ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Observations on Equation-Based Rate Control*

Milan Vojnović

Joint Work with Jean-Yves Le Boudec

March 12, 2002

*IPAM Workshop on Large-Scale Communication Networks: Topology, Routing, Traffic, and Control, University of California, Los Angeles, USA, March 18, 2002. (http://www.ipam.ucla.edu)

The Control that We Study

Data send rate set as:

$$X_n = f(\hat{p}_n)$$

$$X(t) = X_n, \; T_n < t \le T_{n+1}$$

We call it basic control.

- $\{T_n\}$ a point process on $\mathbb{R}$; $S_n \stackrel{def}{=} T_{n+1} - T_n$
- $\bar{p}$, long-run loss event ratio

$$\bar{p} = \lim_{t \to \infty} \frac{\sum_{k>0} 1_{[0,t)}(T_k)}{\int_0^t X(s)ds}$$

- $\hat{p}_n$, estimator of $\bar{p}$ at T_n
- $f : [0,1] \to \mathbb{R}^+$, non-increasing
- f is typically TCP loss-throughput function; it is also function of some statistics of round-trip time; not considered here

The Control that We Study (cont'd)

Comprehensive control – the basic control, plus: if at t, the amount of data sent since the most recent loss event, $\theta(t)$, would increase the value of the estimator $\widehat{\theta}(t)$, then use it as a sample

Sample-Paths:

Basic Control

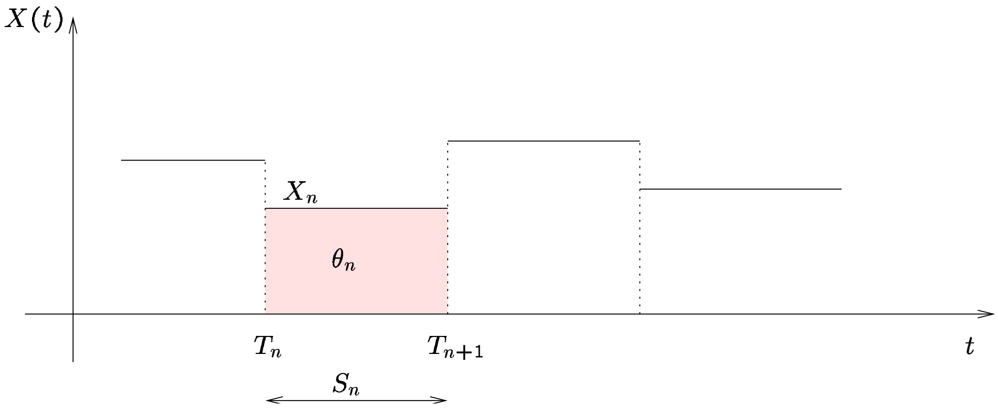

Comprehensive Control

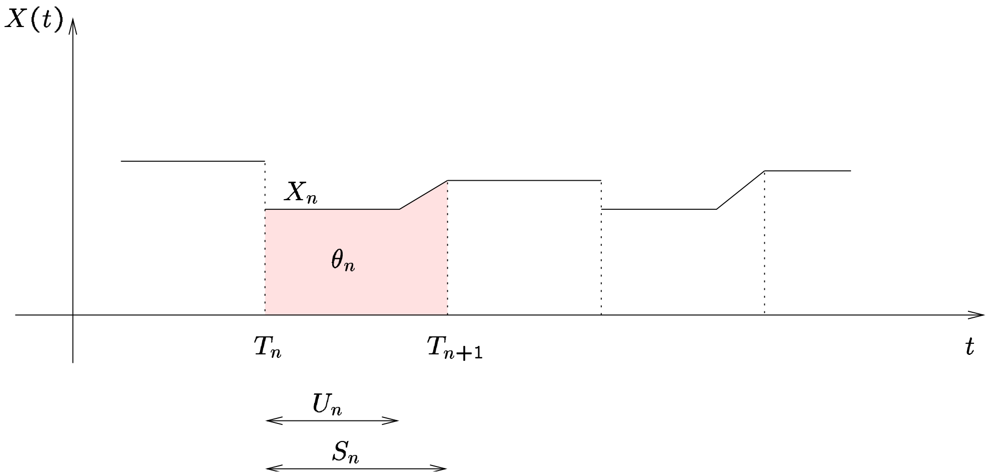

What is the Problem?

P: Is it true

$$\mathbb{E}[X(0)] \leq \mathbb{E}[X_{tcp}(0)] \text{ ?}$$

Here $\mathbb{E}[X_{tcp}(0)]$ is TCP throughput under the same operating conditions. If yes, we say the control is TCP-friendly.

We subdivide **P** into three subproblems:

P1:

$$\mathbb{E}[X(0)] \leq f(\bar{p}) \text{ ?}$$

If yes, we say the control is conservative.

P2: Does it hold

$$\bar{p} \geq \bar{p}_{tcp} \text{ ?}$$

P3: Does TCP satisfy its equation,

$$\mathbb{E}[X_{tcp}(0)] = f(\bar{p}_{tcp}) \text{ ?}$$

Why is the Problem of Interest?

