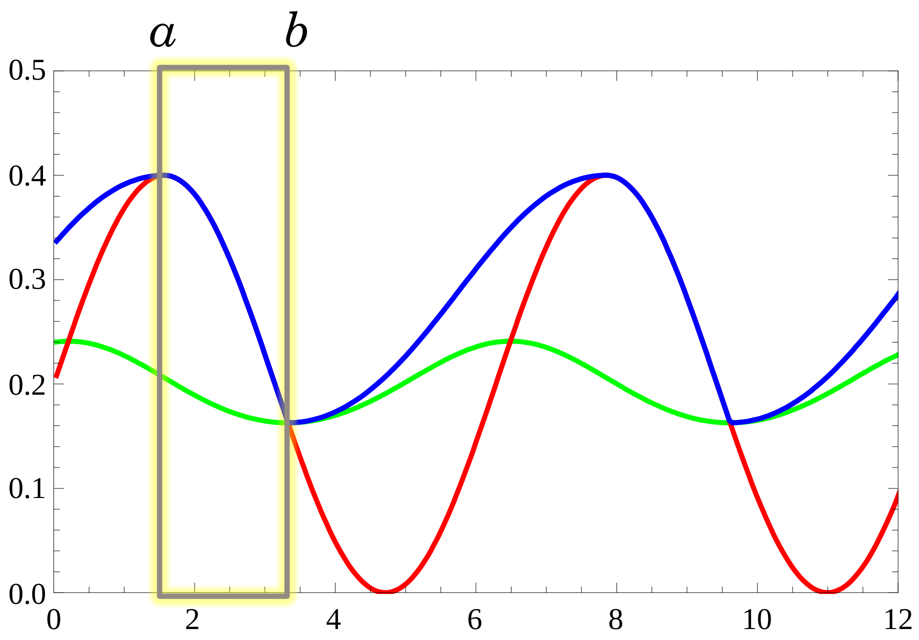

Monotonicity in finite intervals 5

- Proposition:

$$G(x) \text{ is strictly decreasing for all } x \in (a,b)$$

$$\Leftrightarrow$$

$$\Delta^{*}(x) = 0 \ \text{ for all } x \in (a,b)$$

Example 6
Oscillating hazard rate

Gittins index G(x) (rescaled)
optimal service quota $\Delta^{*}(x)$

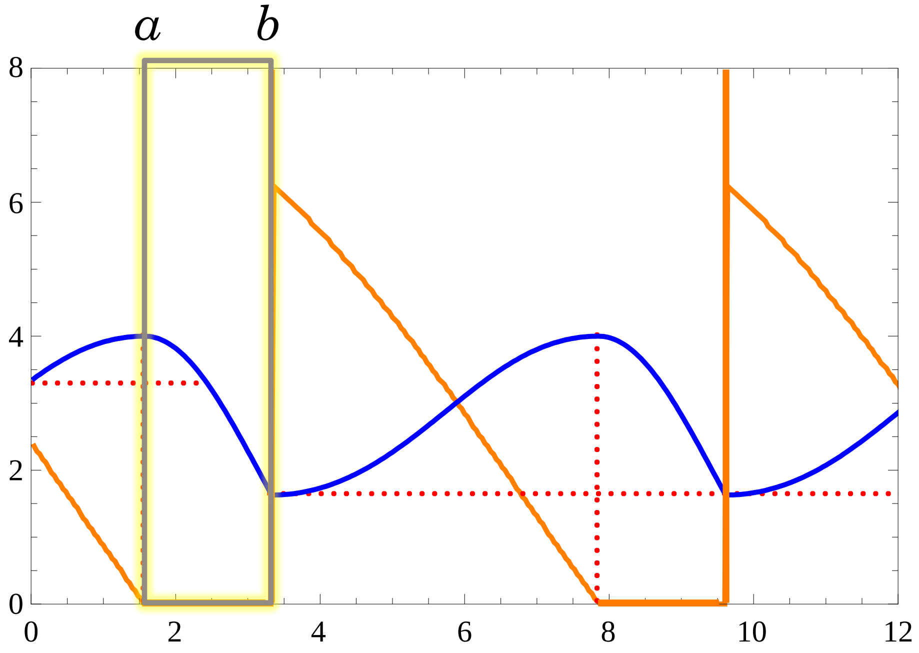

Outline of Part 2

- Introduction
- Gittins index
- Continuity and monotonicity result
- Monotonicity in finite intervals
- Service time distribution classes
- Optimality results
- Summary

Service time distribution classes

- IHR = Increasing Hazard Rate
- DMRL = Decreasing Mean Residual Lifetime
- NBUE = New Better than Used in Expectation

- DHR = Decreasing Hazard Rate
- IMRL = Increasing Mean Residual Lifetime
- NWUE = New Worse than Used in Expectation

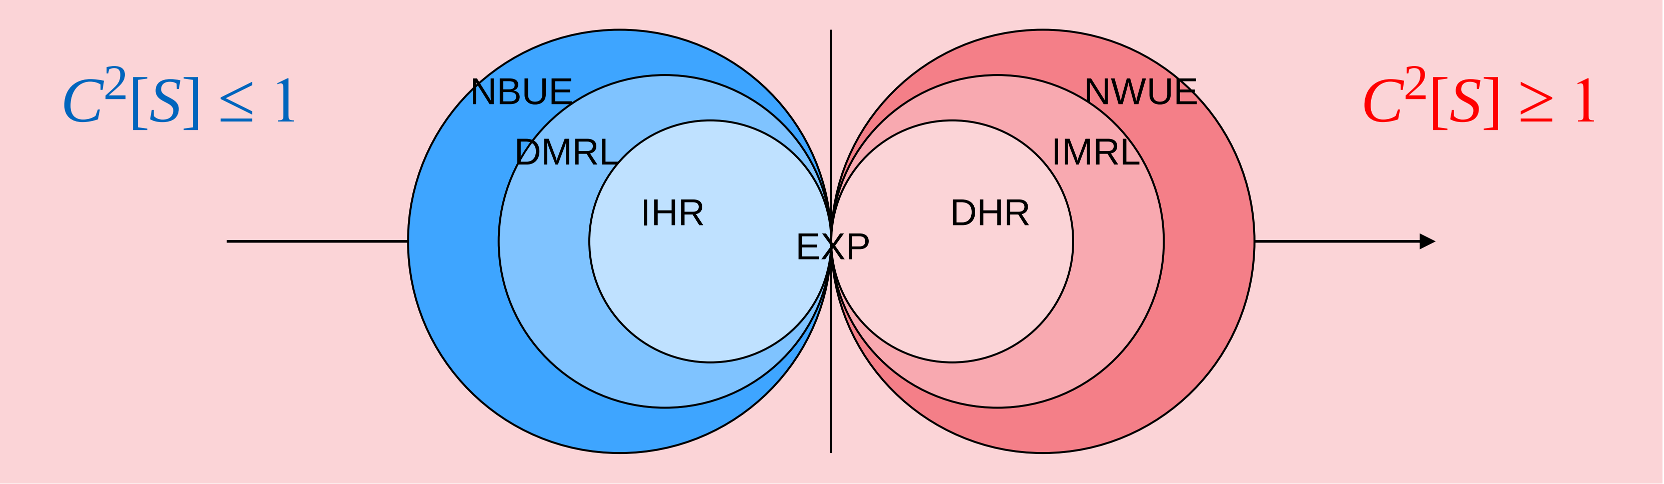

Properties in infinite intervals 1

- Proposition:

$$G(x) \geq G(a) \text{ for all } x \in (a,\infty)$$

$$\Leftrightarrow$$

$$H(x) \geq H(a) \text{ for all } x \in (a,\infty)$$

$$\Leftrightarrow$$

$$G(a) = H(a)$$

Example 6
Oscillating hazard rate

Gittins index G(x)
inverse MRL H(x)
hazard rate h(x)

NOTE!

G(a) = H(a)

G(x) ≥ G(a)

H(x) ≥ H(a)

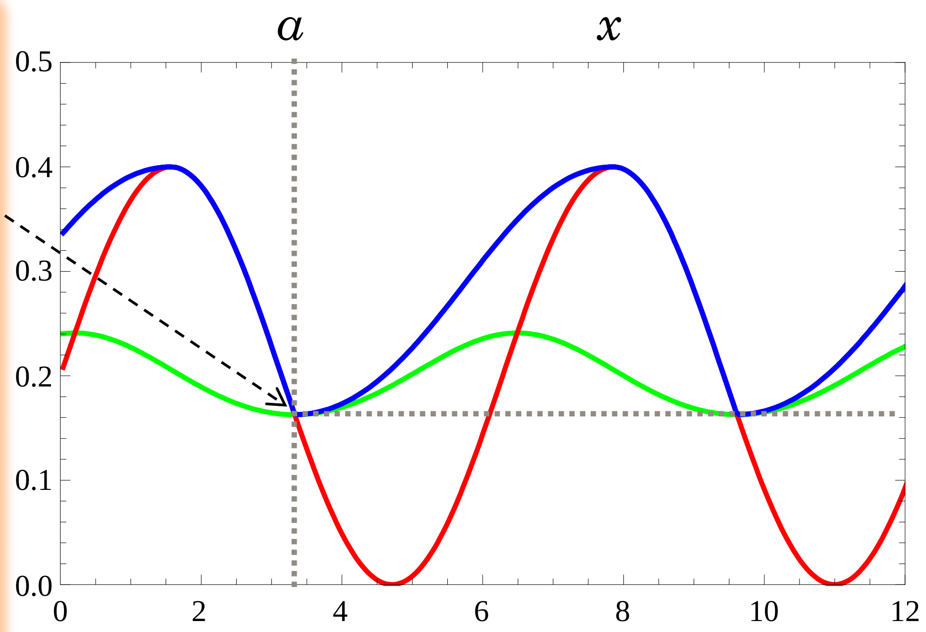

NBUE service times

- Corollary:

$$G(x) \geq G(0) \text{ for all } x$$

$$\Leftrightarrow$$

Service times are NBUE

$$\Leftrightarrow$$

$$G(0) = H(0)$$

Properties in infinite intervals 2

- Proposition:

$$G(x) \text{ is increasing for all } x \in (a, \infty)$$

$$\Leftrightarrow$$

$$H(x) \text{ is increasing for all } x \in (a, \infty)$$

$$\Leftrightarrow$$

$$G(x) = H(x) \quad \text{for all } x \in (a, \infty)$$

Example 5
Decreasing-increasing hazard rate

Gittins index G(x)
inverse MRL H(x)
hazard rate h(x)

NOTE!

