Heavy Tails: Performance Models and Scheduling Disciplines

Part II – Workload Asymptotics for Generalized Processor Sharing Systems

Sem Borst
Bell Labs - CWI - TU/e

ITC-18, Berlin, August 31, 2003

Based on joint work with Onno Boxma, Predrag Jelenković, Michel Mandjes & Miranda van Uitert

Organization

1. Background & motivation
2. Generalized Processor Sharing (GPS)
3. Performance evaluation
4. Model description
5. Workload asymptotics in various scenarios
6. Discussion & conclusion
7. References

Background & motivation

Future Internet expected to support variety of services

Voice and video communications induce far more stringent QoS requirements than typical data applications

Integration of heterogeneous services raises need for differentiated QoS

Packet scheduling provides natural mechanism to achieve differentiated QoS

Scheduling mechanisms should be able to cope with adversarial or erratic traffic characteristics

Packet scheduling may be implemented at various levels

- Individual traffic flows (e.g. IntServ)

- Aggregate traffic flows / service classes (e.g. DiffServ: Expedited Forw. (EF), Assured Forw. (AF), BE)

Packet scheduling may be implemented at various levels

- **Individual traffic flows (e.g. IntServ)**
- **Aggregate traffic flows / service classes (e.g. DiffServ: Expedited Forw. (EF), Assured Forw. (AF), BE)**

Involves trade-off between implementation complexity and degree of service differentiation

- **For scalability reasons, packet scheduling at granularity level of individual flows in core is viewed as impractical**
- **Packet scheduling at aggregate level does not provide strict guarantees to individual flows**

Possible intermediate scenario

- **Fine-grained scheduling at network edge (in particular wireless access and application servers)**

- **Coarse-level or no scheduling in network core**

Generalized Processor Sharing (GPS)

In GPS, each traffic class is assigned some positive weight

Bandwidth is shared among backlogged classes in proportion to respective weight factors

Two crucial properties

- Minimum-rate guarantees, providing flow isolation and preventing starvation effects
- Work conservation, achieving statistical multiplexing gains and thus ensuring efficient bandwidth utilization

GPS includes strict-priority scheduling as special case

Weights offer greater flexibility in service differentiation

However, weights play “double role”, fixing absolute minimum rate as well as relative rate share

These two rate attributes thus appear intertwined

GPS is idealized mechanism, assuming bandwidth is infinitely divisible and can be shared in infinitesimal quanta

In practice, traffic consists of cells or packets, and bandwidth can only be provided in discrete quanta

Various packet-based emulations of GPS proposed, most notably Weighted Fair Queueing (WFQ) and numerous variants (WFQ^{+}, virtual-clock FQ, self-clocked FQ, ..., ...)

Use time-stamping of packets based on 'background simulation' of idealized GPS mechanism

Involve trade-off between implementation complexity and accuracy

WFQ variants also proposed for use in wireless networks

Raises various additional issues related to idiosyncrasies of wireless propagation characteristics

- **Heterogeneity in rate among spatially distributed users (rate shares differ from time shares)**
- **Rate variations (over time)**
- **Transmission errors**

Performance evaluation

Focus on evaluation of performance for given weights

Inverse problem: how to set weights to meet given performance target
[Elwalid & Mitra (1999), Kumaran & Mitra (2000)]

Performance evaluation

Focus on evaluation of performance for given weights

Inverse problem: how to set weights to meet given performance target
[Elwalid & Mitra (1999), Kumaran & Mitra (2000)]

In GPS system, service rate of each class depends on workload of other classes

Interdependence between classes complicates analysis

Exact analysis extremely difficult, motivating derivation of bounds and asymptotics

GPS system is equivalent to coupled-processors model

In coupled-processors model, service rate of each queue depends on whether other queues are empty or not

Latter model has been studied for two-queue case

- **Fayolle & Iasnogorodski (1979) consider exponential service times and reduce analysis of joint queue length distribution to Riemann-Hilbert problem**

- **Cohen & Boxma (1983) extend analysis to general service times and obtain joint workload distribution as solution to boundary-value problem**

Delay bounds

- **Det. delay bounds for leaky-bucket controlled traffic [Parekh & Gallager (1993, 1994)]**

- **Statist. delay bounds for exponentially-bounded traffic [Yaron & Sidi (1994), Yu *et al.* (2003)]**

Workload asymptotics

Main distinctions

- Light-tailed versus heavy-tailed traffic characteristics
- Large-buffer versus many-sources regime
- Exact versus logarithmic asymptotics
- Sample path techniques or large-deviations principles versus Tauberian theorems

Tutorial focuses on exact large-buffer asymptotics for combination of heavy-tailed and light-tailed traffic

- **Logarithmic large-buffer asymp. for light-tailed traffic: Bertsimas, Paschalidis & Tsitsiklis (1999), Massoulié (1999), Zhang *et al.* (1995, 1996, 1997, 1998)**

- **Logarithmic many-sources asymp. for various models: Kotopoulos & Mazumdar (2002)**

- **Logarithmic many-sources asymp. for Gaussian traffic: Mannersalo & Norros (2002), Mandjes & Van Uitert (2003)**

‘Workload’ need *not* be limited to buffer content, but may also include backlog at end-users device

Main commonalities/caveats

- **Infinite-buffer model (no loss)**
  [Jelenković & Momčilović (2001, 2002) consider finite-buffer model]

- **Exogenous traffic (no feedback at ‘workload’ level)**
  [Arvidsson & Karlsson (1999) examine buffer content for TCP/IP]

Main commonalities/caveats (cont'd)

- **Single-node models**
  [networks analyzed in Van Uitert & B (2001), (2002)]

- **Packet-level performance (static population of classes)**
  [dynamic population of users (flow-level performance) gives rise to Discriminatory Processor-Sharing models (B, Van Ooteghem & Zwart (2003))]

Model description

Two classes sharing link of unit rate

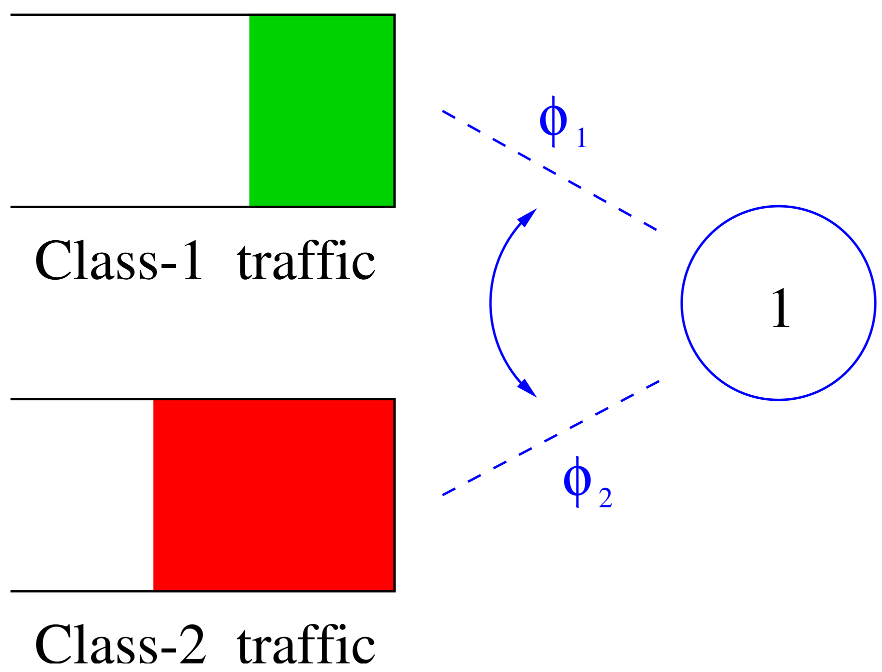

Class i is assigned weight $\phi_i \geq 0$, with $\phi_1 + \phi_2 = 1$

If both classes are backlogged, then class i receives service at rate ϕ_i

If one class is *not* backlogged, then its (excess) capacity is re-allocated to the other class, which then receives service at full link rate

