Comparison of bandwidth-sharing policies in a linear network

Maaike Verloop (CWI)
Urtzi Ayesta (LAAS-CNRS)
and Sem Borst (TU/e, Bell Labs)

ValueTools, 21 October 2008

Bandwidth-sharing networks

- Flow-level modeling of elastic data transfers over the Internet
- Data flows traverse several links on the path from their source to destination
- Link is represented by a node

View network at flow level

- A flow gets simultaneously the same bandwidth in all links along its path

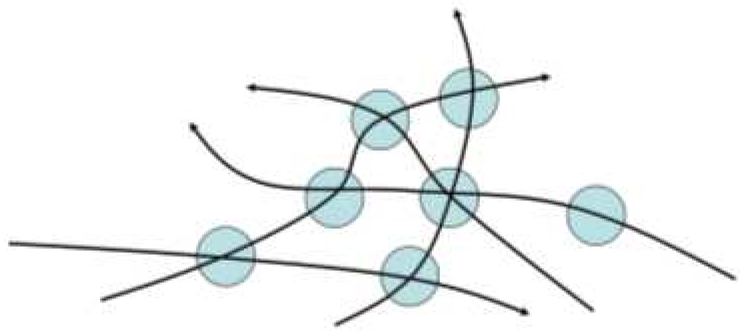

.

Bandwidth-sharing networks

- Flow-level modeling of elastic data transfers over the Internet
- Data flows traverse several links on the path from their source to destination
- Link is represented by a node

View network at flow level

- A flow gets simultaneously the same bandwidth in all links along its path

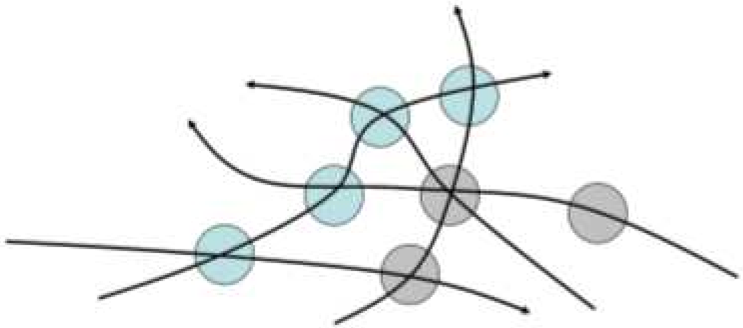

Concentrate on one flow path

Model description: linear network

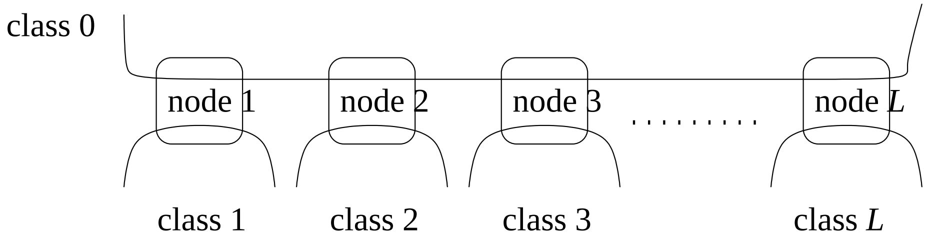

- Class-i users arrive according to a renewal process with mean inter-arrival time $\frac{1}{\lambda_i}$, $i = 0, \ldots, L$
- Generally distributed service requirements with mean $\frac{1}{\mu_i}$
- Load of class i: $\rho_i = \lambda_i / \mu_i$
- $N_i^\pi(t)$: number of class-i users at time t under policy π and $W_i^\pi(t)$: amount of work in class i at time t under policy π
- $s_i^\pi(t, \vec{N}) \in R(t)$ capacity given to class i at time t

$$R(t) = \{\vec{s} : s_0 + s_i \leq C_i(t), \ \ i = 1, \ldots, L\}$$

Bandwidth-sharing mechanisms

A policy π determines how to allocate the capacity of the links to all flows present in the network

Internet: TCP determines implicit rate allocation through congestion control

Some allocations: Weighted α-fair policies, max-min, proportional fair, maximum throughput etc.

- Stability is ensured whenever possible (for $\alpha > 0$, exponentially distributed service requirements and fixed capacities)
- Metrics like delay, throughput, number of users in the system difficult to determine in general

Objective of the talk: Compare the performance under various bandwidth-sharing policies in the linear network

Outline of the talk

- Linear network
  - sample-path comparison result of policies
  - performance: stability and mean number of users
- Special case: single node with two classes
  - monotonicity with respect to the weights for DPS and GPS
- Weighted α-fair policies
  - monotonicity results w.r.t. fairness parameter α and the weights: stability and mean number of users
  - heavy-traffic regime
  - numerical results
- Conclusion and future work

Sample-path comparison result

Property: Let π and $\tilde{\pi}$ be two policies such that

$$s_0^{\pi}(\vec{N}) \leq s_0^{\tilde{\pi}}(\vec{N}),$$

and either $s_0^{\tilde{\pi}}(\vec{N})$ or $s_0^{\tilde{\pi}}(\vec{N})$ is non-increasing w.r.t. $N_i, \ i \neq 0$.