G(x) = H(x)

for all x > a

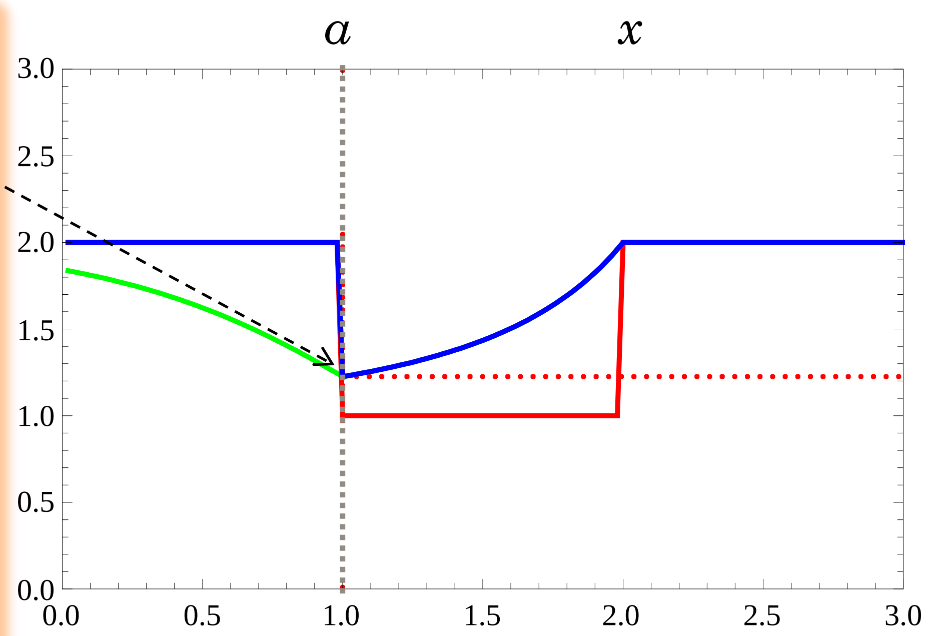

DMRL service times

- Corollary:

$$G(x) \text{ is increasing for all } x$$

$$\Leftrightarrow$$

$$\text{Service times are DMRL}$$

$$\Leftrightarrow$$

$$G(x) = H(x) \text{ for all } x$$

Properties in infinite intervals 3

- Proposition:

$$G(x) \text{ is constant for all } x \in (a,\infty)$$

$$\Leftrightarrow$$

$$H(x) \text{ is constant for all } x \in (a,\infty)$$

$$\Leftrightarrow$$

$$h(x) \text{ is constant for all } x \in (a,\infty)$$

$$\Leftrightarrow$$

$$G(x) = H(x) = h(x) \text{ for all } x \in (a,\infty)$$

EXP service times

- Corollary:

$$G(x) \text{ is constant for all } x$$

$$\Leftrightarrow$$

$$\text{Service times are EXP}$$

$$\Leftrightarrow$$

$$G(x) = H(x) = h(x) \quad \text{for all } x$$

Properties in infinite intervals 4

- Proposition:

$$G(x) \text{ is decreasing for all } x \in (a,\infty)$$

$$\Leftrightarrow$$

$$h(x) \text{ is decreasing for all } x \in (a,\infty)$$

$$\Leftrightarrow$$

$$G(x) = h(x) \;\text{ for all } x \in (a,\infty)$$

Example 4
Increasing-decreasing hazard rate

Gittins index G(x)
inverse MRL H(x)
hazard rate h(x)

NOTE!

G(x) = h(x)

for all x > a

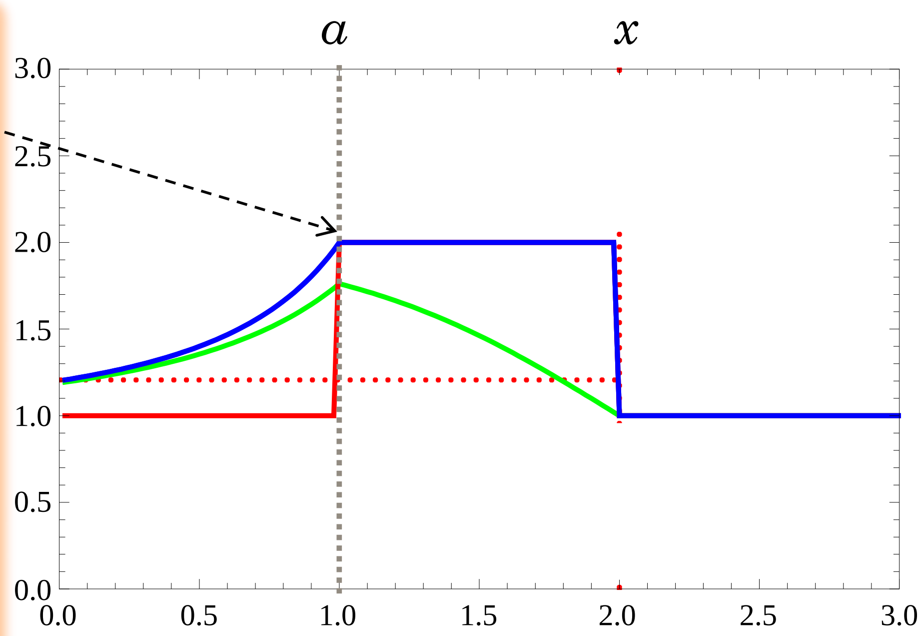

DHR service times

- Corollary:

$$G(x) \text{ is decreasing for all } x$$

$$\Leftrightarrow$$

$$\text{Service times are DHR}$$

$$\Leftrightarrow$$

$$G(x) = h(x) \ \text{ for all } x$$

Outline of Part 2

- Introduction
- Gittins index
- Continuity and monotonicity result
- Monotonicity in finite intervals
- Service time distribution classes
- Optimality results
- Summary

Optimality of the MAS discipline

- Corollary:

$$\text{MAS is optimal} \iff \text{Service times are NBUE}$$

- Note: In this case MAS = SERPT (due to NBUE)

Example 2
Increasing hazard rate

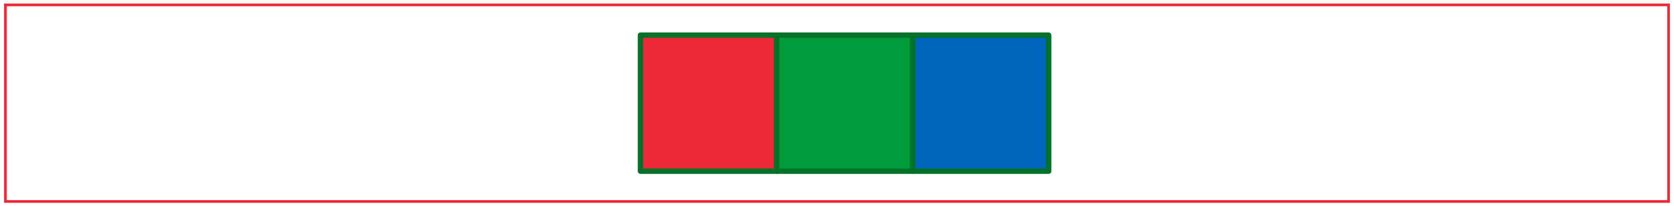

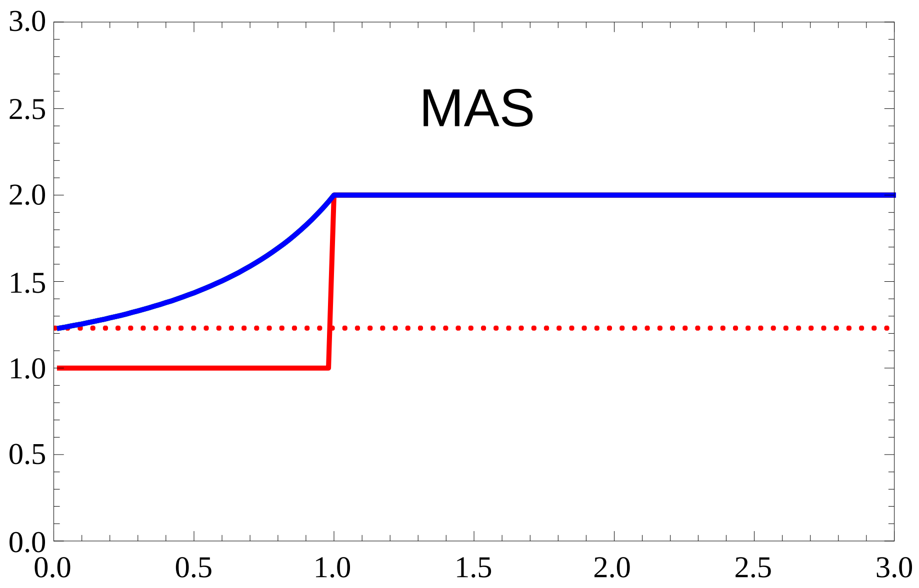

Optimality of the LAS discipline

- Corollary:

$$\text{LAS is optimal} \Leftrightarrow \text{Service times are DHR}$$

- Note: In this case LAS = SERPT (since DHR $\Rightarrow$ IMRL)