If both classes are backlogged, then class i receives service at rate ϕ_i

If one class is *not* backlogged, then its (excess) capacity is re-allocated to the other class, which then receives service at full link rate

Let ρ_i be traffic intensity of class i

Let V_i^{GPS} be stationary workload of class i

Traffic assumptions

Class 1 has 'light-tailed' characteristics, e.g.,

- **G/G/1 input with 'exponentially-bounded' service times**
- **Markov-modulated fluid input**

Traffic assumptions

Class 1 has 'light-tailed' characteristics, e.g.,

- G/G/1 input with 'exponentially-bounded' service times
- Markov-modulated fluid input

Class 2 has 'heavy-tailed' characteristics, e.g.,

- Instantaneous 'heavy-tailed' bursts B_2
- On-Off process with 'heavy-tailed' On-periods A_2 with fraction On-time p_2, peak rate r_2

Theorem [Cohen (1973), Pakes (1975)]

If B_i^r is subexponential, and $\rho_i < c$, then

$$\mathbb{P}\{\mathbf{V}_i^c > x\} \sim \frac{\rho_i}{c - \rho_i} \mathbb{P}\{\mathbf{B}_i^r > x\} \qquad \textbf{as } x \to \infty$$

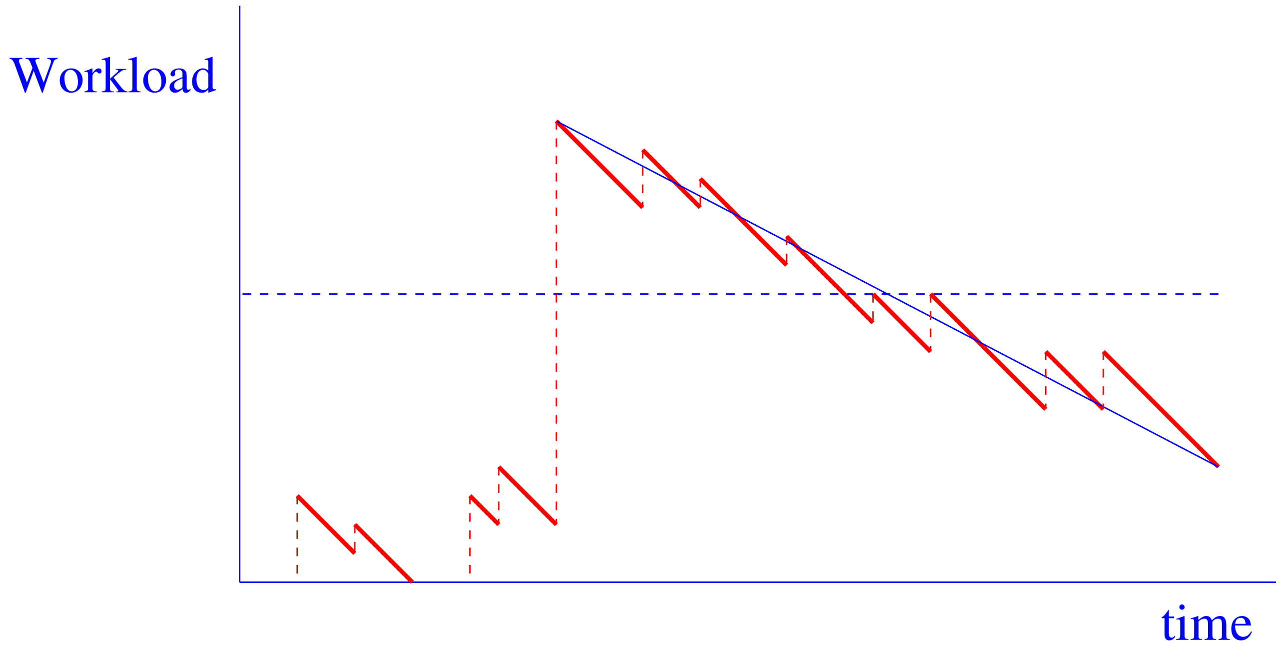

Catastrophe scenario:

Due to SINGLE extremely large burst

Theorem [Jelenković & Lazar (1999)]

If $\mathbf{A}_i^r$ is subexponential, and $\rho_i < c < r_i$, then

$$\mathbb{P}\{\mathbf{V}_i^c > x\} \sim (1 - p_i)\frac{\rho_i}{c - \rho_i}\mathbb{P}\{\mathbf{A}_i^r > x/(r_i - c)\} \qquad \text{as } x \to \infty$$