The property states that higher priority is given to class 0 under policy $\tilde{\pi}$ compared to π.

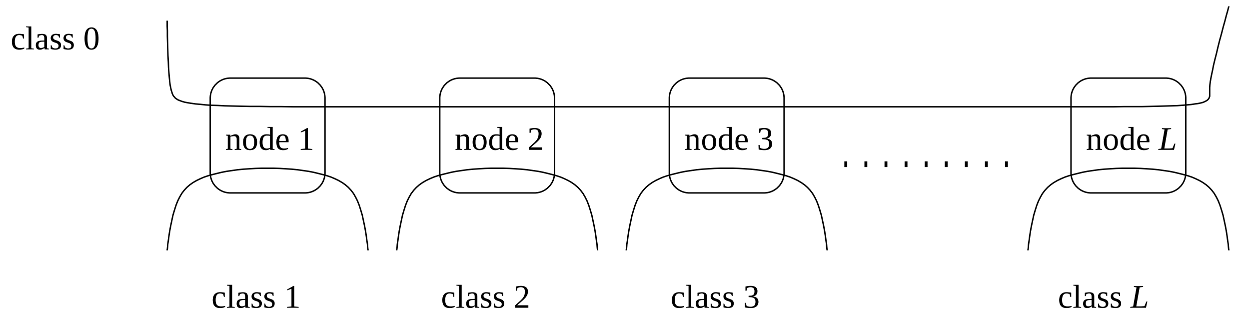

Sample-path comparison result

Property: Let π and $\tilde{\pi}$ be two policies such that

$$s_0^{\pi}(\vec{N}) \leq s_0^{\tilde{\pi}}(\vec{N}),$$

and either $s_0^{\pi}(\vec{N})$ or $s_0^{\tilde{\pi}}(\vec{N})$ is non-increasing w.r.t. $N_i, \;\; i \neq 0$.

Sample-path comparison result

Property: Let π and $\tilde{\pi}$ be two policies such that

$$s_0^{\pi}(\vec{N}) \leq s_0^{\tilde{\pi}}(\vec{N}),$$

and either $s_0^{\pi}(\vec{N})$ or $s_0^{\tilde{\pi}}(\vec{N})$ is non-increasing w.r.t. $N_i, \ i \neq 0$.

Proposition:

- Consider the same realizations of the arrival processes and service requirements.
- Intra-class policy is FCFS.

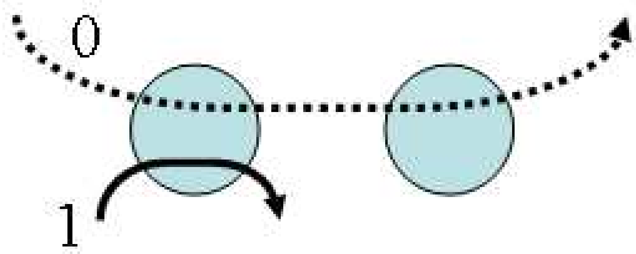

If $\vec{W}^{\pi}(0) = \vec{W}^{\tilde{\pi}}(0)$, then

i) $N_0^{\pi}(t) \geq N_0^{\tilde{\pi}}(t)$ and $W_0^{\pi}(t) \geq W_0^{\tilde{\pi}}(t)$,

ii) $W_0^{\pi}(t) + W_i^{\pi}(t) \geq W_0^{\tilde{\pi}}(t) + W_i^{\tilde{\pi}}(t)$

Relation with stochastic comparison results in [Massey87]

The necessary and sufficient conditions on the policies π and $\tilde{\pi}$ to obtain

$$\{N_0^{\pi}(t)\}_{t\geq 0} \geq_{st} \{N_0^{\tilde{\pi}}(t)\}_{t\geq 0},$$

for any two ordered initial states $N_0^{\pi}(0) \geq N_0^{\tilde{\pi}}(0)$, are

$$s_0^{\pi}(\vec{N}^{\pi}) \leq s_0^{\tilde{\pi}}(\vec{N}^{\tilde{\pi}}) \text{ when } N_0^{\pi} = N_0^{\tilde{\pi}}.$$

For bandwidth-sharing policies in the linear network this often does not hold. Since if $N_i^{\tilde{\pi}} \to \infty$, $i \neq 0$, then $s_0^{\tilde{\pi}}(\vec{N}^{\tilde{\pi}}) \to 0$.