Example 3
Decreasing hazard rate

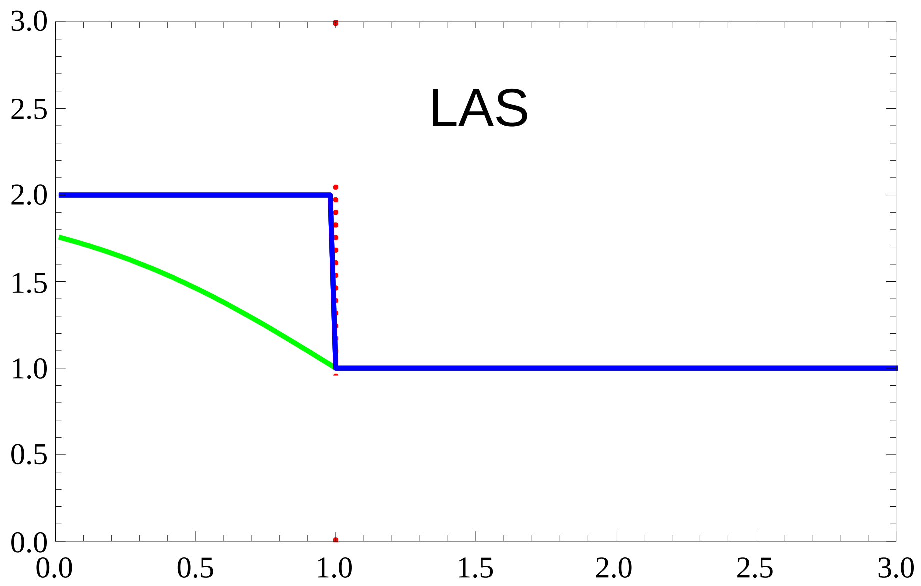

Optimality of the MAS+LAS discipline

- Corollary:

$$\text{Service times are NBUE} + \text{DHR}(k)$$

$$\Rightarrow$$

$$\text{MAS} + \text{LAS}(k^*) \text{ is optimal}$$

- MAS+LAS belongs to MLPS (Multi-Level Processor Sharing) policies, cf. Kleinrock (1976)

Example 4
Increasing-decreasing hazard rate

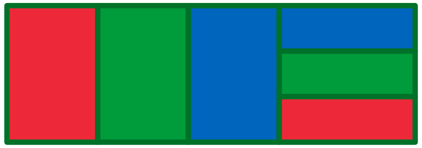

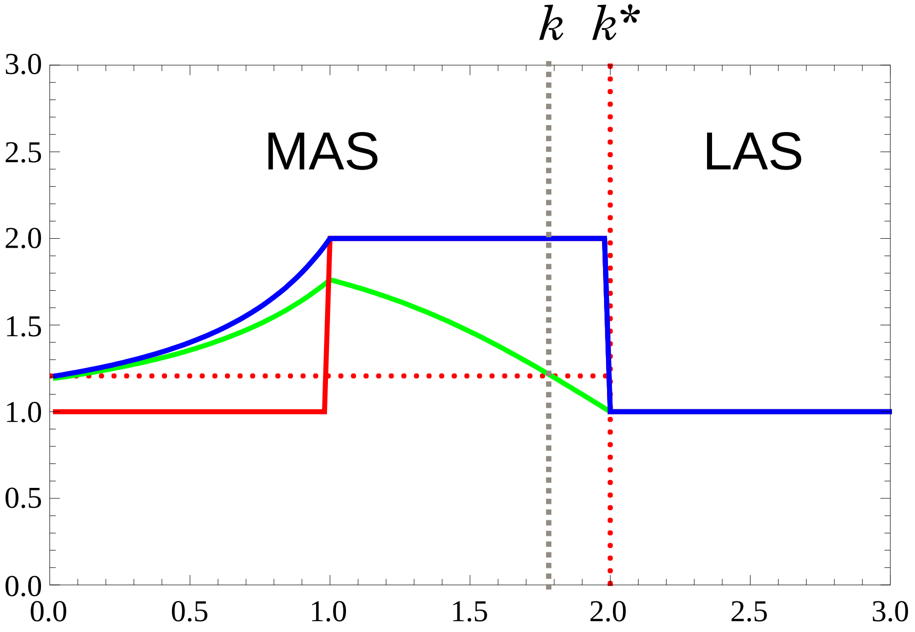

Optimality of the LAS+MAS discipline

- Corollary:

$$\text{Service times are DHR} + \text{IHR}(k),$$
$$h(0) \geq H(\infty)$$
$$\Rightarrow$$
$$\text{LAS} + \text{MAS}(k^*) \text{ is optimal}$$

- LAS+MAS belongs to MLPS (Multi-Level Processor Sharing) policies, cf. Kleinrock (1976)

Example 5
Decreasing-increasing hazard rate

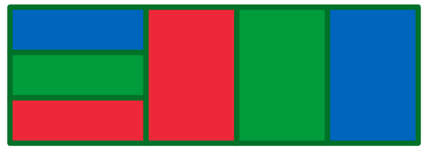

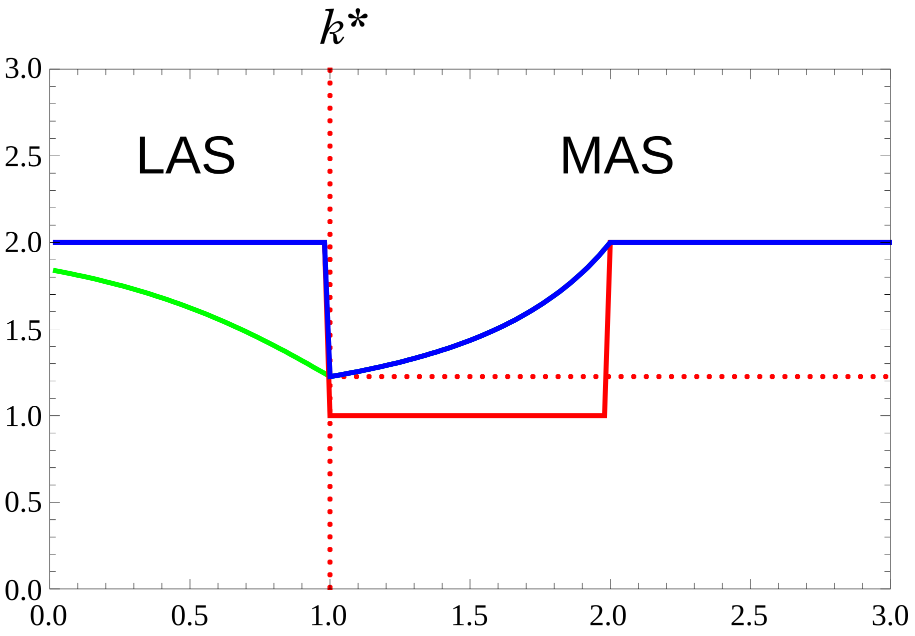

Transient system 1

- Assume: $h(x)$ continuous and piecewise monotonic
- Corollary:

$$\text{Hazard rate } h(x) \text{ is first increasing}$$
$$\Rightarrow$$
$$\text{MAS} + \text{LAS} + \text{MAS} + \ldots (k_1^*, k_2^*, \ldots) \text{ is optimal}$$
$$\text{for the transient system}$$

- MAS+LAS+MAS+… belongs to MLPS (Multi-Level Processor Sharing) policies, cf. Kleinrock (1976)