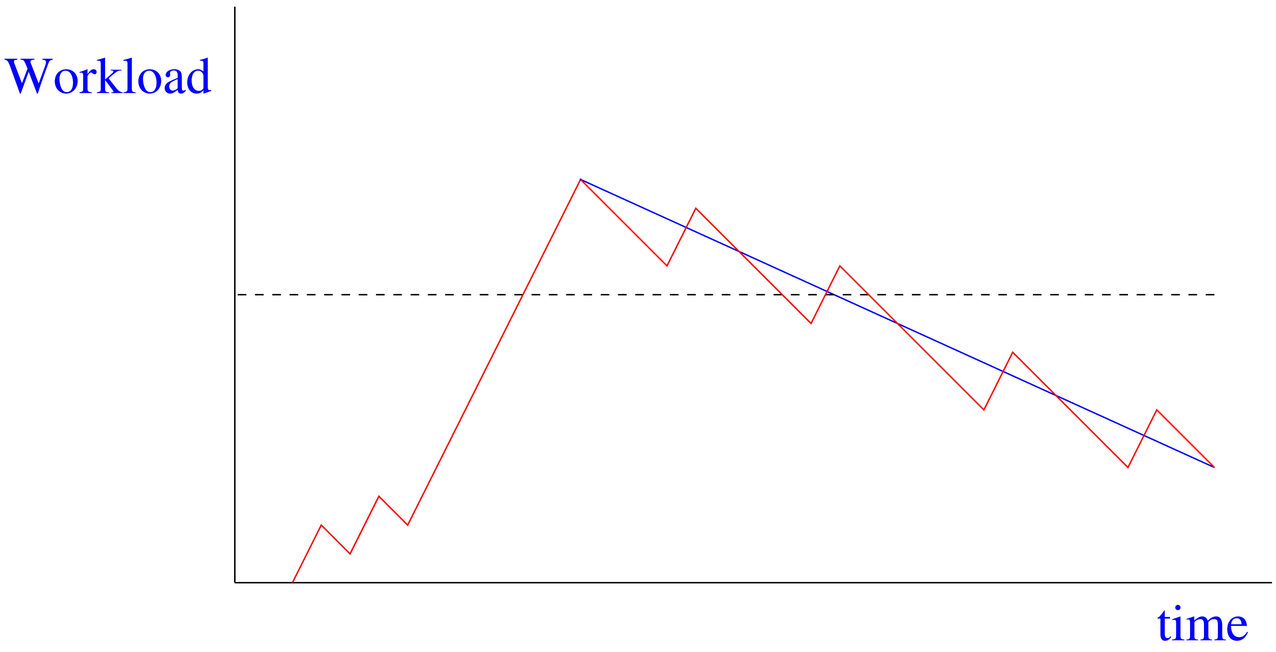

Due to SINGLE extremely long On-period

In contrast, class-1 builds up large workload level in gradual manner

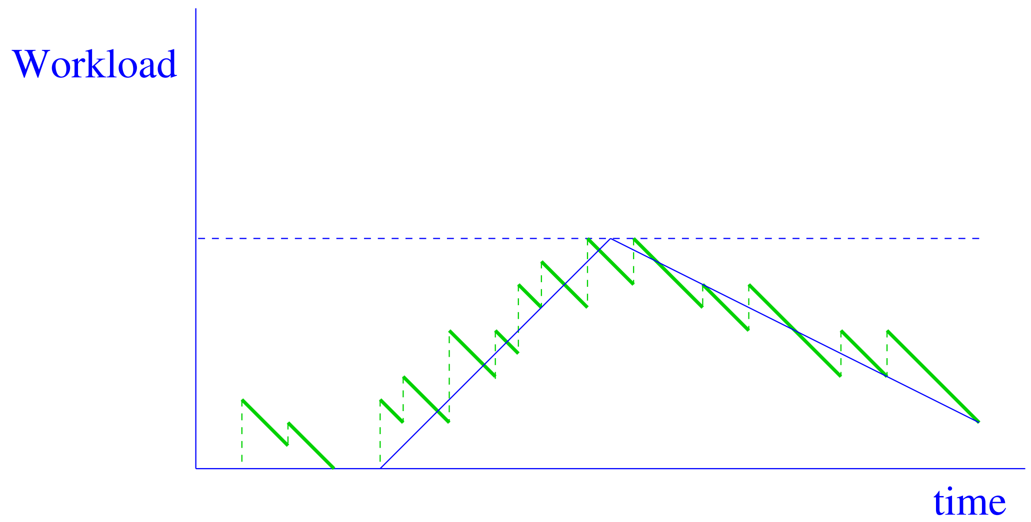

Conspiracy scenario:

Combination of MANY relatively large bursts and MANY relatively short interarrival times

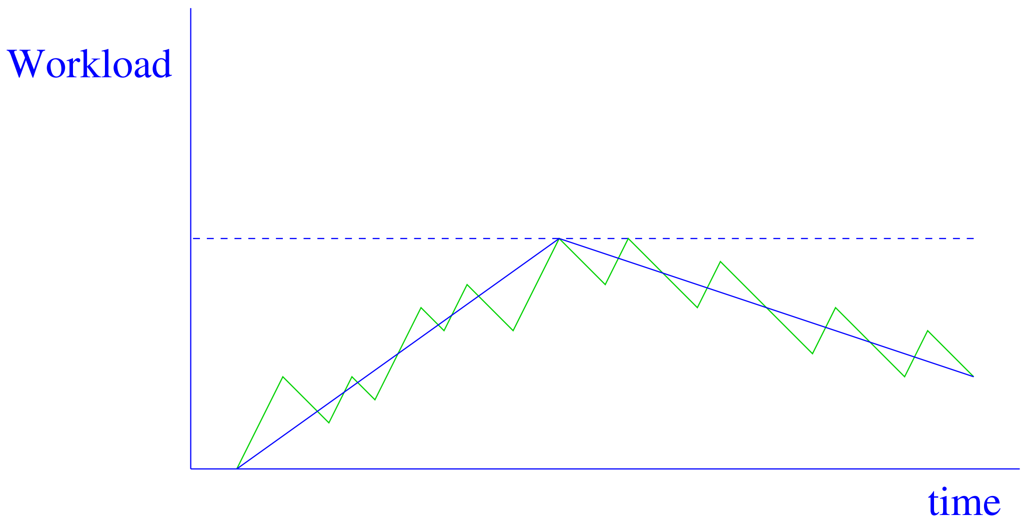

Combination of MANY relatively long On-periods and MANY relatively short Off-periods

Workload asymptotics in various scenarios

Class-2 workload behavior

Case I: $\rho_1 < \phi_1$, $\rho_2 < \phi_2$

Catastrophe scenario:

- **Class 2 generates large burst (or long On-period)**
- **Class 1 generates traffic at rate** $\rho_1 < \phi_1$
- **Class 2 is effectively served at rate** $1 - \rho_1$

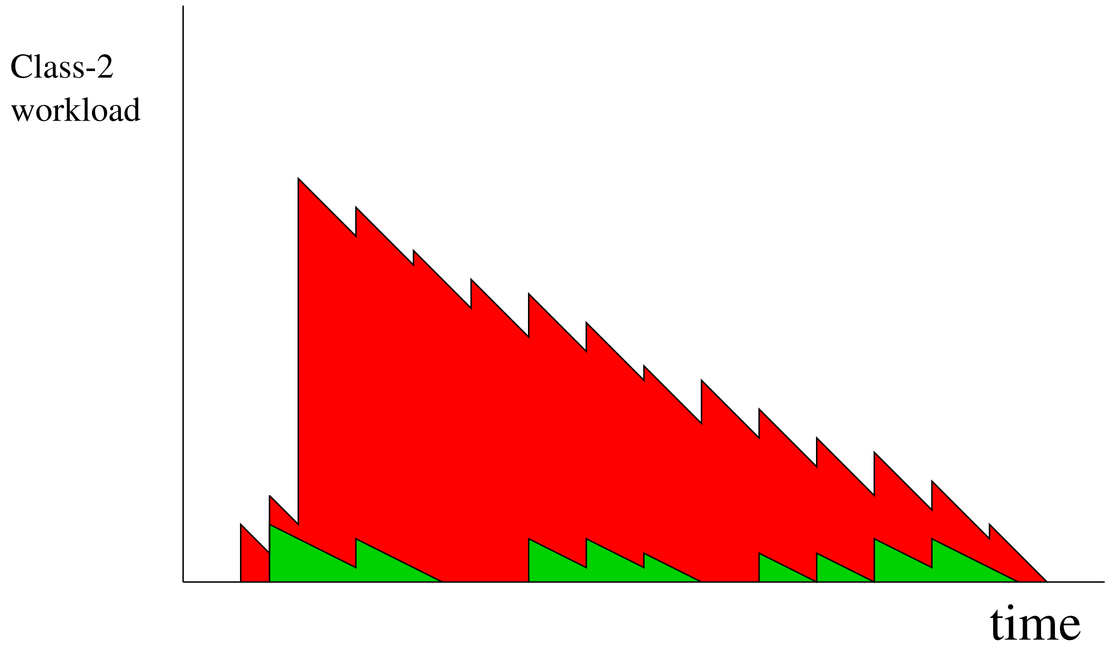
Class-2
workload
time

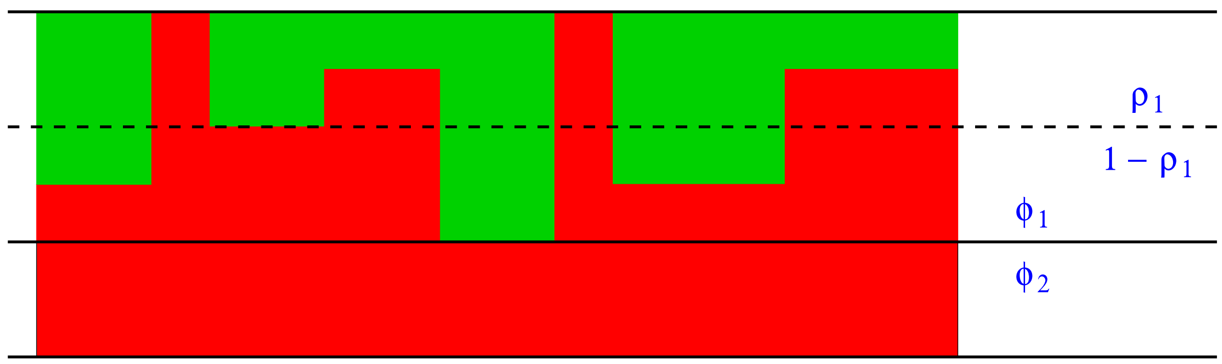
ρ_1
$1 - \rho_1$
ϕ_1
ϕ_2

Theorem

If A_2^r or B_2^r is regularly varying, $\rho_1 < \phi_1$ and $\rho_2 < \phi_2$, then

$$\mathbb{P}\{\mathbf{V}_2^{GPS} > x\} \sim \mathbb{P}\{\mathbf{V}_2^{1-\rho_1} > x\} \qquad \text{as } x \to \infty$$