Performance

From the sample-path comparison result, we obtain the following

Proposition:

1. Stability: If policy π is stable, then policy $\tilde{\pi}$ is stable
2. Exponentially distributed service requirements:

   If
   $$\sum_{i=1}^{L} c_i \mu_i \leq c_0 \mu_0,$$
   then
   $$\sum_{i=0}^{L} c_i \mathbb{E}(N_i^{\pi}(t)) \geq \sum_{i=0}^{L} c_i \mathbb{E}(N_i^{\tilde{\pi}}(t)), \quad \forall\, t \geq 0.$$

Application: Single node with time-varying capacity and two classes

- GPS allocation

$$s_i^{GPS(\phi)}(\vec{N}) = C(t) \cdot \frac{\phi_i}{\phi_1 \mathbf{1}_{(N_1>0)} + \phi_2 \mathbf{1}_{(N_2>0)}}, \quad i = 1,2$$

Intra-class policy is FCFS

- DPS allocation

$$s_i^{DPS(\phi)}(\vec{N}) = C(t) \cdot \frac{\phi_i N_i}{\phi_1 N_1 + \phi_2 N_2}, \quad i = 1,2$$

Intra-class policy is PS

Remark: For exponentially distributed service requirements the stochastic behavior of the system does not depend on the non-anticipating policy (like PS, FCFS, . . .) chosen.

Application: DPS and GPS

Proposition: Let $\phi_1 < \tilde{\phi}_1$. Consider the same realizations of the arrival processes and service requirements for both processes.

- GPS: for generally distributed service requirements:

$$W_1^{GPS(\phi)}(t) \geq W_1^{GPS(\tilde{\phi})}(t) \text{ and } N_1^{GPS(\phi)}(t) \geq N_1^{GPS(\tilde{\phi})}(t).$$

The opposite inequalities hold for class 2.

Application: DPS and GPS

Proposition: Let $\phi_1 < \tilde{\phi}_1$. Consider the same realizations of the arrival processes and service requirements for both processes.

- GPS: for generally distributed service requirements:

  $$W_1^{GPS(\phi)}(t) \geq W_1^{GPS(\tilde{\phi})}(t) \text{ and } N_1^{GPS(\phi)}(t) \geq N_1^{GPS(\tilde{\phi})}(t).$$

- DPS: for exponentially distributed service requirements:

  $$\{W_1^{DPS(\phi)}(t)\}_t \geq_{st} \{W_1^{DPS(\tilde{\phi})}(t)\}_t, \text{ and}$$

  $$\{N_1^{DPS(\phi)}(t)\}_t \geq_{st} \{N_1^{DPS(\tilde{\phi})}(t)\}_t.$$

The opposite inequalities hold for class 2.

Application: DPS and GPS (cont)

Proposition: Let $\phi_1 < \tilde{\phi}_1$.

Assume exponentially distributed service requirements with $c_1\mu_1 \geq c_2\mu_2$.

Then for all $t \geq 0$

$$\sum_{i=1}^{2} c_i \mathbb{E}(N_i^{GPS(\phi)}(t)) \geq \sum_{i=1}^{2} c_i \mathbb{E}(N_i^{GPS(\tilde{\phi})}(t))$$

and

$$\sum_{i=1}^{2} c_i \mathbb{E}(N_i^{DPS(\phi)}(t)) \geq \sum_{i=1}^{2} c_i \mathbb{E}(N_i^{DPS(\tilde{\phi})}(t))$$

Application to linear network: Weighted α-fair policies

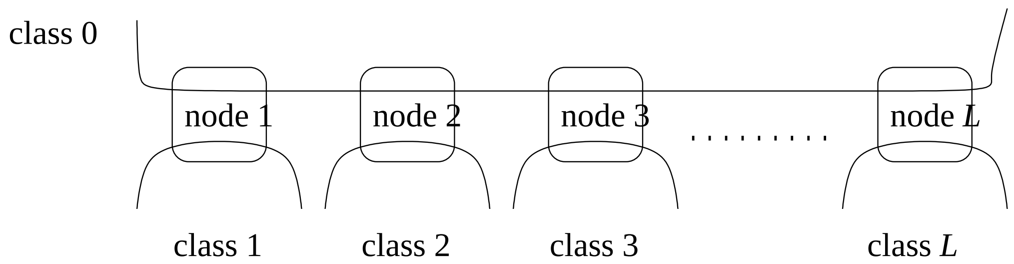

- The weighted-α fair allocation is the solution to the following optimization problem:

$$\begin{array}{ll} \max_{\vec{s}\in R(t)} \sum_{i=0}^{L} w_i N_i \left(\frac{s_i}{N_i}\right)^{1-\alpha} / (1-\alpha) & \text{if } \alpha \neq 1 \\ \max_{\vec{s}\in R(t)} \sum_{i=0}^{L} w_i N_i \log s_i & \text{if } \alpha = 1. \end{array}$$

- Intra-class policy is Processor Sharing.
  Assume exponentially distributed service requirements.
  $\longrightarrow$ results for general setting with FCFS can be used

Weighted α-fair policies (cont.)

Varying α we obtain: $\alpha \to 0$ maximum throughput, $\alpha = 1$ Proportional Fairness, $\alpha \to \infty$ max-min fairness, $\alpha = 2$ approximates TCP.

