Individual quality-of-service guarantees and call admission

Marty Reiman
Bell Laboratories

and
Adam Shwartz
Technion and Bell Laboratories

BRIMS, 20 August 1997

The context:

- ATM traffic
- Many users
- Statistical multiplexing
- Quality of service: blocking, delay
- Precise definitions and optimality
- Calculating an admission policy

Outline:

- A traffic model
- Definitions of QoS
- Optimal policies and mathematical programs
- Admission with identical users
- Mixed criteria optimization

The model:

- ATM link is fluid pipe, capacity R

- J classes of connections,
  Poisson arrivals $\lambda_i,\ i = 1, \dots, J$
  $\mathbf{k}(t) = \{k_i(t),\ i = 1, \dots, J\}$ users

- fast on-off transitions: call of type i is "on" w.p. p_i. Hence binomial distribution of "on" sources: Probability of $\mathbf{n} = \{n_i(t),\ i = 1, \dots, J\}$ users transmitting

$$\psi(\mathbf{k}, \mathbf{n}) = \prod_{i=1}^{J} \binom{k_i}{n_i} p_i^{n_i} (1 - p_i)^{k_i - n_i}$$

- Type i call terminates when "off" w.p. q_i.

- Total cell (=fluid) rate $\lambda_a(\mathbf{n}) = \sum_{i=1}^{J} n_i \nu_i$

Probability of $\mathbf{n}$ users transmitting is $\psi(\mathbf{k}, \mathbf{n})$.
Total cell (=fluid) rate is $\lambda_a(\mathbf{n})$.

- Cell loss rate for type i in state $\mathbf{k}$

$$b_i(\mathbf{k}) = \sum_{n_1=0}^{k_1} \cdots \sum_{n_J=0}^{k_J} \psi(\mathbf{k}, \mathbf{n}) \left[\lambda_a(\mathbf{n}) - R\right]^{+} \frac{n_i \nu_i}{\lambda_a(\mathbf{n})}$$

- Type i has guaranteed cell loss f_i:

$$g_i^0 = b_i(\mathbf{k}) - f_i k_i p_i \nu_i \quad \text{or} \quad g_i^1 = \frac{b_i(\mathbf{k})}{k_i} - f_i p_i \nu_i$$

  is the instantaneous violation/margin of the Cell-Loss Guarantee for type i.

An admission policy is chosen to
minimize blocking probability
(steady state, weighted) subject to
Guarantees on Cell Losses (each type).

Definitions of Cell-Level QoS: recall
g_i measures instantaneous type-i Cell-Loss.

- Conservative: $g_i(\mathbf{k}) \leq 0$, all $\mathbf{k}$.

- Average: $\sum_{\mathbf{k}} g_i(\mathbf{k})\pi(\mathbf{k}) \leq 0$
  where $\pi(\mathbf{k})$ is the steady state distribution.

Drawbacks:
Conservative does not allow burstiness
Average is user-meaningless, optimistic.

Individual QoS:
Follow a new connection;
  enforce guarantee until end of connection.
This depends on transient behavior!
Restrict to stationary, threshold policies:
Admit iff inside "admission region."

Optimal Policies, Mathematical Programs.

MDP's: discrete time. At time t:
if at state x and action a is chosen, then
Probability of ending in state y is $p(y|x,a)$
A reward of $c(x,a)$ is collected.
SMDP: random transition time,
statistics depend on (x,a,y).

A decision rule $\sigma(a|y)$ specifies
probability of using action a when in state y.

Stationary policy is identified with decision rule.

Average reward under σ (think blocking)

$$V(\sigma) = \lim_{T\to\infty} \mathbb{E}_y^\sigma \frac{1}{T} \sum_{t=1}^{T} c(x_t, a_t)$$

is to be maximized over policies σ.

Average reward under σ

$$
\begin{aligned}
V(\sigma) &= \lim_{T\to\infty} \mathbb{E}_y^\sigma \frac{1}{T}\sum_{t=1}^{T} c(x_t, a_t) \\
&= \lim_{T\to\infty} \frac{1}{T}\sum_{x,a}\sum_{t=1}^{T} c(x,a)\mathbb{P}^\sigma(x_t = x, a_t = a) \\
&= \sum_{x,a}\left(\lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}\mathbb{P}^\sigma(x_t = x, a_t = a)\right) c(x,a) \\
&\stackrel{\text{def}}{=} \sum_{x,a} z^\sigma(x,a)c(x,a).
\end{aligned}
$$

$z^\sigma(x,a)$ is probability (steady state) to be in state x and use action a. So,

$$
\sigma(a|x) = \frac{z^\sigma(x,a)}{\sum_b z^\sigma(x,b)} \ .
$$

So, σ determines z^σ through limits, z (satisfying (4)–(6)) determines σ^z.