Reduced-load equivalence (RLE):
Class-2 workload roughly behaves as in isolated system with rate $1 - \rho_1$

Similar behavior has been shown for total workload in queues fed by mixture of heavy-tailed and light-tailed input [Agrawal, Nain & Makowski (1999), Zwart, B & Mandjes (2001)]

Note: here *independent* of class-1 traffic characteristics

Sample path lower bound

$$V_i^{GPS}(t) \geq V_i^{1-\rho_{-i}+\delta}(t) - \underbrace{U_{-i}^{\rho_{-i}-\delta}(t) - \sum_{j \neq i} V_j^{\phi_j}(t)}_{\text{"small correction terms"}}$$

Proof

Sample path wise,

$$
\begin{aligned}
V_i^{GPS}(t) &= V^{GPS}(t) - \sum_{j \neq i} V_j^{GPS}(t) \\
&\overset{Min-rate\ \ guarantee}{\geq} V^{GPS}(t) - \sum_{j \neq i} V_j^{\phi_j}(t) \\
&\overset{Work-conservation}{=} \sup_{0 \leq s \leq t} \{A(s,t) - (t-s)\} - \sum_{j \neq i} V_j^{\phi_j}(t) \\
&\geq \sup_{0 \leq s \leq t} \{A_i(s,t) - (1-\theta)(t-s)\} \\
&- \sup_{0 \leq s \leq t} \{\theta(t-s) - A_{-i}(s,t)\} - \sum_{j \neq i} V_j^{\phi_j}(t) \\
&= V_i^{1-\theta}(t) - U_{-i}^{\theta}(t) - \sum_{j \neq i} V_j^{\phi_j}(t)
\end{aligned}
$$

Then take $\theta = \rho_{-i} - \delta$

Sample path upper bound

$$V_i^{GPS}(t) \leq \min\{V_i^{\phi_i}(t), V_i^{1-\rho_{-i}-\delta}(t) + \underbrace{V_{-i}^{\rho_{-i}+\delta}(t)}_{\text{“correction term”}}\}$$

Proof

Sample path wise,

$$
\begin{aligned}
V_i^{GPS}(t) &\leq V^{GPS}(t) \\
&\overset{Work-conservation}{=} \sup_{0 \leq s \leq t} \{A(s,t) - (t-s)\} \\
&\leq \sup_{0 \leq s \leq t} \{A_i(s,t) - (1-\theta)(t-s)\} \\
&+ \sup_{0 \leq s \leq t} \{A_{-i}(s,t) - \theta(t-s)\} \\
&= V_i^{1-\theta}(t) + V_{-i}^{\theta}(t)
\end{aligned}
$$

Also,

$$V_i^{GPS}(t) \overset{Min-rate\ \ guarantee}{\leq} V_i^{\phi_i}(t)$$

Then take $\theta = \rho_{-i} + \delta$

Want to show

If A_2^r **or** B_2^r **is regularly varying,** $\rho_1 < \phi_1$ **and** $\rho_2 < \phi_2$**, then**

$$\mathbb{P}\{\mathbf{V}_2^{GPS} > x\} \sim \mathbb{P}\{\mathbf{V}_2^{1-\rho_1} > x\} \qquad \text{as } x \to \infty$$

Want to show

If A_2^r **or** B_2^r **is regularly varying,** $\rho_1 < \phi_1$ **and** $\rho_2 < \phi_2$**, then**

$$\mathbb{P}\{\mathbf{V}_2^{GPS} > x\} \sim \mathbb{P}\{\mathbf{V}_2^{1-\rho_1} > x\} \qquad \text{as } x \to \infty$$

Proof (sketch)

From sample path lower bound, for any $\delta > 0$ **and** y**,**

$$\mathbb{P}\{\mathbf{V}_2^{GPS} > x\} \geq \mathbb{P}\{\mathbf{V}_2{}^{1-\rho_1+\delta} > x + y\}\mathbb{P}\{\mathbf{U}_1^{\rho_1-\delta} + \mathbf{V}_1^{\phi_1} < y\}$$

Want to show

If A_2^r **or** B_2^r **is regularly varying,** $\rho_1 < \phi_1$ **and** $\rho_2 < \phi_2$**, then**

$$\mathbb{P}\{\mathbf{V}_2^{GPS} > x\} \sim \mathbb{P}\{\mathbf{V}_2^{1-\rho_1} > x\} \qquad \text{as } x \to \infty$$

Proof (sketch)

From sample path lower bound, for any $\delta > 0$ **and** y**,**

$$\mathbb{P}\{\mathbf{V}_2^{GPS} > x\} \geq \mathbb{P}\{\mathbf{V}_2{}^{1-\rho_1+\delta} > x + y\}\mathbb{P}\{\mathbf{U}_1^{\rho_1-\delta} + \mathbf{V}_1^{\phi_1} < y\}$$

From sample path upper bound, for any $\delta > 0$ **and** y**,**

$$\mathbb{P}\{\mathbf{V}_2^{GPS} > x\} \leq \mathbb{P}\{\mathbf{V}_2{}^{1-\rho_1-\delta} > x - y\} + \mathbb{P}\{\mathbf{V}_2^{\phi_2} > x\}\mathbb{P}\{\mathbf{V}_1^{\rho_1+\delta} > y\}$$

Want to show

If A_2^r or B_2^r is regularly varying, $\rho_1 < \phi_1$ and $\rho_2 < \phi_2$, then

$$\mathbb{P}\{V_2^{GPS} > x\} \sim \mathbb{P}\{V_2^{1-\rho_1} > x\} \qquad \text{as } x \to \infty$$

Proof (sketch)

From sample path lower bound, for any $\delta > 0$ and y,

$$\mathbb{P}\{V_2^{GPS} > x\} \geq \mathbb{P}\{{V_2}^{1-\rho_1+\delta} > x + y\}\mathbb{P}\{U_1^{\rho_1-\delta} + V_1^{\phi_1} < y\}$$

From sample path upper bound, for any $\delta > 0$ and y,

$$\mathbb{P}\{V_2^{GPS} > x\} \leq \mathbb{P}\{{V_2}^{1-\rho_1-\delta} > x - y\} + \mathbb{P}\{V_2^{\phi_2} > x\}\mathbb{P}\{V_1^{\rho_1+\delta} > y\}$$

Show that, for $y \to \infty$, $\delta \downarrow 0$, both bounds behave as

$$\mathbb{P}\{V_2^{1-\rho_1} > x\}$$

Requires that A_2^r or B_2^r is regularly varying

Class-2 workload behavior (cont'd)

Case II: $\rho_1 > \phi_1$, $\rho_2 < \phi_2$

Catastrophe scenario:

- **Class 2 generates large burst (or long On-period)**
- **Class 1 generates traffic at rate $\rho_1 > \phi_1$, but only receives service at rate ϕ_1**
- **Class 2 is effectively served at rate $\phi_2 = 1 - \phi_1$**