- There exist congestion control algorithms that realize α-fair policies in a decentralized way
- For fixed capacities and $\alpha > 0$, stability is ensured whenever possible [BM01]

Application: Weighted α-fair policies (cont)

For a given α and weights $w = (w_0, w_1, \ldots, w_L)$ denote the policy by $\pi^{\alpha,w}$ and the allocation vector by $\vec{s}^{(\alpha,w)}(\vec{N})$.

We have

(i) $s_0^{(\alpha,w)}(\vec{N})$ is non-increasing in N_i, $i = 1, ..., L$

(ii) $s_0^{(\beta,w)}(\vec{N}) \leq s_0^{(\gamma,w)}(\vec{N})$, if $\beta \leq \gamma$

(iii) $s_0^{(\alpha,w)}(\vec{N}) \leq s_0^{(\alpha,\tilde{w})}(\vec{N})$, if $\frac{w_0}{w_i} \leq \frac{\tilde{w}_0}{\tilde{w}_i}$

Hence, the property holds for $\pi^{\beta,w}$ and $\pi^{\gamma,\tilde{w}}$, with $\beta \leq \gamma$ and $\frac{w_0}{w_i} \leq \frac{\tilde{w}_0}{\tilde{w}_i}$

We obtain insights into the performance of these policies in linear networks

Stability results for the linear network

Exponentially distributed service requirements

Corollary: Let $\beta \leq \gamma$ and $\frac{w_0}{w_i} \leq \frac{\tilde{w}_0}{\tilde{w}_i}$, $i = 1, \ldots, L$.

If policy $\pi^{\beta,w}$ is stable, then policy $\pi^{\gamma,\tilde{w}}$ is stable.

Stability results for the linear network

Exponentially distributed service requirements

Corollary: Let $\beta \leq \gamma$ and $\frac{w_0}{w_i} \leq \frac{\tilde{w}_0}{\tilde{w}_i}$, $i = 1, \ldots, L$.

If policy $\pi^{\beta,w}$ is stable, then policy $\pi^{\gamma,\tilde{w}}$ is stable.

In [Liu et al., 2007] it is shown that the stability region is decreasing in α ($w_i = 1$).

Stability results for the linear network

Exponentially distributed service requirements

Corollary: Let $\beta \leq \gamma$ and $\frac{w_0}{w_i} \leq \frac{\tilde{w}_0}{\tilde{w}_i}$, $i = 1, \ldots, L$.

If policy $\pi^{\beta,w}$ is stable, then policy $\pi^{\gamma,\tilde{w}}$ is stable.

In [Liu et al., 2007] it is shown that the stability region is decreasing in α ($w_i = 1$).

Corollary: Assume Poisson arrivals and that $(C_1(t), \ldots, C_L(t))$ can be in a finite number of states and evolves as a stationary and ergodic process, with $\overline{C}_i$ the average of $C_i(t)$.

Policy $\pi^{\alpha,w}$ with $w_i \leq w_0, i = 1, \ldots, L$ is stable if $\rho_0 + \rho_i < \overline{C}_i, \forall i$

Mean number of users: linear network

Proposition: Assume exponentially distributed service requirements with $\sum_{i=1}^{L} c_i \mu_i \leq c_0 \mu_0$. If $\beta \leq \gamma$ and $\frac{w_0}{w_i} \leq \frac{\tilde{w}_0}{\tilde{w}_i}$, then

$$\sum_{i=0}^{L} c_i \mathbb{E}(N_i^{\pi^{\beta,w}}(t)) \geq \sum_{i=0}^{L} c_i \mathbb{E}(N_i^{\pi^{\gamma,\tilde{w}}}(t)), \ \ \forall \, t \geq 0.$$

Mean number of users: linear network

Proposition: Assume exponentially distributed service requirements with $\sum_{i=1}^{L} c_i\mu_i \leq c_0\mu_0$. If $\beta \leq \gamma$ and $\frac{w_0}{w_i} \leq \frac{\tilde{w}_0}{\tilde{w}_i}$, then

$$\sum_{i=0}^{L} c_i \mathbb{E}(N_i^{\pi^{\beta,w}}(t)) \geq \sum_{i=0}^{L} c_i \mathbb{E}(N_i^{\pi^{\gamma,\tilde{w}}}(t)), \quad \forall\, t \geq 0.$$

Natural choice for weights: $c_0 = L,\ c_i = 1,\ i = 1, \ldots, L$.

- The condition $\sum_{i=1}^{L} c_i\mu_i \leq c_0\mu_0$ becomes

$$\frac{1}{L}\sum_{i=1}^{L} \mu_i \leq \mu_0,$$

  i.e. departure rate of class 0 is larger than or equal to the average departure rate of classes $1, \ldots, L$

Uncovered case

If $\sum_{i=1}^{L} c_i\mu_i > c_0\mu_0$, then no such result holds.