Example 6
Oscillating hazard rate

$$h(x) = \frac{1 + \sin x}{5}$$

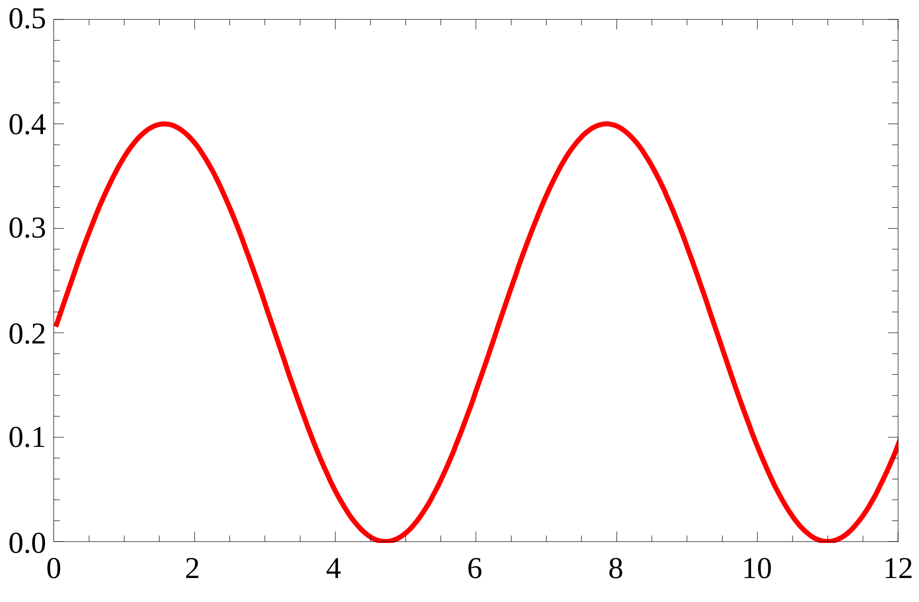

Example 6
Oscillating hazard rate

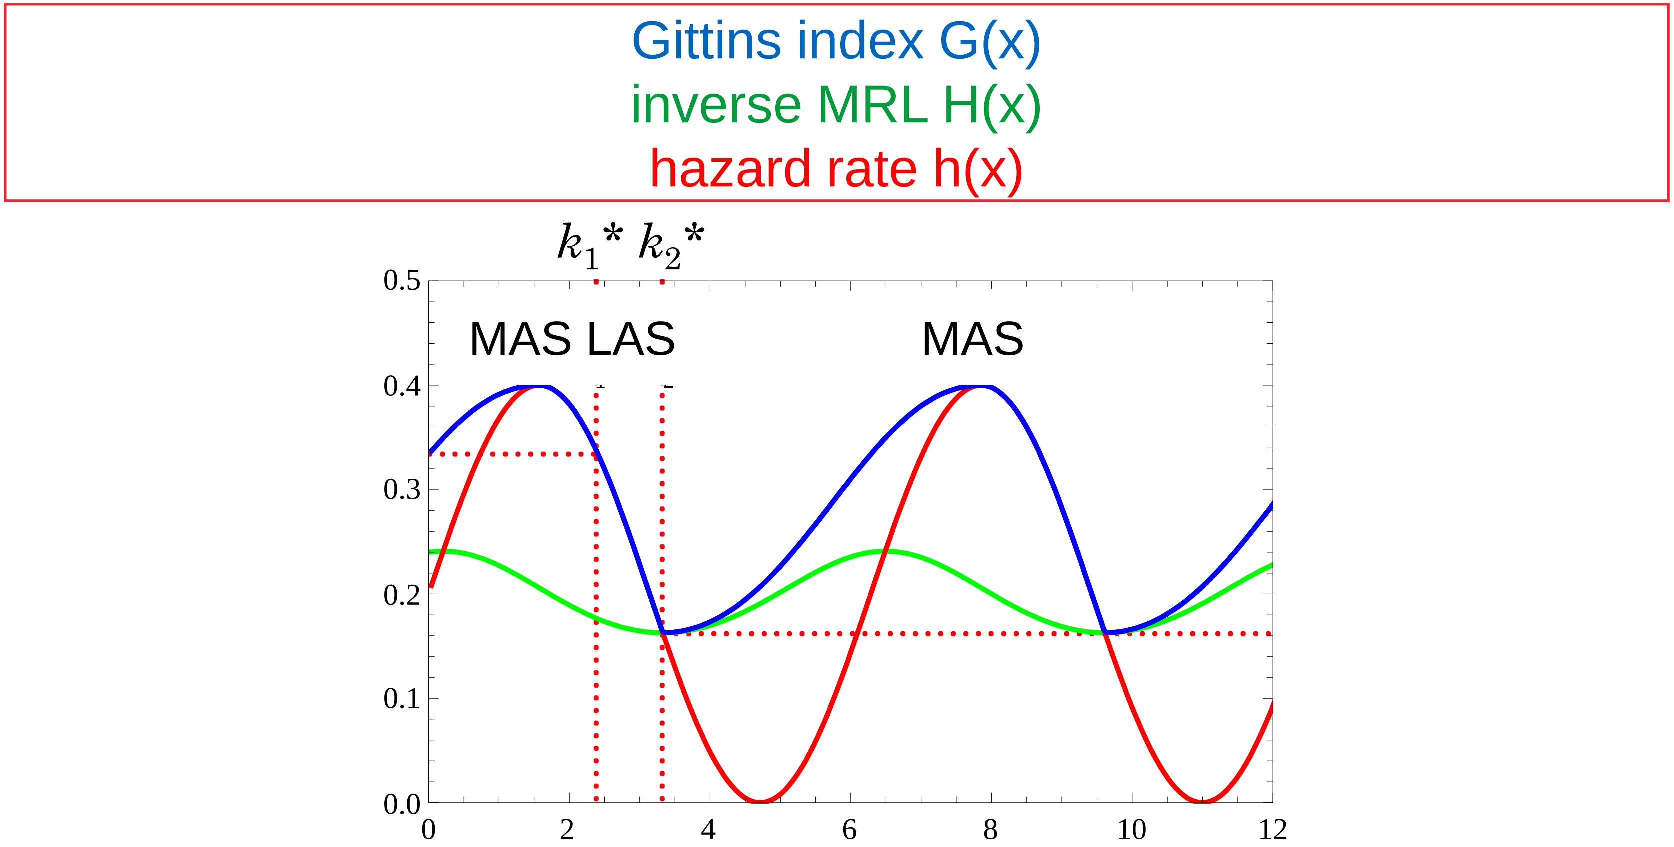

Transient system 2

- Assume: $h(x)$ continuous and piecewise monotonic
- Corollary:

$$\text{Hazard rate } h(x) \text{ is first decreasing}$$
$$\Rightarrow$$
$$\text{LAS} + \text{MAS} + \text{LAS} + \ldots(k_1{}^*, k_2{}^*, \ldots) \text{ is optimal}$$
$$\text{for the transient system}$$

- LAS+MAS+LAS+… belongs to MLPS (Multi-Level Processor Sharing) policies, cf. Kleinrock (1976)

Example 7
Oscillating hazard rate

$$h(x) = \frac{1 - \sin x}{5}$$

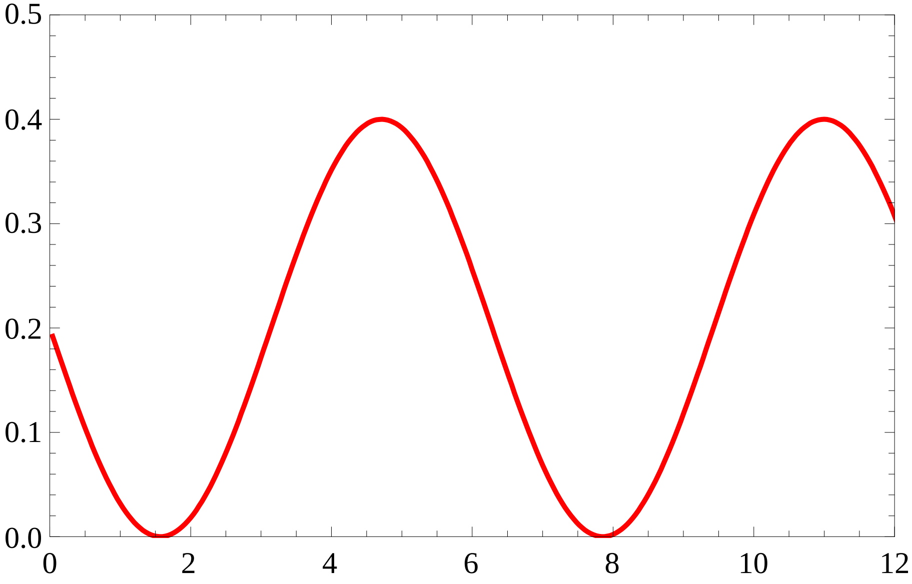

Example 7
Oscillating hazard rate

Gittins index G(x)
inverse MRL H(x)
hazard rate h(x)

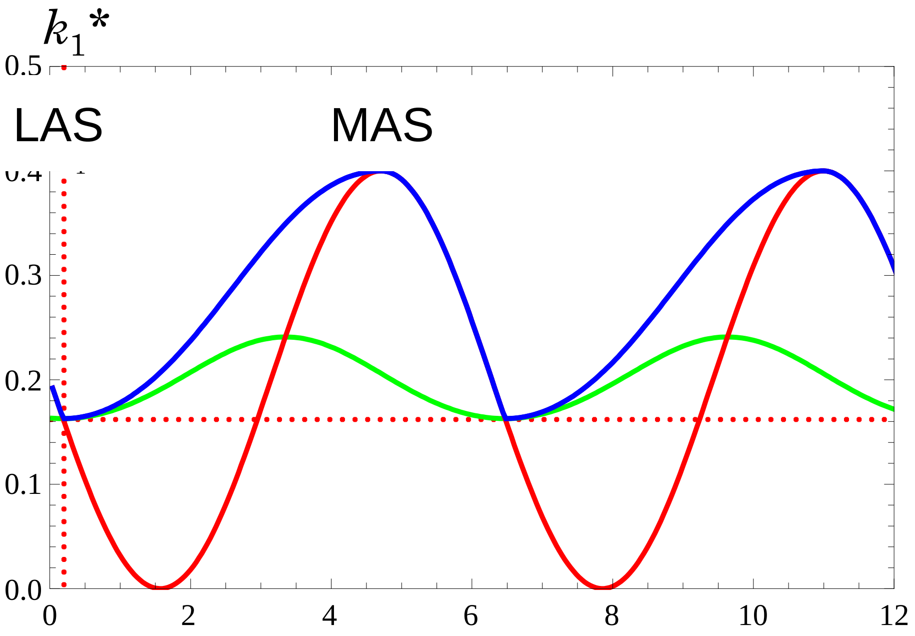

Outline of Part 2

- Introduction
- Gittins index
- Continuity and monotonicity result
- Monotonicity in finite intervals
- Service time distribution classes
- Optimality results
- Summary

Summary of Part 2

Mean delay minimization in M/G/1

NA:

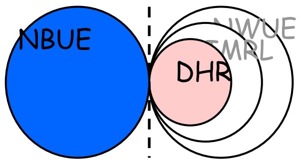

GI optimal

NBUE ⇔ MAS optimal
DHR ⇔ LAS optimal
NBUE+DHR ⇒ MAS+LAS optimal
DHR+IHR ⇒ LAS+MAS optimal

Related contributions

- S. Aalto, U. Ayesta and R. Righter,
  On the Gittins index in the M/G/1 queue,
  Queueing Systems, 2009

- S. Aalto, U. Ayesta and R. Righter,
  Properties of the Gittins index with application,
  Probability in the Engineering and Informational Sciences, 2011

Part 3
Trade-off between size-based and opportunistic scheduling

Outline of Part 3

- Introduction
- Time-scale separation
- Optimal flow-level operating policy
- Examples
- Optimal time-slot-level scheduler
- Summary

Research problem

- Downlink data transmission in a wireless cellular system
- Traffic = elastic flows
  - file transfers using TCP
- Scheduling decisions in each time slot
  - time scale of milliseconds
- Traffic dynamics in a much longer time scale
  - time scale of seconds/minutes
- Optimal time-slot-level scheduler for flow-level performance?

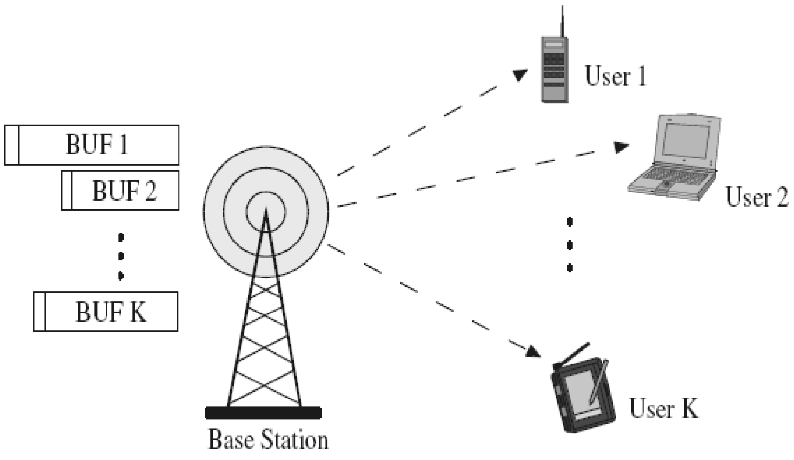

Source: Hu et al. (2004)

Flow-level performance

- Performance is expressed as flow-level delay
  - Mean flow delay describes
    how long, on the average, it takes to transfer a file

- Importance of the time scale
  - Users do not care about time-slot or packet-level delays,
    but the flow-level delay, i.e., the total time to transfer a file

- Flow-level models try to characterize the system at the time scale where users experience the performance

Time-slot-level schedulers

- Channel-aware schedulers
  - Channel conditions varying randomly for each user
  - Scheduling based on channel information
  - Scheduler may prefer users with a good channel
  - Opportunistic scheduling
  - Examples: MR, PF