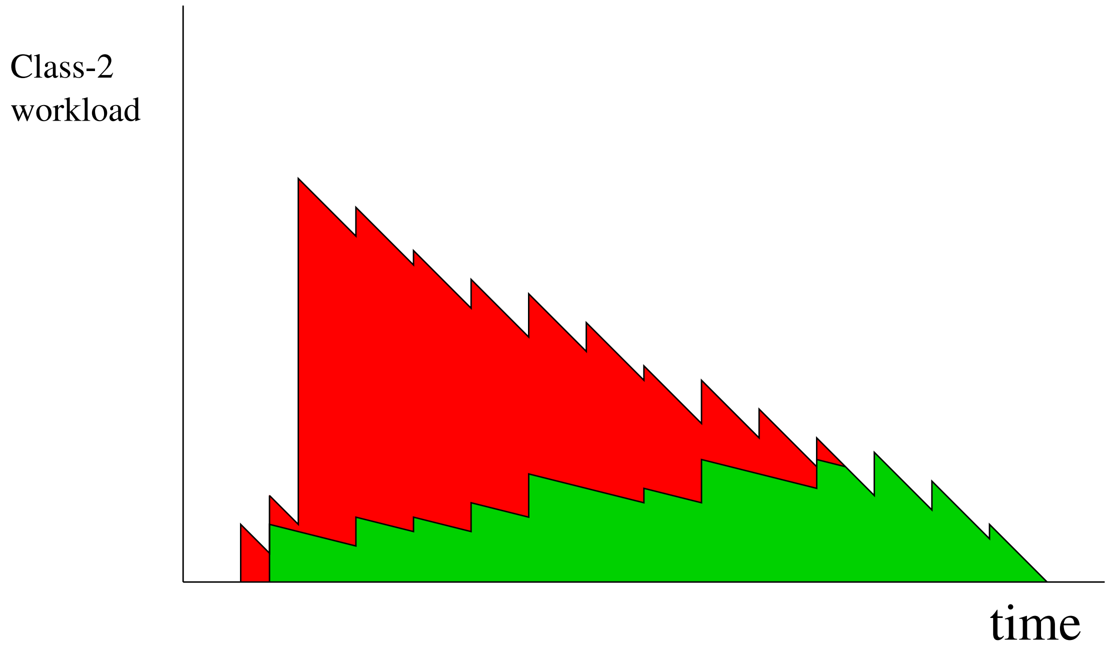
Class-2
workload
time

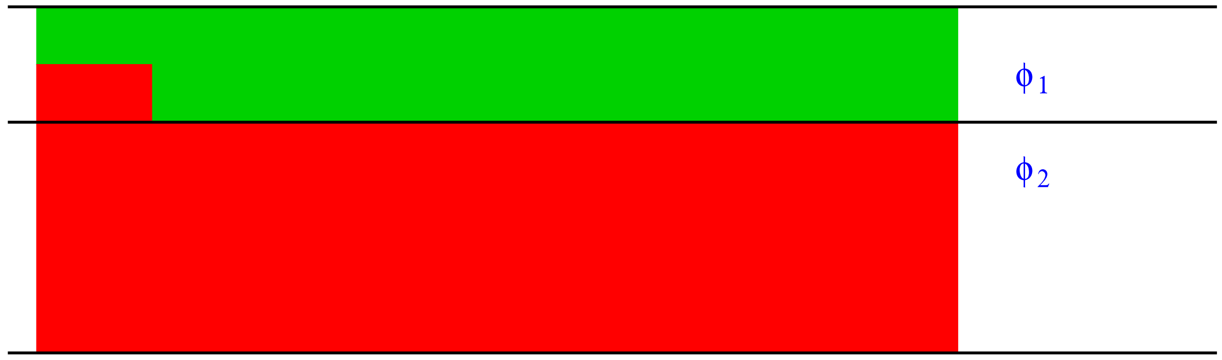
ϕ_1
ϕ_2

Theorem

If A_2^r or B_2^r is regularly varying, $\rho_1 > \phi_1$, and $\rho_2 < \phi_2$, then

$$\mathbb{P}\{\mathbf{V}_2^{GPS} > x\} \sim \mathbb{P}\{\mathbf{V}_2^{\phi_2} > x\} \qquad \text{as } x \to \infty$$

Reduced-weight equivalence (RWE):
Class-2 workload roughly behaves as in isolated system with rate ϕ_2

Qualitatively similar to reduced-*load* equivalence in previous case

Note: *independent* of class-1 traffic characteristics

Class-2 workload behavior (cont'd)

Case III: $\rho_1 < \phi_1$, $\rho_2 > \phi_2$

Catastrophe scenario:

- **Class 2 generates large burst (or long On-period)**
- **Class 1 generates traffic at rate $\rho_1 < \phi_1$**
- **Class 2 is effectively served at rate $1 - \rho_1$**

Theorem

If A_2^r **or** B_2^r **is regularly varying,** $\rho_1 < \phi_1$**, and** $\rho_2 > \phi_2$**, then**

$$\mathbb{P}\{\mathbf{V}_2^{GPS} > x\} \sim \mathbb{P}\{\mathbf{V}_2^{1-\rho_1} > x\} \qquad \textbf{as } x \to \infty$$

Reduced-load equivalence (RLE):
Class-2 workload roughly behaves as in isolated system with rate $1 - \rho_1$

Qualitatively similar as in previous two cases

However, in contrast to previous two cases, now *it is* crucial that class-1 traffic is 'lighter'-tailed than class-2 traffic

Class-1 workload

Case I: $\rho_1 > \phi_1$, $\rho_2 < \phi_2$

Catastrophe scenario:

- **Class 2 generates large burst (or long On-period)**
- **Enters long busy period, and claims service rate ϕ_2 for duration of busy period**
- **Leaves only service rate $\phi_1 = 1 - \phi_2$ for class 1**
- **Class 1 generates traffic at rate $\rho_1 > \phi_1$**
- **Class-1 workload builds up at rate $\rho_1 - \phi_1 > 0$ for duration of class-2 busy period**

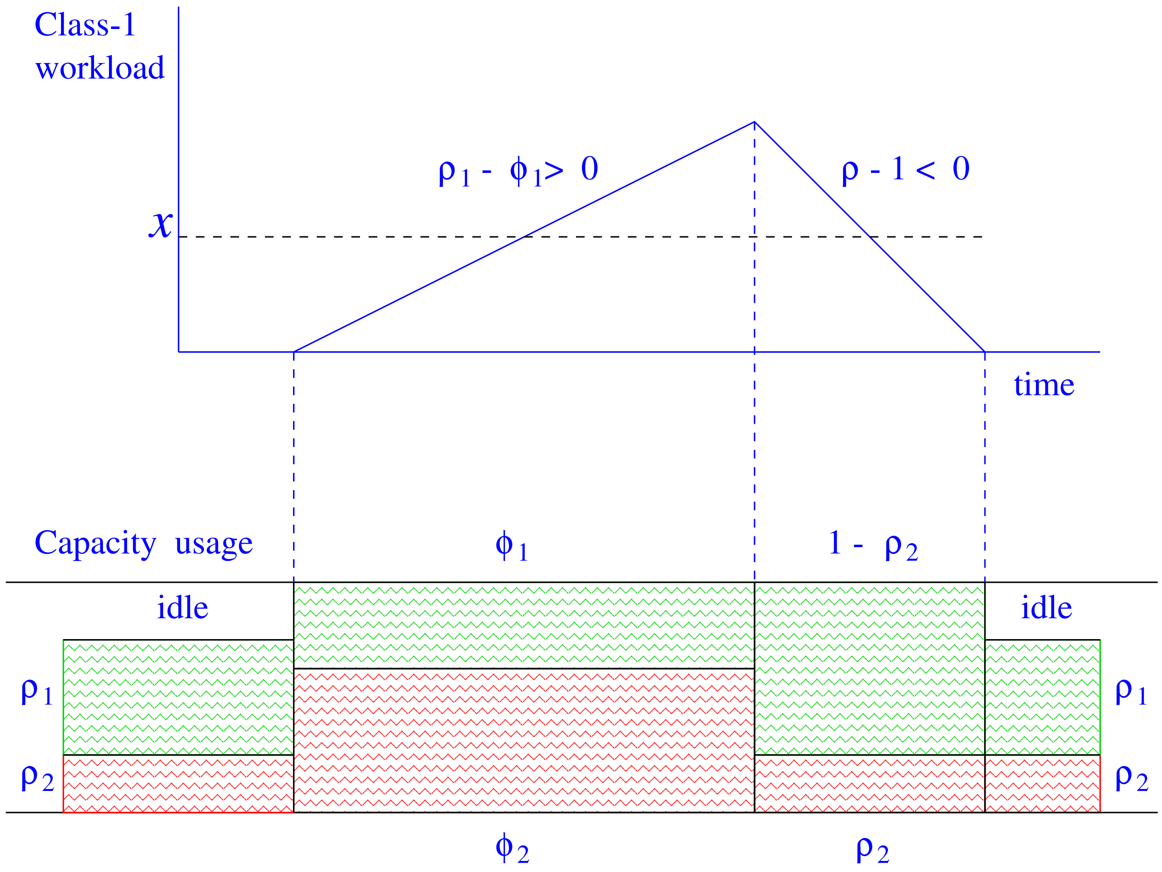
Class-1
workload
ρ1 - φ1> 0
ρ - 1 < 0
x
time
Capacity usage
φ1
1 - ρ2
idle
idle
ρ1
ρ2
ρ1
ρ2
φ2
ρ2