Trade-off:

- More preference to classes $1, \ldots, L$
  $\rightarrow$ increases the instantaneous departure rate

  $\rightarrow$ uses the capacity of the network less efficiently

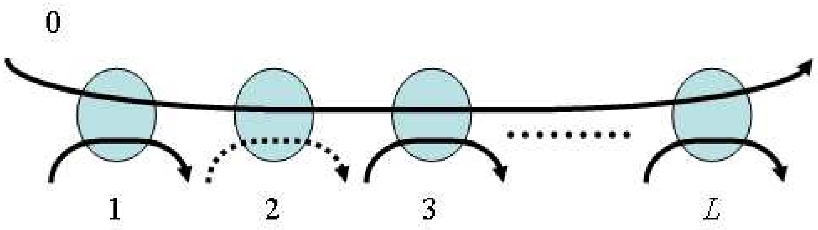

Heavy-traffic regime

Consider the diffusion scaled processes

$$\hat{n}_i^{k,(\alpha)}(t) := \frac{N_i^{\pi^{\alpha,\vec{1}}}(kt)}{\sqrt{k}},\ i = 0, 1, 2$$

$$\hat{v}_i^{k,(\alpha)}(t) = \hat{n}_0^{k,(\alpha)}(t)/\mu_0 + \hat{n}_i^{k,(\alpha)}(t)/\mu_i,\ i = 1, 2.$$

$\hat{v}_i^{k,(\alpha)}(t) \approx$ total workload in node i

Heavy-traffic regime

Consider the diffusion scaled processes

$$\hat{n}_i^{k,(\alpha)}(t) := \frac{N_i^{\pi^{\alpha,\vec{1}}}(kt)}{\sqrt{k}},\ i = 0, 1, 2$$

$$\hat{v}_i^{k,(\alpha)}(t) = \hat{n}_0^{k,(\alpha)}(t)/\mu_0 + \hat{n}_i^{k,(\alpha)}(t)/\mu_i,\ i = 1, 2.$$

$\hat{v}_i^{k,(\alpha)}(t) \approx$ total workload in node i

In [Kang et al, 2007] it is conjectured that $\hat{\vec{v}}^{k,(\alpha)}(t) \xrightarrow{d} \hat{\vec{v}}^{(\alpha)}(t)$, $k \to \infty$, with $\hat{\vec{v}}^{(\alpha)}(t)$ a semimartingale reflecting Brownian motion living in a workload cone.

Heavy-traffic regime

Consider the diffusion scaled processes

$$\hat{n}_i^{k,(\alpha)}(t) \;:=\; \frac{N_i^{\pi^{\alpha,\vec{1}}}(kt)}{\sqrt{k}},\ i = 0, 1, 2$$

$$\hat{v}_i^{k,(\alpha)}(t) \;=\; \hat{n}_0^{k,(\alpha)}(t)/\mu_0 + \hat{n}_i^{k,(\alpha)}(t)/\mu_i,\ i = 1, 2.$$

$\hat{v}_i^{k,(\alpha)}(t) \approx$ total workload in node i

In [Kang et al, 2007] it is conjectured that $\hat{\vec{v}}^{k,(\alpha)}(t) \xrightarrow{d} \hat{\vec{v}}^{(\alpha)}(t)$, $k \to \infty$, with $\hat{\vec{v}}^{(\alpha)}(t)$ a semimartingale reflecting Brownian motion living in a workload cone.

Two-node linear network: workload cone is independent of α

Heavy-traffic regime (cont)

Proposition: Assume $\rho_i + \rho_0 = C_i, \; i = 1, 2.$

- If $c_1\mu_1 + c_2\mu_2 < c_0\mu_0$, then $\mathbb{E}(\sum_{i=0}^{2} c_i \hat{n}_i^{(\alpha)}(t))$ is strictly decreasing in α.
- If $c_1\mu_1 + c_2\mu_2 = c_0\mu_0$, then $\mathbb{E}(\sum_{i=0}^{2} c_i \hat{n}_i^{(\alpha)}(t))$ is constant in α.
- If $c_1\mu_1 + c_2\mu_2 > c_0\mu_0$, then $\mathbb{E}(\sum_{i=0}^{2} c_i \hat{n}_i^{(\alpha)}(t))$ is strictly increasing in α.

$$\sum_{i=0}^{L} c_i \hat{n}_i^{(\alpha)}(t) \stackrel{d}{=} \frac{c_0\mu_0 - \sum_{i=1}^{2} c_i\mu_i}{\mu_0} \hat{n}_0^{(\alpha)}(t) + \sum_{i=1}^{2} c_i \mu_i \hat{v}_i^{(\alpha)}(t).$$