- Size-based schedulers
  - Scheduling based on flow size information
  - Scheduler may prefer users with a short flow
  - Example: SRPT

Fundamental trade-off

- Opportunistic scheduling
  - Aggregate mean service rate increases with the number of users (opportunistic gain, multiuser diversity gain)
  - However, a user with a long remaining service requirement blocks the other users
- SRPT
  - The number of flows is reduced efficiently
  - However, opportunistic gain is lost due to suboptimal channel

Combining opportunistic and size-based scheduling

- Tsybakov (2003)
  - Dynamic programming approach (time-slot scale)
- Hu et al. (2004)
  - Heuristic approach: TAOS (time-slot scale)
- Lassila and Aalto (2008)
  - Another heuristic approach: SRPT-P (time-slot scale)
- Ayesta et al. (2010), Jacko (2011)
  - Age-based information, Markovian system (time-slot scale)
- Sadiq and de Veciana (2010)
  - Time-scale separation (flow scale)
  - Transient system
  - Optimality result for nested polymatroids
  - Cf. optimality of SRPT-FM, Raj et al. (2004)

SRPT-FM

- SRPT-FM = Shortest Remaining Processing Time on the Fastest Machine
- Pinedo (1995)
- Theorem: SRPT-FM minimizes the mean delay in heterogeneous parallel server systems for a batch of jobs (without any new arrivals)

Outline of Part 3

- Introduction
- Time-scale separation
- Optimal flow-level operating policy
- Examples
- Optimal time-slot-level scheduler
- Summary

Time-scale separation

- $R(t) = (R_1(t), \ldots, R_k(t))$ = rate vector in time slot t
- $R_i(t)$ = instantaneous rate of user i
- Assume: $R_i(t)$ is a stationary and ergodic process
- Assume: Scheduling policy $\pi \in \Pi_k$ is stationary
- Define: The long-term throughput for user i:

$$\theta_i^\pi = \sum_{\mathbf{r}} r_i p_i^\pi(\mathbf{r}) P\{R(t) = \mathbf{r}\}$$

- Define: The (opportunistic) capacity region:

$$C_k = \{(\theta_1^\pi, \ldots, \theta_k^\pi) \in \Re_+^k : \pi \in \Pi_k\}$$

Model

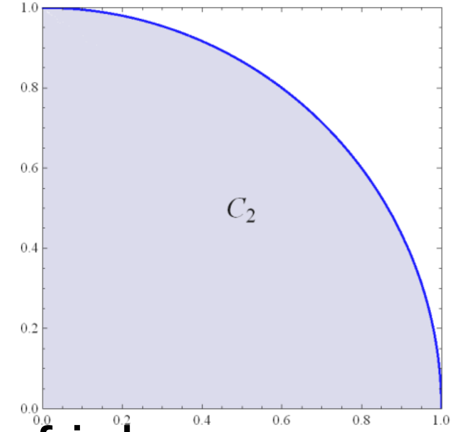

- Service system where the service capacity is adjustable depending on the current number of jobs
- When there are k jobs with sizes

$$s_1 \geq \ldots \geq s_k$$

choose a rate vector

$$\mathbf{c}_k = (c_{k1}, \ldots, c_{kk}) \in C_k$$

and serve job i with rate c_{ki}

- Assume: Capacity regions C_k compact and symmetric

Example: Alpha-ball

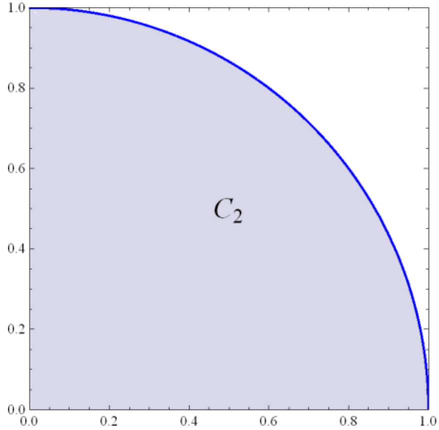

- Let $\alpha \geq 1$. Capacity regions:

$$C_k = \{\mathbf{c}_k \geq 0 : \sum_{j=1}^{k} c_{kj}^{\alpha} \leq 1\}$$

Outline of Part 3

- Introduction
- Time-scale separation
- Optimal flow-level operating policy
- Examples
- Optimal time-slot-level scheduler
- Summary

Optimal scheduling problem (transient system)

- Assume that there are n jobs in the system at time 0
- What is the optimal way to make the system empty?
- Objective: Minimize the mean delay (or flow time)
- Define: Flow time (or total completion time) for policy ϕ

$$T^{\phi} = \sum_{i=1}^{n} t_i^{\phi}$$

where t_i is the completion time of job i

- Define: Operating policies

$$\Phi_n = \{\phi = (\mathbf{c}_1, \ldots, \mathbf{c}_n) : \mathbf{c}_k \in C_k \text{ for all } k\}$$

Trivial case: One job

- Define:

$$G_1^* = \frac{1}{c_1^*}, \quad c_1^* = \max_{c_1 \in C_1} c_1$$

- Now

$$T^* = \min_{\phi \in \Phi_1} T^{\phi} = s_1 G_1^*, \quad \phi^* = (\mathbf{c}_1^*)$$

Simple case: Two jobs

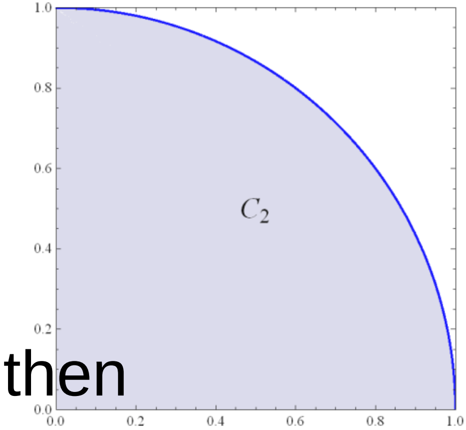

- If job 2 (i.e., the shorter one) completes first, then

$$T^{\phi} = 2\frac{s_2}{c_{22}} + (s_1 - \frac{s_2}{c_{22}} c_{21})\frac{1}{c_1^*} = \frac{s_2}{c_{22}}(2 - \frac{c_{21}}{c_1^*}) + \frac{s_1}{c_1^*}$$

- Otherwise

$$T^{\phi} = 2\frac{s_1}{c_{21}} + (s_2 - \frac{s_1}{c_{21}} c_{22})\frac{1}{c_1^*} = \frac{s_1}{c_{21}}(2 - \frac{c_{22}}{c_1^*}) + \frac{s_2}{c_1^*}$$

- Let us minimize (a function not depending on sizes!)

$$g(\mathbf{c}_2) = \frac{1}{c_{22}}(2 - \frac{c_{21}}{c_1^*}), \quad \mathbf{c}_2 \in C_2$$

Simple case: Two jobs (cont.)

- Geometric interpretation

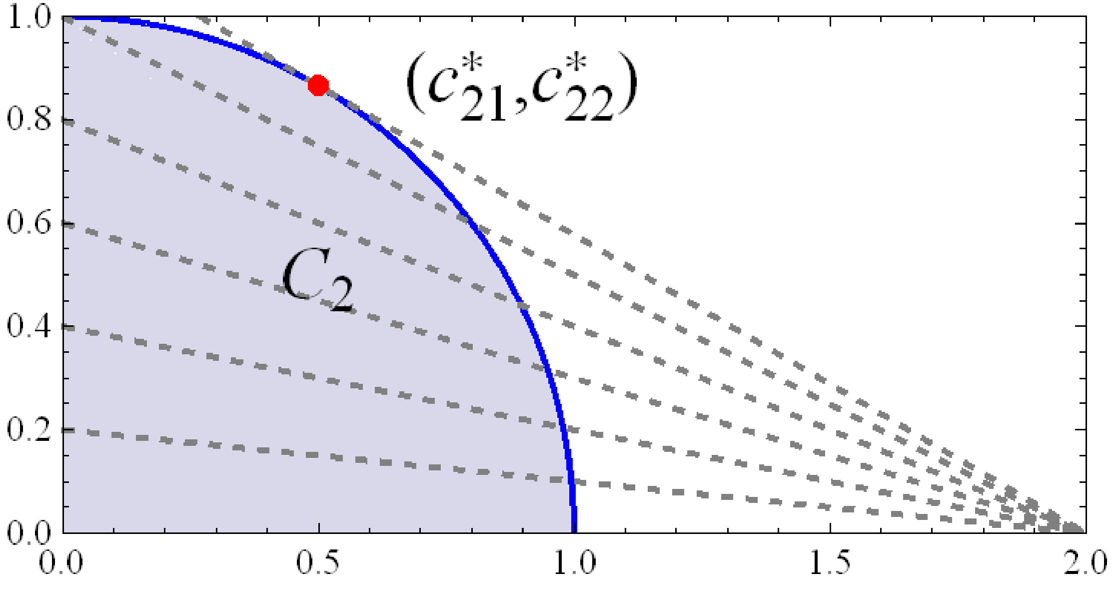

Simple case: Two jobs (cont.)

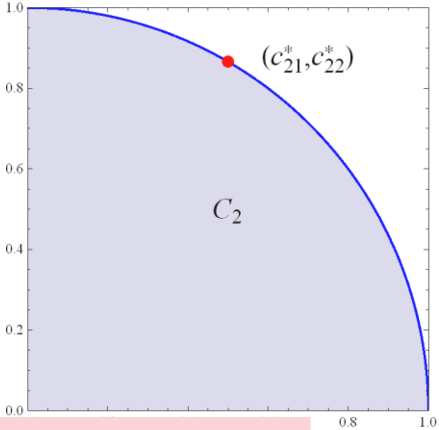

- Define:

$$G_2^* = g(\mathbf{c}_2^*) = \min_{\mathbf{c}_2 \in C_2} g(\mathbf{c}_2)$$

- Result: **If**

$$G_1^* < G_2^*$$

then (due to the symmetry property!)

$$T^* = \min_{\phi \in \Phi_2} T^\phi = s_2 G_2^* + s_1 G_1^*,\ \ \phi^* = (\mathbf{c}_1^*, \mathbf{c}_2^*),\ \ c_{21}^* \le c_{22}^*$$

Simple case: Two jobs (cont.)