Theorem

If B_2^r is regularly varying, $\rho_1 > \phi_1$ and $\rho_2 < \phi_2$, then

$$\mathbb{P}\{\mathbf{V}_1^{GPS} > x\} \sim \frac{\phi_2 - \rho_2}{\phi_2} \frac{\rho_2}{1 - \rho_1 - \rho_2} \mathbb{P}\{\mathbf{P}_2^r > \frac{x}{\rho_1 - \phi_1}\},$$

with P_2^r residual class-2 busy period when served at rate ϕ_2

Induced burstiness (IB):
Class-1 workload behaves as that of heavy-tailed On-Off process with as On-periods the class-2 busy periods, and inherits ill-behaved class-2 characteristics

Class-1 workload behavior (cont'd)

Case II: $\rho_1 < \phi_1$, $\rho_2 < \phi_2$

Class 1 remains stable even when class 2 is backlogged, so previous catastrophe scenario can no longer occur

Class 1 too must show abnormal activity in order for large workload to build up

Recall class 1 in isolation builds up large workload in gradual manner by deviating from its normal traffic intensity for long period

Conspiracy scenario:

- **Class 1 shows similar abnormal activity as in isolation, raising its traffic intensity to $\widehat{\rho}_1 > \phi_1$ for period $\frac{x}{\widehat{\rho}_1 - \phi_1}$**

- **During that period, class 2 remains constantly backlogged, leaving service rate $\phi_1 = 1 - \phi_2$ for class 1**

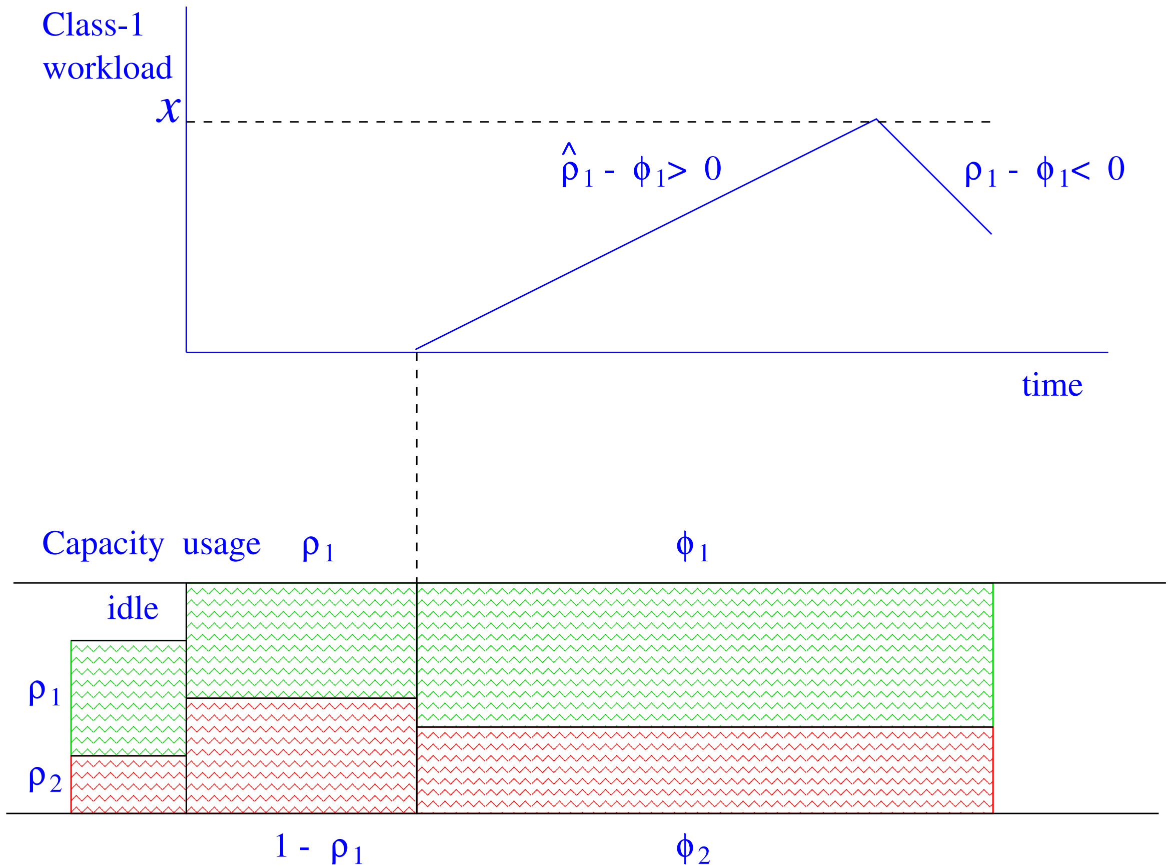
Class-1
workload
x
$\hat{\rho}_1 - \phi_1 > 0$
$\rho_1 - \phi_1 < 0$
time
Capacity usage
ρ_1
ϕ_1
idle
ρ_1
ρ_2
$1 - \rho_1$
ϕ_2

Theorem

If B_2^r is regularly varying, $\rho_1 < \phi_1$ and $\rho_2 < \phi_2$, then

$$\mathbb{P}\{\mathbf{V}_1^{GPS} > x\} \sim \mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\}\mathbb{P}\{\mathrm{T}_2 > \frac{x}{\hat{\rho}_1 - \phi_1}\},$$

with T_2 'drain' time of class 2 when served at rate ϕ_2 with initial workload $\mathrm{V}_2^{1-\rho_1}$

Reduced-weight equivalence (RWE):
but now *major* contribution from deviant class-2 behavior

Similar behavior has been shown for total workload in queues fed by mixture of heavy-tailed and light-tailed input [B & Zwart (2000)] and various related models [Boxma, Deng & Zwart (2002), Boxma & Kurkova (2000)]

$$\mathbb{P}\{\mathbf{V}_1^{GPS} > x\} \sim \mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\}\mathbb{P}\{\mathbf{T}_2 > \frac{x}{\hat{\rho}_1 - \phi_1}\}$$