Numerical results I

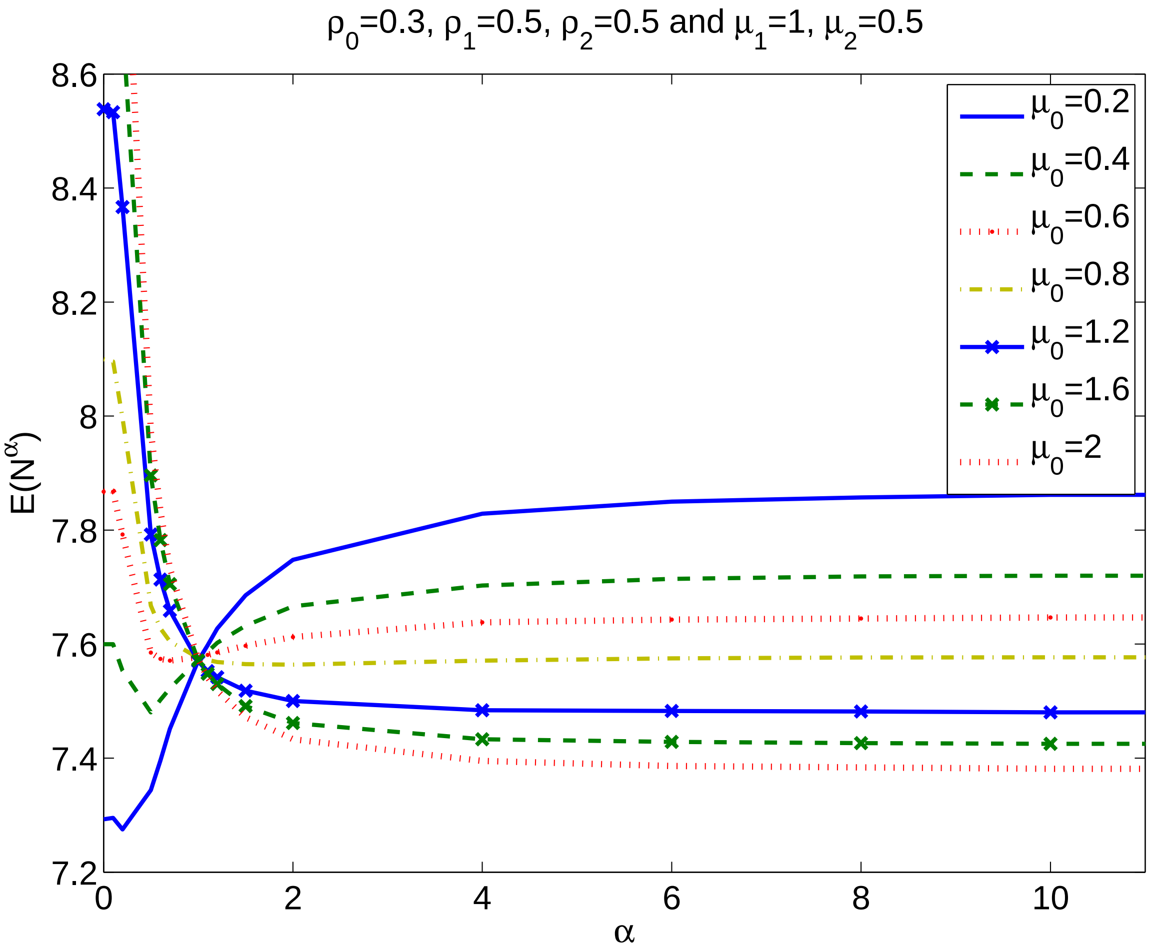

Numerical results II

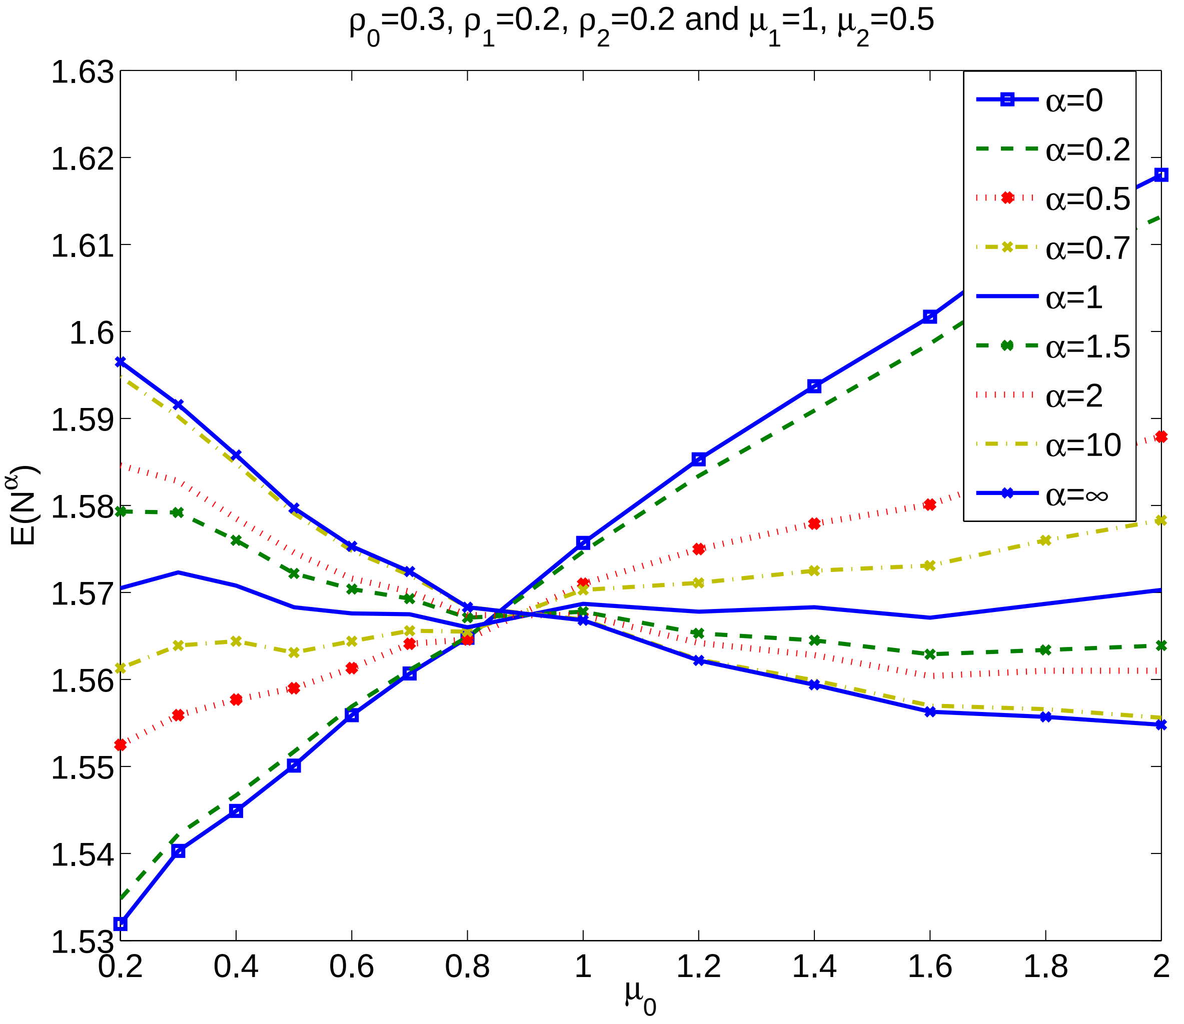

Conclusion and future work

- Single server with more than 2 classes
- Extend to different topologies like star or grid network
- Monotonicity in μ_0.

Comparison of bandwidth-sharing policies in a linear network

Maaike Verloop (CWI)
Urtzi Ayesta (LAAS-CNRS)
and Sem Borst (TU/e, Bell Labs)

ValueTools, 21 October 2008

Intuition

Let $S_i^\pi(t) := \int_{u=0}^{t} s_i^\pi(\vec{N}^\pi(u))\mathrm{d}u$, be the cumulative amount of service received by class i during $[0, t)$.

Property: Class 0 is given more priority under policy π than under $\tilde{\pi}$

(i') $S_0^\pi(t) \leq S_0^{\tilde{\pi}}(t)$

(ii') $S_0^\pi(t) + S_i^\pi(t) \leq S_0^{\tilde{\pi}}(t) + S_i^{\tilde{\pi}}(t)$, since π gives more priority to class 0 and hence makes better use of the available capacity of the network.

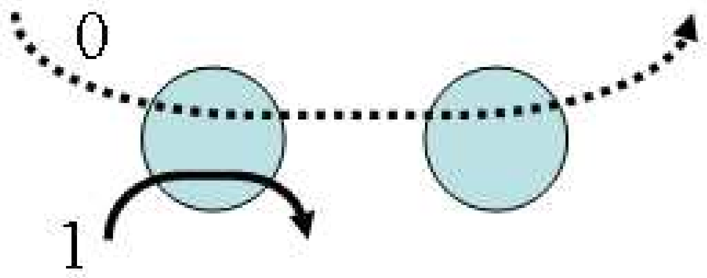

Not trivial:

Giving higher priority to class 0 implies that classes 1 and 2 will contain more users. Hence, class 0 receives less service later on.

Intuition

Let $S_i^\pi(t) := \int_{u=0}^t s_i^\pi(\vec{N}^\pi(u))\mathrm{d}u$, be the cumulative amount of service received by class i during $[0,t)$.

Property: Class 0 is given more priority under policy π than under $\tilde{\pi}$

(i') $S_0^\pi(t) \leq S_0^{\tilde{\pi}}(t)$

(ii') $S_0^\pi(t) + S_i^\pi(t) \leq S_0^{\tilde{\pi}}(t) + S_i^{\tilde{\pi}}(t)$

Intuition

Let $S_i^\pi(t) := \int_{u=0}^t s_i^\pi(\vec{N}^\pi(u))\mathrm{d}u$, be the cumulative amount of service received by class i during $[0,t)$.

Property: Class 0 is given more priority under policy π than under $\tilde{\pi}$

(i') $S_0^\pi(t) \leq S_0^{\tilde{\pi}}(t)$

(ii') $S_0^\pi(t) + S_i^\pi(t) \leq S_0^{\tilde{\pi}}(t) + S_i^{\tilde{\pi}}(t)$

Hence,

(i) $N_0^\pi(t) \geq N_0^{\tilde{\pi}}(t),$ and $W_0^\pi(t) \geq W_0^{\tilde{\pi}}(t)$,

(ii) $W_0^\pi(t) + W_i^\pi(t) \geq W_0^{\tilde{\pi}}(t) + W_i^{\tilde{\pi}}(t)$.