- Justification:

$$
\begin{aligned}
T^{\phi} &\geq \min\{s_2 g(c_{21}, c_{22}) + s_1 G_1^*, s_1 g(c_{22}, c_{21}) + s_2 G_1^*\} \\
&\geq \min\{s_2 G_2^* + s_1 G_1^*, s_1 G_2^* + s_2 G_1^*\} \\
&= s_2 G_2^* + s_1 G_1^* \qquad [\text{since } G_2^* > G_1^*] \\
T^{\phi^*} &= s_2 g(c_{21}^*, c_{22}^*) + s_1 G_1^* \qquad [\text{since } c_{22}^* \geq c_{21}^*] \\
&= s_2 G_2^* + s_1 G_1^*
\end{aligned}
$$

Simple case: Two jobs (cont.)

- Required additional result:

$$\frac{1}{c_{22}^{*}}\left(2-\frac{c_{21}^{*}}{c_{1}^{*}}\right) \le \frac{1}{c_{21}^{*}}\left(2-\frac{c_{22}^{*}}{c_{1}^{*}}\right) \Leftrightarrow$$

$$c_{21}^{*}\left(2-\frac{c_{21}^{*}}{c_{1}^{*}}\right) \le c_{22}^{*}\left(2-\frac{c_{22}^{*}}{c_{1}^{*}}\right) \Leftrightarrow$$

$$(c_{22}^{*}-c_{21}^{*})\left(2-\frac{c_{21}^{*}}{c_{1}^{*}}-\frac{c_{22}^{*}}{c_{1}^{*}}\right) \ge 0 \Leftrightarrow$$

$$c_{22}^{*}(c_{22}^{*}-c_{21}^{*})(G_{2}^{*}-G_{1}^{*}) \ge 0$$

Simple case: Two jobs (cont.)

- Equivalent condition:

$$G_2^* > G_1^* \iff$$

$$c_{21} + c_{22} < 2 \cdot c_1^*$$

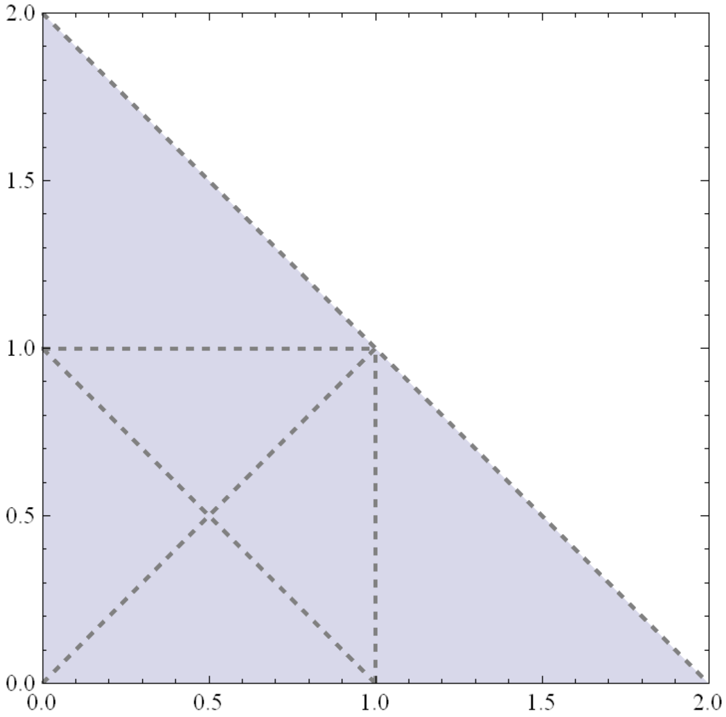

- Suffient condition: nested capacity regions
- Note: However, capacity regions are not required to be nested

General case: n jobs

- Define (recursively):

$$G_k^* = \min_{\mathbf{c}_k \in C_k} g_k(\mathbf{c}_k), \quad g_k(\mathbf{c}_k) = \frac{1}{c_{kk}}\left(k - \sum_{i=1}^{k-1} c_{ki} G_i^*\right)$$

- Theorem 1: **If**

$$G_1^* < \ldots < G_n^*$$

then

$$T^* = \min_{\phi \in \Phi_n} T^{\phi} = \sum_{k=1}^{n} s_k G_k^*, \quad \phi^* = (\mathbf{c}_1^*, \ldots, \mathbf{c}_n^*)$$

General case: n jobs (cont.)

- In addition,

$$c_{k1}^{*} \le \ldots \le c_{kk}^{*} \quad \text{for all } k$$

- Thus, the optimal policy applies the SRPT-FM principle:
  - the shortest job is served with the highest rate,
  - the second shortest job is served with the second highest rate,
  - etc.

- Note also that the optimal rate vector does not depend on the absolute sizes (only on their order)

Outline of Part 3

- Introduction
- Time-scale separation
- Optimal flow-level operating policy
- Examples
- Optimal time-slot-level scheduler
- Summary

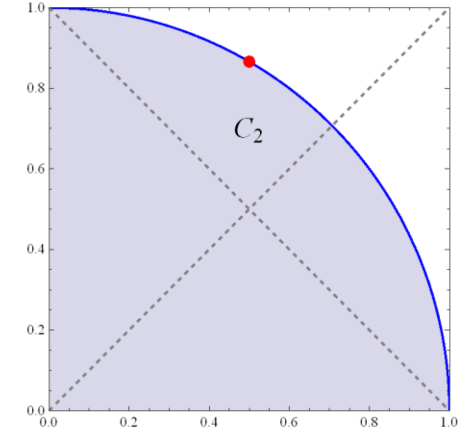

Example: Alpha-ball

- Let $\alpha \geq 1$ and consider capacity regions

$$C_k = \{\mathbf{c}_k \geq 0 : \sum_{j=1}^{k} c_{kj}^{\alpha} \leq 1\}$$

- Now

$$G_k^* = \left(k^{\frac{\alpha}{\alpha-1}} - (k-1)^{\frac{\alpha}{\alpha-1}} \right)^{\frac{\alpha-1}{\alpha}} \quad \text{(increasing in } k\text{)}$$

$$c_{kj}^* = \left(\frac{G_j^*}{k} \right)^{\frac{1}{\alpha-1}} \quad \text{(increasing in } j\text{)}$$

Alpha = 1.0 (single-server queue)