Theorem 1. *Stationary policies σ are in one-to-one correspondance with occupation measures $z = \{z(x,a)\}$ with*

$$z(x,a) \geq 0 \tag{4}$$

$$\sum_{x,a} z(x,a) = 1 \tag{5}$$

$$\sum_{x,a} z(x,a)[\delta_{xy} - p(y|x,a)] = 0 \qquad \textit{all } y \tag{6}$$

and they are related by

$$\sigma^z(a|x) = \frac{z^\sigma(x,a)}{\sum_b z^\sigma(x,b)} \ .$$

Moreover, σ maximizes $V(\sigma)$ if and only if z^σ solves the linear program

$$\max \quad \sum_{x,a} z(x,a)c(x,a)$$

subject to the linear constraints (4)–(6).

Impose a constraint (Individual cell-loss):
total cost from x until some random time τ

$$\tilde{V}_x(\sigma) = \mathbb{E}_x^\sigma \sum_{t=1}^{\tau} r(x_t, a_t)\,.$$

Fix x. Assume τ does not depend on σ.
Define a new process: at τ jump back to x.
Then by renewal arguments

$$\tilde{V}_x(\sigma) = \lim_{T\to\infty} \tilde{\mathbb{E}}_x^\sigma \frac{1}{T} \sum_{t=1}^{\infty} \tilde{r}(x_t, a_t)\ .$$

By Theorem, finding σ to maximize V subject to a constraint on $\tilde{V}_x$ is equivalent to finding variables $(z, \tilde{z})$ so that
each satisfies the linear constraints,
each is related to σ by formula, and
σ maximizes a linear cost:

Theorem 2. *The problem*

$$\max \quad V(\sigma) \quad \textit{subject to} \quad \tilde{V}_x(\sigma) \le f$$

is equivalent to the mathematical program

$$\max \quad \sum_{x,a} z^1(x,a)c(x,a)$$

subject to ($i = 1, 2$)

$$\sum_{x,a} z^2(x,a)\tilde{r}(x,a) \le f \tag{12}$$

$$z^i(x,a) \ge 0 \tag{13}$$

$$\sum_{x,a} z^i(x,a) = 1 \tag{14}$$

$$\sum_{x,a} z^i(x,a)[\delta_{xy} - p^i(y|x,a)] = 0 \tag{15}$$

for all y, and for all x, a,

$$\frac{z^1(x,a)}{\sum_b z^1(x,b)} = \frac{z^2(x,a)}{\sum_b z^2(x,b)} \ . \tag{16}$$

Nonlinear non-convex mathematical program.
Note different transition probabilities p^i!

Remarks on Mathematical Program.

Theoretically:
a general method to incorporate constraints.
can be used to solve Individual QoS problem.

Individual QoS: variable z^x for each state!
mathematical progam is HUGE.

Continuous-time version more complicated.

Solutions of mathematical program are not always deterministic policies, due to constraints.

May be usefull for structural results.
practically useless.

QoS with a single class.

Minimize (steady state) blocking, subject to Cell-loss Guarantee at level f
Find a Threshold (Stationary Deterministic) K

$$g^0 = b(k) - fkp\nu \quad \text{or} \quad g^1 = \frac{b(k)}{k} - fp\nu$$

$$K_C = \max\{k : g(k) \le 0\} \qquad \text{Conservative}$$

$$K_A = \max\left\{k : \sum_{m=0}^{k} g(m)\pi^k(m) \le 0\right\}$$

Average

Individual criterion: QoS guarantee f should hold for new call accepted at each state x

$$K_I = \max\left\{k : \sum_{m=0}^{k} g(m)\tilde{\pi}_x^k(m) \le 0 \quad \text{for all } x \le k\right\}$$
$$= \max\{k : V(k,x) \le 0 \quad \text{for all } x \le k\}$$

$\tilde{\pi}_x$ derived from process that terminates when a “marked” call that entered at x terminates.

Structure of Optimal Policies

Individual criterion: minimize blocking under QoS guarantee.

Lemma 3. *Any feasible policy blocks beyond some K_{max}. Consequently, $K_I < \infty$.*

Theorem 4. $K_C \le K_I \le K_A$.

Note that $K_A = \infty$ for small λ, but $K_I < \infty$.

Computing optimal policies.

Theorem 5. $V(k,k) \le 0$ *implies* $V(k,x) \le 0$ *for* $x < k$. *Hence*

$$\begin{aligned} K_I &= \max\left\{k : \sum_{m=0}^{k} g(m)\tilde{\pi}_x^k(m) \le 0 \quad \textit{for all } x \le k\right\} \\ &= \max\left\{k : \sum_{m=0}^{k} g(m)\tilde{\pi}_x^k(m) \le 0\right\} \\ &= \max\{k : V(k,k) \le 0\} \end{aligned}$$

Computing optimal policies.

- Compute $V(k,k)$ via second-order difference equation with boundary conditions.
- Reduce to equation with initial conditions
- Increase k from K_C until $V(k,k) > 0$.
- Computation time negligible.

Numerical examples: $R = 45$, $f = 10^{-9}$. λ chosen in cases 1,3 so that blocking probability is 0.01.

μ	ν	p	λ	K_C	K_I	K_A
0.1	6.0	0.025	1.125	14	16	20
0.1	6.0	0.025	1 .05	14	16	∞
1.0	1.5	0.1	94.0	100	105	111
1.0	1.5	0.1	85.0	100	107	∞