First term represents upper bound for class 1 based on minimum-rate guarantee ϕ_1,
and captures deviant behavior of class 1 itself

Second term reflects that class 2 must remain backlogged long enough for class-1 workload to build up,
and provides measure for gains from sharing surplus capacity with class 2

$$\mathbb{P}\{\mathbf{V}_1^{GPS} > x\} \sim \mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\}\mathbb{P}\{\mathbf{T}_2 > \frac{x}{\hat{\rho}_1 - \phi_1}\}$$

First term represents upper bound for class 1 based on minimum-rate guarantee ϕ_1,
and captures deviant behavior of class 1 itself

Second term reflects that class 2 must remain backlogged long enough for class-1 workload to build up,
and provides measure for gains from sharing surplus capacity with class 2

General decompositional form holds irrespective of detailed traffic characteristics of two classes

Specific form of two terms however *does* depend on detailed properties, in particular whether class 2 generates instantaneous or fluid input

Instantaneous input

$$\mathbb{P}\{\mathbf{T}_2 > x\} \sim \frac{\rho_1}{1 - \rho_1 - \rho_2} \mathbb{P}\{\mathbf{B}_2^r > (\phi_2 - \rho_2)x\}$$

Class 2 must remain backlogged for period of length x

Normally generates traffic at rate ρ_2

Receives service at rate ϕ_2 while class-1 workload builds up

Instantaneous input (cont'd)

Class 2 needs to make up for 'deficit' amount $(\phi_2 - \rho_2)x$

Enjoys service at rate $1 - \rho_1$ **before that**

Most likely scenario: initial $\mathrm{V}_2^{1-\rho_1}$ **exceeds** $(\phi_2 - \rho_2)x$ **(due to earlier large burst), which occurs with probability**

$$\mathbb{P}\{\mathbf{V}_2^{1-\rho_1} > (\phi_2 - \rho_2)x\} \sim \frac{\rho_2}{1 - \rho_1 - \rho_2}\mathbb{P}\{\mathbf{B}_2^r > (\phi_2 - \rho_2)x\}$$

Fluid input

Similar yet slightly more involved scenario

Instantaneous input (cont'd)

Class 2 needs to make up for 'deficit' amount $(\phi_2 - \rho_2)x$

Enjoys service at rate $1 - \rho_1$ before that

Most likely scenario: initial $V_2^{1-\rho_1}$ exceeds $(\phi_2 - \rho_2)x$ (due to earlier large burst), which occurs with probability

$$\mathbb{P}\{\mathbf{V}_2^{1-\rho_1} > (\phi_2 - \rho_2)x\} \sim \frac{\rho_2}{1-\rho_1-\rho_2}\mathbb{P}\{\mathbf{B}_2^r > (\phi_2 - \rho_2)x\}$$

Fluid input

Similar yet slightly more involved scenario

Class-1 workload behavior (cont'd)

Case III: $\rho_1 < \phi_1$, $\rho_2 > \phi_2$

Now class 2 remains constantly backlogged with probability O(1) while class-1 workload builds up

$$\mathbb{P}\{\mathbf{V}_1^{GPS} > x\} \sim K_2 \mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\} \qquad \textbf{as } x \to \infty$$

Constant K_2 is difficult to determine

Reduced-weight equivalence (RWE):
but now *minor* contribution from deviant class-2 behavior

Discussion & conclusion

Various scenarios for qualitative behavior

- Reduced-load equivalence (RLE):
  class receives total rate reduced by load of other class

- Reduced-weight equivalence – no effort (RWE-0):
  class gets total rate reduced by weight of other class;
  other class shows average behavior (prob. 1)

- Reduced-weight equivalence – minor effort (RWE-1):
  class gets total rate reduced by weight of other class;
  other class shows minor deviant behavior (prob. O(1))

- **Reduced-weight equivalence – major effort (RWE-2):** class gets total rate reduced by weight of other class; other class shows major deviant behavior (prob. o(1))

- **Induced burstiness (IB):** class affected by other class, and inherits ill-behaved traffic characteristics

Classification of qualitative behavior

Qualitative behavior Q_1	$\rho_1 < \phi_1$ $\rho_2 < \phi_2$	$\rho_1 < \phi_1$ $\rho_2 > \phi_2$	$\rho_1 > \phi_1$ $\rho_2 < \phi_2$	$\rho_1 > \phi_1$ $\rho_2 > \phi_2$
Q_1 **HT,** Q_2 **LT**	**RLE** $\uparrow$	**RWE-0** $\uparrow$	**RLE** $\uparrow$	**unstable**
Q_1 **HT,** Q_2 **HT** Q_1 **'heavier' than 2**	**RLE** $\uparrow$	**RWE-0** $\uparrow$	**RLE**	**unstable**
Q_1 **HT,** Q_2 **HT** Q_1 **'lighter' than 2**	**RLE**	**RWE-0**	**IB** $\downarrow$	**unstable**
Q_1 **LT,** Q_2 **HT**	**RWE-2**	**RWE-1**	**IB**	**unstable**

References

Agrawal, R., Makowski, A.M., Nain, Ph. (1999). On a reduced load equivalence for fluid queues under subexponentiality. Queueing Systems 33, 5–41.

Anantharam, V. (1999). Scheduling strategies and long-range dependence. Queueing Systems 33, 73–89.

Arvidsson, A. Karlsson, P. (1999). On traffic models for TCP/IP. In: Teletraffic Engineering in a Competitive World, Proc. ITC-16, Edinburgh, UK, eds. P. Key, D. Smith (North-Holland, Amsterdam), 457–466.

Bertsimas, D., Paschalidis, I.Ch., Tsitsiklis, J.N. (1999). Large deviations analysis of the Generalized Processor Sharing policy. Queueing Systems 32, 319–349.

Borst, S.C., Boxma, O.J., Jelenković, P.R. (1999). Generalized Processor Sharing with long-tailed traffic sources. In: Teletraffic Engineering in a Competitive World, Proc. ITC-16, Edinburgh, UK, eds. P. Key, D. Smith (North-Holland, Amsterdam), 345–354.

Borst, S.C., Boxma, O.J., Jelenković, P.R. (1999). Induced burstiness in Generalized Processor Sharing queues with long-tailed traffic flows. In: Proc. 37th Annual Allerton Conference on Communication, Control, and Computing, 316–325.

Borst, S.C., Boxma, O.J., Jelenković, P.R. (2000). Asymptotic behavior of Generalized Processor Sharing with long-tailed traffic sources. In: Proc. Infocom 2000, Tel-Aviv, Israel, 912–921.

Borst, S.C., Boxma, O.J., Jelenković, P.R. (2000). Coupled processors with regularly varying service times. In: Proc. Infocom 2000, Tel-Aviv, Israel, 157–164.

Borst, S.C., Boxma, O.J., Jelenković, P.R. (2003). Reduced-load equivalence and induced burstiness in GPS queues with long-tailed traffic flows. Queueing Systems 43, 273–306.

Borst, S.C., Boxma, O.J., Van Uitert, M.J.G. (2001). Two coupled queues with heterogeneous traffic. In: Teletraffic Engineering in the Internet Era, Proc. ITC-17, Salvador da Bahia, eds. J.M. de Souza, N.L.S. da Fonseca, E.A. de Souza e Silva (North-Holland, Amsterdam), 1003–1014.

Borst, S.C., Boxma, O.J., Van Uitert, M.J.G. (2003). The asymptotic workload behavior of two coupled queues. Queueing Systems 43, 81–102.

Borst, S.C., Mandjes, M., Van Uitert, M.J.G. (2002). Generalized Processor Sharing queues with heterogeneous traffic classes. Adv. Appl. Prob., to appear. Shortened version in: Proc. Infocom 2002, New York, USA, 74–83.