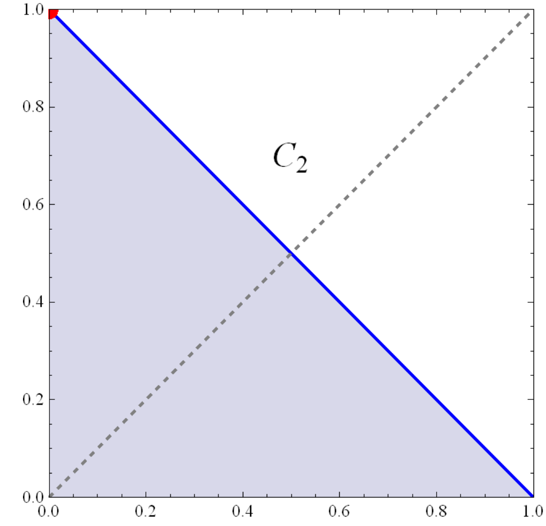

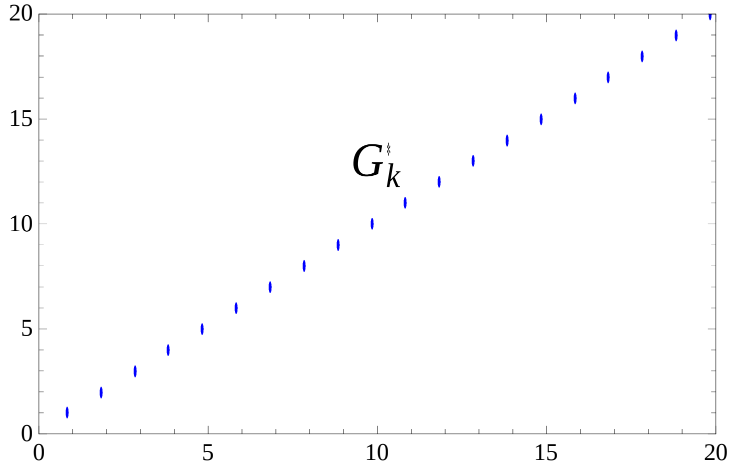

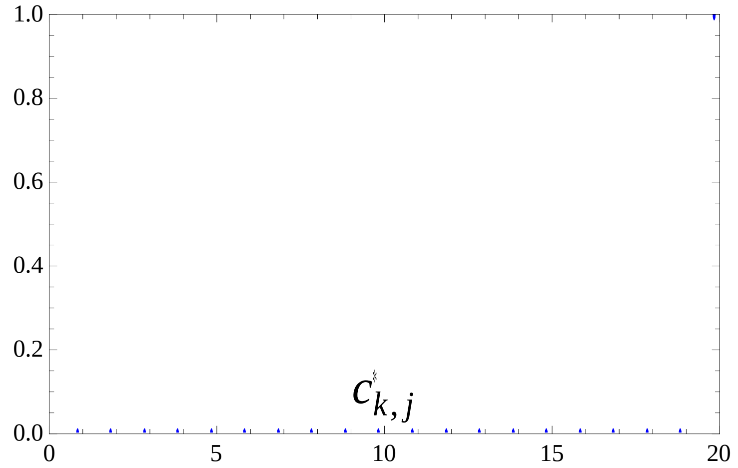

Alpha = 1.2

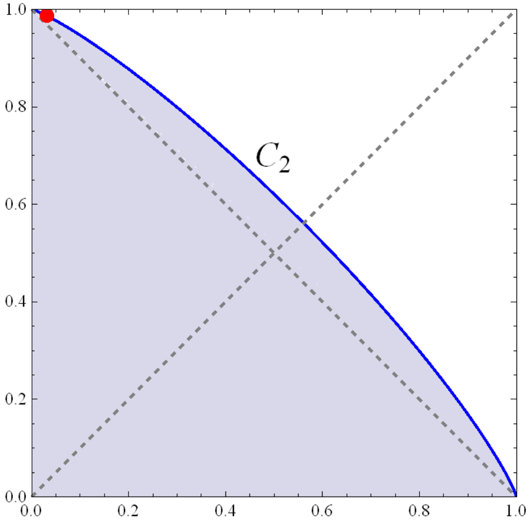

1.0
0.8
0.6
0.4
0.2
0.0
0.0
0.2
0.4
0.6
0.8
1.0
C_2

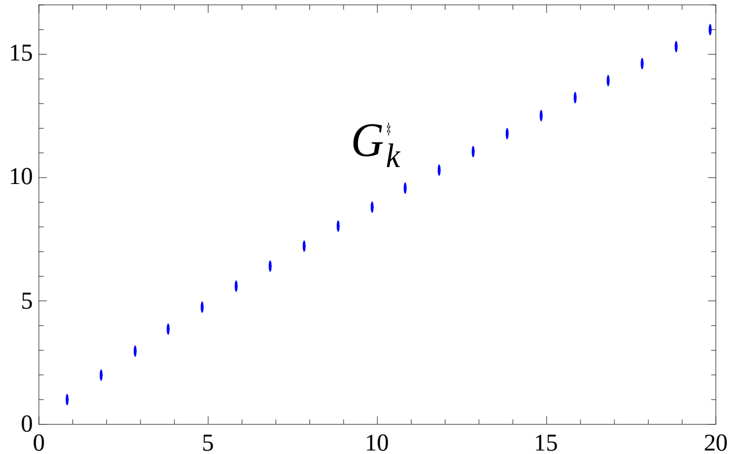

15
10
5
0
0
5
10
15
20
$G_k^{\ddagger}$

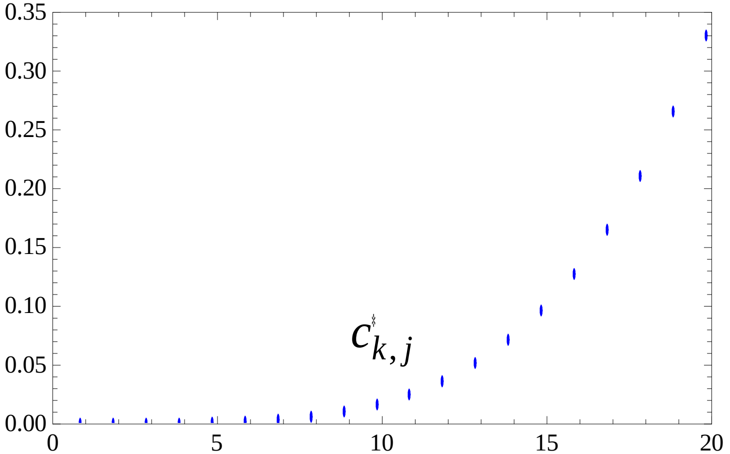

0.35
0.30
0.25
0.20
0.15
0.10
0.05
0.00
0
5
10
15
20
$c_{k,j}^{\ddagger}$

Alpha = 2.0

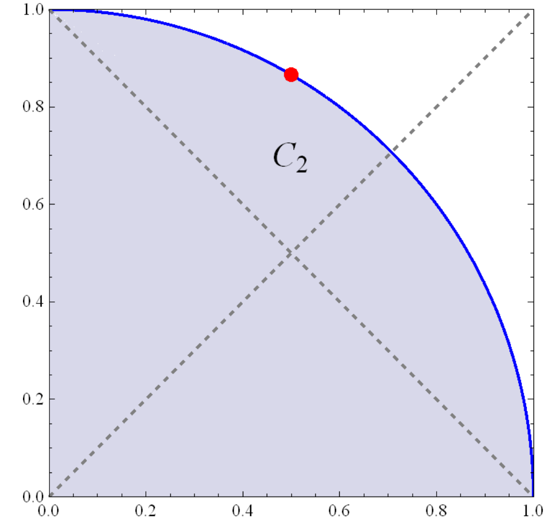

1.0
0.8
0.6
0.4
0.2
0.0
0.0
0.2
0.4
0.6
0.8
1.0
C_2

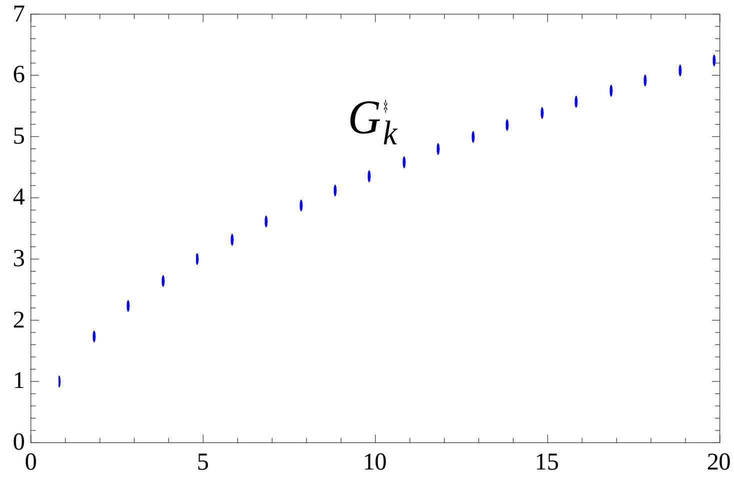

7
6
5
4
3
2
1
0
0
5
10
15
20
$G_k^\natural$

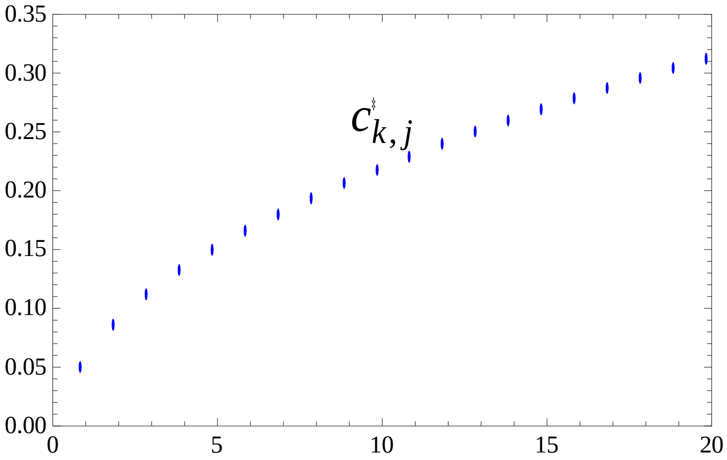

0.35
0.30
0.25
0.20
0.15
0.10
0.05
0.00
0
5
10
15
20
$c_{k,j}^\natural$

Alpha = 5.0

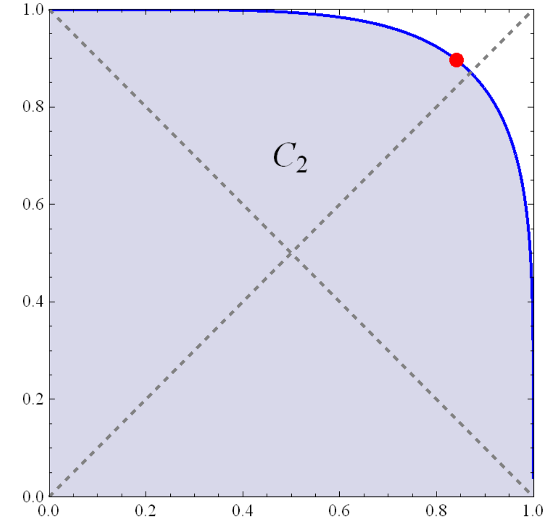

1.0
0.8
0.6
0.4
0.2
0.0
C_2
0.0
0.2
0.4
0.6
0.8
1.0

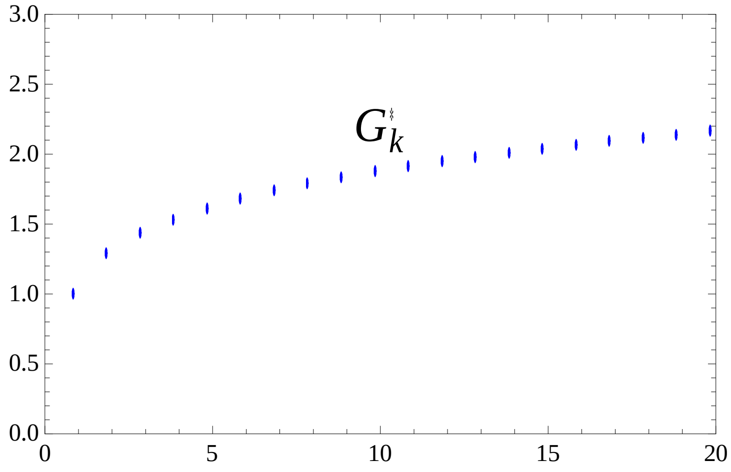

3.0
2.5
2.0
1.5
1.0
0.5
0.0
$G_k^\ddagger$
0
5
10
15
20

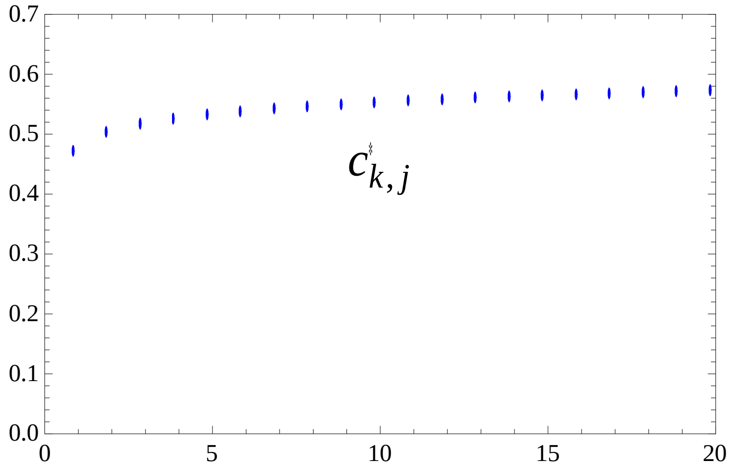

0.7
0.6
0.5
0.4
0.3
0.2
0.1
0.0
$c_{k,j}^\ddagger$
0
5
10
15
20

Alpha = infinite (infinite-server queue)