Proof:

- $N_0^{\pi}(t) \geq N_0^{\tilde{\pi}}(t) \quad \longrightarrow \quad \mathbb{E}(N_0^{\pi}(t)) \geq \mathbb{E}(N_0^{\tilde{\pi}}(t))$

Proof:

- $N_0^{\pi}(t) \geq N_0^{\tilde{\pi}}(t) \quad \longrightarrow \quad \mathbb{E}(N_0^{\pi}(t)) \geq \mathbb{E}(N_0^{\tilde{\pi}}(t))$

- $W_0^{\pi}(t) + W_i^{\pi}(t) \geq W_0^{\tilde{\pi}}(t) + W_i^{\tilde{\pi}}(t)$

  $\longrightarrow \mathbb{E}(W_0^{\pi}(t)) + \mathbb{E}(W_i^{\pi}(t)) \geq \mathbb{E}(W_0^{\tilde{\pi}}(t)) + \mathbb{E}(W_i^{\tilde{\pi}}(t))$

Proof:

- $N_0^{\pi}(t) \geq N_0^{\tilde{\pi}}(t) \quad \rightarrow \quad \mathbb{E}(N_0^{\pi}(t)) \geq \mathbb{E}(N_0^{\tilde{\pi}}(t))$
- $W_0^{\pi}(t) + W_i^{\pi}(t) \geq W_0^{\tilde{\pi}}(t) + W_i^{\tilde{\pi}}(t)$

  $\rightarrow \mathbb{E}(W_0^{\pi}(t)) + \mathbb{E}(W_i^{\pi}(t)) \geq \mathbb{E}(W_0^{\tilde{\pi}}(t)) + \mathbb{E}(W_i^{\tilde{\pi}}(t))$
- Intra-class policy is FCFS and exponentially service req.
  $\rightarrow \mathbb{E}(W_i(t)) = \frac{1}{\mu_i}\mathbb{E}(N_i(t))$, and hence

$$\frac{1}{\mu_0}\mathbb{E}(N_0^{\pi}(t)) + \frac{1}{\mu_i}\mathbb{E}(N_i^{\pi}(t)) \geq \frac{1}{\mu_0}\mathbb{E}(N_0^{\tilde{\pi}}(t)) + \frac{1}{\mu_i}\mathbb{E}(N_i^{\tilde{\pi}}(t)).$$

We have $\mathbb{E}(N_0^{\pi}(t)) \geq \mathbb{E}(N_0^{\tilde{\pi}}(t))$ and

$$\frac{1}{\mu_0}\mathbb{E}(N_0^{\pi}(t)) + \frac{1}{\mu_i}\mathbb{E}(N_i^{\pi}(t)) \geq \frac{1}{\mu_0}\mathbb{E}(N_0^{\tilde{\pi}}(t)) + \frac{1}{\mu_i}\mathbb{E}(N_i^{\tilde{\pi}}(t)).$$

We have $\mathbb{E}(N_0^{\pi}(t)) \geq \mathbb{E}(N_0^{\tilde{\pi}}(t))$ and

$\frac{1}{\mu_0}\mathbb{E}(N_0^{\pi}(t)) + \frac{1}{\mu_i}\mathbb{E}(N_i^{\pi}(t)) \geq \frac{1}{\mu_0}\mathbb{E}(N_0^{\tilde{\pi}}(t)) + \frac{1}{\mu_i}\mathbb{E}(N_i^{\tilde{\pi}}(t))$.

Since $\sum_{i=1}^{L} c_i\mu_i \leq c_0\mu_0$, we obtain

$$\sum_{i=0}^{L} c_i\mathbb{E}(N_i^{\pi}(t))$$

$$= \frac{c_0\mu_0 - \sum_{i=1}^{L} c_i\mu_i}{\mu_0}\mathbb{E}(N_0^{\pi}(t)) + \sum_{i=1}^{L} c_i\mu_i \left(\frac{1}{\mu_0}\mathbb{E}(N_0^{\pi}(t)) + \frac{1}{\mu_i}\mathbb{E}(N_i^{\pi}(t))\right)$$

$$\geq \frac{c_0\mu_0 - \sum_{i=1}^{L} c_i\mu_i}{\mu_0}\mathbb{E}(N_0^{\tilde{\pi}}(t)) + \sum_{i=1}^{L} c_i\mu_i \left(\frac{1}{\mu_0}\mathbb{E}(N_0^{\tilde{\pi}}(t)) + \frac{1}{\mu_i}\mathbb{E}(N_i^{\tilde{\pi}}(t))\right)$$

$$= \sum_{i=0}^{L} c_i\mathbb{E}(N_i^{\tilde{\pi}}(t)).$$