Borst, S.C., Zwart, A.P. (2000). A reduced-peak equivalence for queues with a mixture of light-tailed and heavy-tailed input flows. SPOR-Report 2000-04, Eindhoven University of Technology. Adv. Appl. Prob., to appear.

Boxma, O.J., Deng, Q., Zwart, A.P. (2002). Waiting-time asymptotics for the M/G/2 queue with heterogeneous servers. Queueing Systems 40, 5–31.

Boxma, O.J., Kurkova, I.A. (2000). The M/M/1 queue in a heavy-tailed random environment. Statistica Neerlandica 54, 221–236.

Cohen, J.W. (1973). Some results on regular variation for distributions in queueing and fluctuation theory. J. Appl. Prob. 10, 343–353.

Delas, S., Mazumdar, R.R., Rosenberg, C.P. (2002). Tail asymptotics for HOL priority queues handling a large number of independent stationary sources. Queueing Systems 40, 183–204.

Dumas, V., Simonian, A. (2000). Asymptotic bounds for the fluid queue by subexponential on/off sources. Adv. Appl. Prob. 32, 244–255.

Dupuis, P., Ramanan, K. (1998). A Skorokhod problem formulation and large deviation analysis of a processor sharing model. Queueing Systems 28, 109–124.

Elwalid, A.I., Mitra, D. (1999). Design of Generalized Processor Sharing schedulers which statistically multiplex heterogeneous QoS classes. In: Proc. Infocom '99, New York, USA, 1220–1230.

Fayolle, G., Iasnogorodski, R. (1979). Two coupled processors: the reduction to a Riemann-Hilbert problem. Z. Wahr. verw. Geb. 47, 325–351.

Fayolle, G., Mitrani, I., Iasnogorodski, R. (1980). Sharing a processor among many job classes. J. ACM 27, 519–532.

Guillemin, F., Mazumdar, R.R., Dupuis, A., Boyer, J. (2003). Analysis of the fluid weighted fair queueing system. J. Appl. Prob. 40, 180–199.

Jelenković, P.R., Lazar, A.A. (1999). Asymptotic results for multiplexing subexponential on-off processes. Adv. Appl. Prob. 31, 394–421.

Jelenković, P.R., Momčilović, P. (2001). Network multiplexer with Generalized Processor Sharing and heavy-tailed on-off flows. In: Teletraffic Engineering in the Internet Era, Proc. ITC-17, Salvador da Bahia, eds. J.M. de Souza, N.L.S. da Fonseca, E.A. de Souza e Silva (North-Holland, Amsterdam), 719–730.

Jelenković, P.R., Momčilović, P. (2002). Finite buffer queue with Generalized Processor Sharing and heavy-tailed input processes. Comp. Netw. 40, 433–443.

Kotopoulos, C., Likhanov, N., Mazumdar, R.R. (2001). Asymptotic analysis of the GPS system fed by heterogeneous long-tailed sources. In: Proc. Infocom 2001, Anchorage AK, USA, 299–308.

Kotopoulos, C., Mazumdar, R.R. (2002). Buffer occupancy and delay asymptotics in multi-buffered systems with Generalized Processor Sharing handling a large number of independent traffic streams. Preprint, Purdue University.

Kotopoulos, C., Mazumdar, R.R. (2002). Many sources asymptotics for a 2-buffer system with Generalized Processor Sharing. Preprint, Purdue University.

Kumaran, K., Margrave, G.E., Mitra, D., Stanley, K.R. (2000). Novel techniques for the design and control of Generalized Processor Sharing schedulers for multiple QoS classes. In: Proc. Infocom 2000, Tel-Aviv, Israel, 932–941.

Lelarge, M. (2001). Asymptotic behavior of Generalized Processor Sharing queues under subexponential hypothesis. Report RR-4339, INRIA Rocquencourt.

Mandjes, M.R.H., Van Uitert, M.J.G. (2003). Large deviations for complex buffer architectures: the short-range dependent case. Work in progress.

Mandjes, M.R.H., Van Uitert, M.J.G. (2003). Sample-path large deviations for Generalized Processor Sharing queues with Gaussian inputs. CWI Report.

Mannersalo, P, Norros, I. (2002). GPS schedulers and Gaussian traffic. In: Proc. Infocom 2002, New York, USA, 1660–1667.

Massoulié, L. (1999). Large deviations estimates for polling and weighted fair queueing service systems. Adv. Perf. Anal. 2, 103–128.

De Meyer, A., Teugels, J.L. (1980). On the asymptotic behaviour of the distribution of the busy period and service time in M/G/1. J. Appl. Prob. 17, 802–813.

Pakes, A.G. (1975). On the tails of waiting-time distributions. J. Appl. Prob. 12, 555–564.

Parekh, A.K., Gallager, R.G. (1993). A Generalized Processor Sharing approach to flow control in integrated services networks: the single-node case. IEEE/ACM Trans. Netw. 1, 344–357.

Parekh, A.K., Gallager, R.G. (1994). A Generalized Processor Sharing approach to flow control in integrated services networks: the multiple node case. IEEE/ACM Trans. Netw. 2, 137–150.

Paschalidis, I.Ch. (1999). Class-specific quality of service guarantees in multimedia communication networks. Automatica 35, 1951–1968.

Pereira, F.M., Fonseca, N.L.S., Arantes, D.S. (2002). On the performance of Generalized Processor Sharing servers under long-range dependent traffic. Comp. Netw. 40, 413–431.

Toutain, F. (1998). Decoupled Generalized Processor Sharing: a fair queueing principle for adaptive multimedia applications. In: Proc. Infocom '98, San Francisco, USA, 291–298.

Van Uitert, M.J.G. (2003). Generalized Processor Sharing queues. Ph.D. thesis, Eindhoven University of Technology.

Van Uitert, M.J.G., Borst, S.C. (2000). A reduced-load equivalence for Generalised Processor Sharing networks with long-tailed input flows. Queueing Systems 41, 123–163. Shortened version in: Proc. Infocom 2001, Alaska, 269–278.

Yaron, O., Sidi, M. (1994). Generalized Processor Sharing networks with exponentially bounded burstiness arrivals. Proc. Infocom '94, Toronto, Canada, 628–634.

Yaron, O., Sidi, M. (1994). Generalized Processor Sharing networks with exponentially bounded burstiness arrivals. J. High Speed Netw. 3, 375–387.

Yu, X., Thng, I. L.-J., Jiang, Y., Qiao, C. (2003). Individual queueing process in GPS and PGPS with LRD traffic inputs. Preprint.

Zhang, Z.-L. (1996). Large deviations and the Generalized Processor Sharing scheduling for a two-queue system. Queueing Systems 26, 229–245.

Zhang, Z.-L. (1998). Large deviations and the Generalized Processor Sharing scheduling for a multiple-queue system. Queueing Systems 28, 349–376.

Zhang, Z.-L., Liu, Z., Towsley, D., Kurose, J. (1997). Call admission control under Generalized Processor Sharing. T-elecommunication Systems 7, 125–152.

Zhang, Z.-L., Towsley, D., Kurose, J. (1995). Statistical analysis of the Generalized Processor Sharing discipline. IEEE J. Sel. Areas Commun. 13, 1071–1080.

Zwart, A.P. (1999). Tail asymptotics for the busy period in the GI/G/1 queue. Math. Oper. Res. 26, 485–493.

Zwart, A.P., Borst, S.C., Mandjes, M. (2000). Exact asymptotics for fluid queues fed by multiple heavy-tailed On-Off flows. Ann. Appl. Prob., to appear. Shortened version in: Proc. Infocom 2001, Anchorage AK, USA, 279–288.