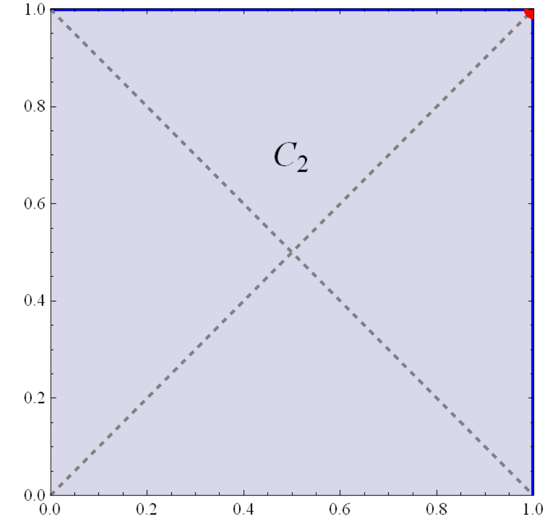

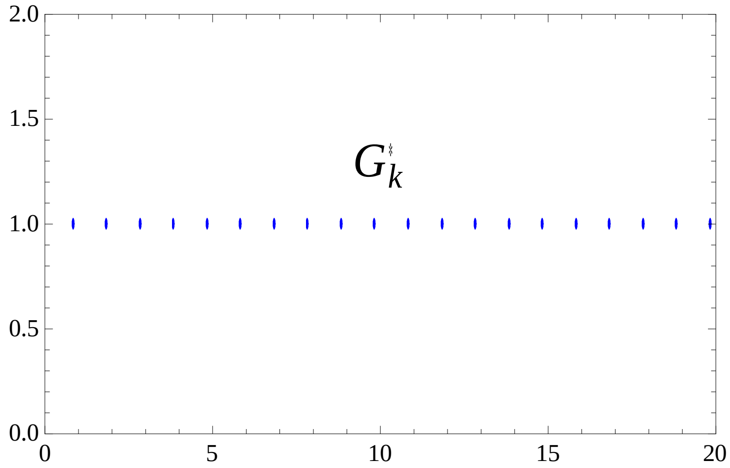

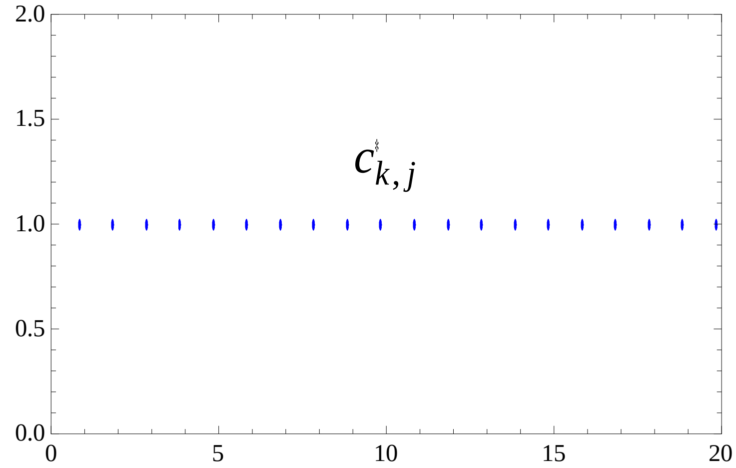

Outline of Part 3

- Introduction
- Time-scale separation
- Optimal flow-level operating policy
- Examples
- Optimal time-slot-level scheduler
- Summary

Summary thus far

- Assumptions:
  - Abstract capacity regions (time-scale separation)
  - Transient system
- Results:
  - Optimality result for compact and symmetric capacity regions
    - includes nested polymatroids (cf. Sadiq and de Veciana (2010))
    - requires an implicit condition related to capacity regions
  - Optimal rate vectors for each phase
    - applying the SRPT-FM principle
- Open questions:
  - Is it possible to make the implicit condition explicit?
  - Is it possible to implement the optimal policy at time-slot scale?

Time-slot-level model

- $R(t) = (R_1(t),\ldots, R_k(t))$ = rate vector in time slot t
- $R_i(t)$ = instantaneous rate of user i
- Assume: $R_i(t)$ is a stationary and ergodic process taking values in a finite set
- Assume: Processes $R_i(t)$ are IID (symmetric case)

Time-slot-level schedulers

- Assume: Scheduling policy $\pi \in \Pi_k$ is stationary
- Define: The long-term throughput for user i:

$$\theta_i^{\pi} = \sum_{\mathbf{r}} r_i p_i^{\pi}(\mathbf{r}) P\{R(t) = \mathbf{r}\}$$

- Define: The (opportunistic) capacity region:

$$C_k = \{(\theta_1^{\pi}, \ldots, \theta_k^{\pi}) \in \Re_+^k : \pi \in \Pi_k\}$$

- Note: Capacity regions are compact and symmetric

Weight-based schedulers

- Define: Weight-based scheduler $\pi \in \Pi_k$ allocates time slot t to user i^* for which

$$w_{i*}R_{i*}(t) = \max_i w_i R_i(t)$$

where w_i are the weights related to the scheduler

- Example: MR (which is the same as PF in our case) $w_i = 1$ for all i

Connection between the two time scales

- Proposition 1:

$$E[\max_i w_i R_i(t)] = \max_{\mathbf{c}_k \in C_k} \sum_i w_i c_{ki}$$

  – Proof is straightforward:

$$\begin{aligned}\max_{\mathbf{c}_k} \sum_i w_i c_{ki} &= \max_\pi \sum_{\mathbf{r}} \sum_i w_i r_i p_i^\pi(\mathbf{r}) P\{R(t) = \mathbf{r}\} \\ &= \sum_{\mathbf{r}} (\max_i w_i r_i) P\{R(t) = \mathbf{r}\} \\ &= E[\max_i w_i R_i(t)]\end{aligned}$$

Recall the optimal scheduling problem (transient system)

- Assume that there are n jobs in the system at time 0
- What is the optimal way to make the system empty?
- Objective: Minimize the mean delay (or flow time)
- Define: Flow time (or total completion time) for policy ϕ

$$T^{\phi} = \sum_{i=1}^{n} t_i^{\phi}$$

where t_i is the completion time of job i

- Define: Operating policies

$$\Phi_n = \{\phi = (\mathbf{c}_1, \ldots, \mathbf{c}_n) : \mathbf{c}_k \in C_k \text{ for all } k\}$$

Recall the recursion for G* (based on the flow-level model)

- Define (recursively):

$$G_k^* = \min_{\mathbf{c}_k \in C_k} g_k(\mathbf{c}_k), \quad g_k(\mathbf{c}_k) = \frac{1}{c_{kk}}\left(k - \sum_{i=1}^{k-1} c_{ki} G_i^*\right)$$

- Open problem 1: Is it possible to show that in our case

$$G_1^* < \ldots < G_n^*$$

- Open problem 2: If so, how to implement the optimal operating policy with a time-slot-level scheduler so that

$$\theta_i^{\pi_k^*} = c_{ki}^* \quad \text{for all } k, i$$

Key property

- Proposition 2:

$$E[\max_i G_i^* R_i] = \max_{\mathbf{c}_k \in C_k} \sum_i G_i^* c_{ki} = \sum_i G_i^* c_{ki}^* = k$$

  - Proof by induction

Alternative recursion for G* (based on the time-slot-level model)

- Define (recursively):

$$f_k(a) = \int_0^\infty \left(1 - P\{aR_k \le r\} \prod_{i=1}^{k-1} P\{G_i^* R_i \le r\}\right) dr$$

$$G_k^* = f_k^{-1}(k) \quad \text{(well-defined since } f_k \text{ increasing)}$$

  – Based on the equation:

$$E[\max_{i=1,\ldots,k} G_i^* R_i] = \int_0^\infty \left(1 - \prod_{i=1}^{k} P\{G_i^* R_i \le r\}\right) dr = k$$

Key result

- Proposition 3:

$$G_1^* < \ldots < G_n^*$$

  - Proof by induction
  - Idea briefly on the following slide
- Corollary: Solution of the optimal scheduling problem

$$T^* = \min_{\phi \in \Phi_n} T^{\phi} = \sum_{k=1}^{n} s_k G_k^*, \quad \phi^* = (\mathbf{c}_1^*, \ldots, \mathbf{c}_n^*)$$

$$c_{k1}^* \leq \ldots \leq c_{kk}^* \quad \text{for all } k$$

Idea of the proof

- Define:

$$X_k = \max_{i=1,\dots,k} G_i^* R_i$$

$$h_{k+1}(a) = E[(aR_{k+1} - X_k) \cdot 1_{\{aR_{k+1} > X_k\}}]$$

- Easily: $h_{k+1}(a)$ is non-decreasing and satisfies

$$h_{k+1}(G_{k+1}^*) = E[X_{k+1} - X_k] = (k+1) - k = 1$$

- It remains to show that

$$h_{k+1}(G_k^*) < 1$$

Optimal time-slot-level scheduler for flow-level performance

- Theorem 2: The optimal operating policy ϕ^* can be implemented by a sequence of weight-based schedulers π_k defined by weight vectors

$$\mathbf{w}_k = (G_1^*, \ldots, G_k^*)$$

  - Proof based on Propositions 1 and 2
- Summary: The optimal time-slot-level scheduler allocates time slot t to user i^* for which

$$G_{i^*}^* R_{i^*}(t) = \max_i G_i^* R_i(t)$$

Outline of Part 3

- Introduction
- Time-scale separation
- Optimal flow-level operating policy
- Examples
- Optimal time-slot-level scheduler
- Summary

Summary of Part 3

- Assumptions:
  - Stationary and ergodic rate processes
  - Symmetric case (rate processes IID for different users)
  - Transient system (a batch of jobs without new arrivals)
- Results:
  - Optimality result based on a time-scale separation argument
  - Optimal flow-level rate vectors for each phase
  - Optimal time-slot-level scheduler constructed
- Open questions:
  - Optimal scheduler for the asymmetric case (with non-IID users)?
  - Optimal scheduler for the dynamic system (with new arrivals)?

Related contributions

- S. Aalto, A. Penttinen, P. Lassila and P. Osti,
  On the optimal trade-off between SRPT and opportunistic scheduling, in *ACM Sigmetrics* 2011

- S. Aalto, A. Penttinen, P. Lassila and P. Osti,
  Optimal size-based opportunistic scheduler for wireless systems, *Queueing Systems*, 2012 (to appear)

Final remarks

The End