- Several equation-based rate controls proposed with various definitions of $\{T_n\}$, $\widehat{p}_n$

  Ex. See:

  `http://www.psc.edu/networking/tcp_friendly.html`

  `http://www.icir.org/tfrc`

- Common objective:
  - smooth send rate dynamics, but still responsive to sustained congestion
  - TCP-friendly: the long-run throughput less than or equal to TCP throughput under the same operating conditions

Two Additional Assumptions

(A1) $\{T_n\}$ is point process of loss events

(A2) $\widehat{p}_n = \frac{1}{\widehat{\theta}_n}$, $\widehat{\theta}_n = \sum_{l=1}^{L} w_l \theta_{n-l}$

- $(w_l)_{l=1}^{L}$ positive-valued, $\sum_{l=1}^{L} w_l = 1$
- θ_n is amount of data sent in $[T_n, T_{n+1})$; we call it as in TFRC: loss-event interval

Ob. $1/\widehat{p}_n$ is unbiased estimator of $1/\bar{p}$

Some Functions f

SQRT:

$$f(p) = \frac{1}{c_1 R \sqrt{p}}$$

PFTK-standard:

$$f(p) = \frac{1}{c_1 R \sqrt{p} + Q \min[1, c_2 \sqrt{p}](p + 32p^3)}$$

PFTK-simplified:

$$f(p) = \frac{1}{c_1 R \sqrt{p} + Q c_2 (p^{3/2} + 32p^{7/2})}$$

- R: expected round-trip time
- Q: expected retransmit timeout
- c_1, c_2: positive-valued constants

We Will See

1. Throughput Representation

2. What Makes the Basic Control Conservative or Not?

   - For the basic control: two sets of conditions for either conservative or non-conservative control

   - Suggestions of the analytical results, validation by numerical and ns-2 experiments

3. We expect $\bar{p} \geq \bar{p}_{tcp}$.

4. It may be $\mathbb{E}[X(0)] > \mathbb{E}[X_{tcp}(0)]$, even though $\mathbb{E}[X(0)] \leq f(\bar{p})$, and $\bar{p} \geq \bar{p}_{tcp}$.

5. Conclusion

Throughput Representation

Palm inversion formula:

$$\mathbb{E}[X(0)] = \lambda \mathbb{E}_T[\int_0^{S_0} X(s)ds]$$

- $\lambda = 1/\mathbb{E}_T[S_0]$
- $\mathbb{E}_T$ is expectation w.r.t. Palm probability (given a point at 0; $T_0 = 0$)

We suppose stability, i.e. $\{X_n\}$ and $\{S_n\}$ are stationary ergodic

Throughput Representation (cont'd)

Basic control:

$$\mathbb{E}[X(t)] = \frac{\mathbb{E}[\theta_n]}{\mathbb{E}[\frac{\theta_n}{f(1/\widehat{\theta}_n)}]}$$

Comprehensive control (PFTK-simplified):

$$\mathbb{E}[X(t)] \leq \frac{\mathbb{E}[\theta_n]}{\mathbb{E}[\frac{\theta_n}{f(1/\widehat{\theta}_n)}] - \mathbb{E}[V_n 1_{\widehat{\theta}_{n+1} > \widehat{\theta}_n}]}$$

$$V_n = \frac{1}{w_1}\Big[-2c_1 R(\widehat{\theta}_{n+1}^{\frac{1}{2}} - \widehat{\theta}_n^{\frac{1}{2}}) + 2c_2 Q(\widehat{\theta}_{n+1}^{-\frac{1}{2}} - \widehat{\theta}_n^{-\frac{1}{2}}) - \\ + \frac{2}{5}c_3 Q(\widehat{\theta}_{n+1}^{-\frac{5}{2}} - \widehat{\theta}_n^{-\frac{5}{2}}) + (\widehat{\theta}_{n+1} - \widehat{\theta}_n)\frac{1}{f(1/\widehat{\theta}_n)}\Big]$$

Ob. From the joint law of $\theta_n, \ldots, \theta_{n-L}$ one may compute the throughput

First Set of Sufficient Conditions for Conservativeness

Assume

(F1) $\frac{1}{f(1/x)}$ is convex with x

(C1) $\mathrm{Cov}[\theta_n, \widehat{\theta}_n] \leq 0$

Then, the basic control is conservative.

Moreover,

$$\mathbb{E}[X(t)] \leq f(\bar{p}) \frac{1}{1 + \frac{f'(\bar{p})\bar{p}^3}{f(\bar{p})} \mathrm{Cov}[\theta_n, \widehat{\theta}_n]}$$

Ob. $\mathrm{Cov}[\theta_n, \widehat{\theta}_n] = \sum_{l=1}^{L} w_l \mathrm{Cov}[\theta_n, \theta_{n-l}]$

Ex. (C1) is indeed true for $\{\theta_n\}$ renewal process (i.i.d.)

For SQRT and PFTK-simplified, $\frac{1}{f(1/x)}$ is Convex

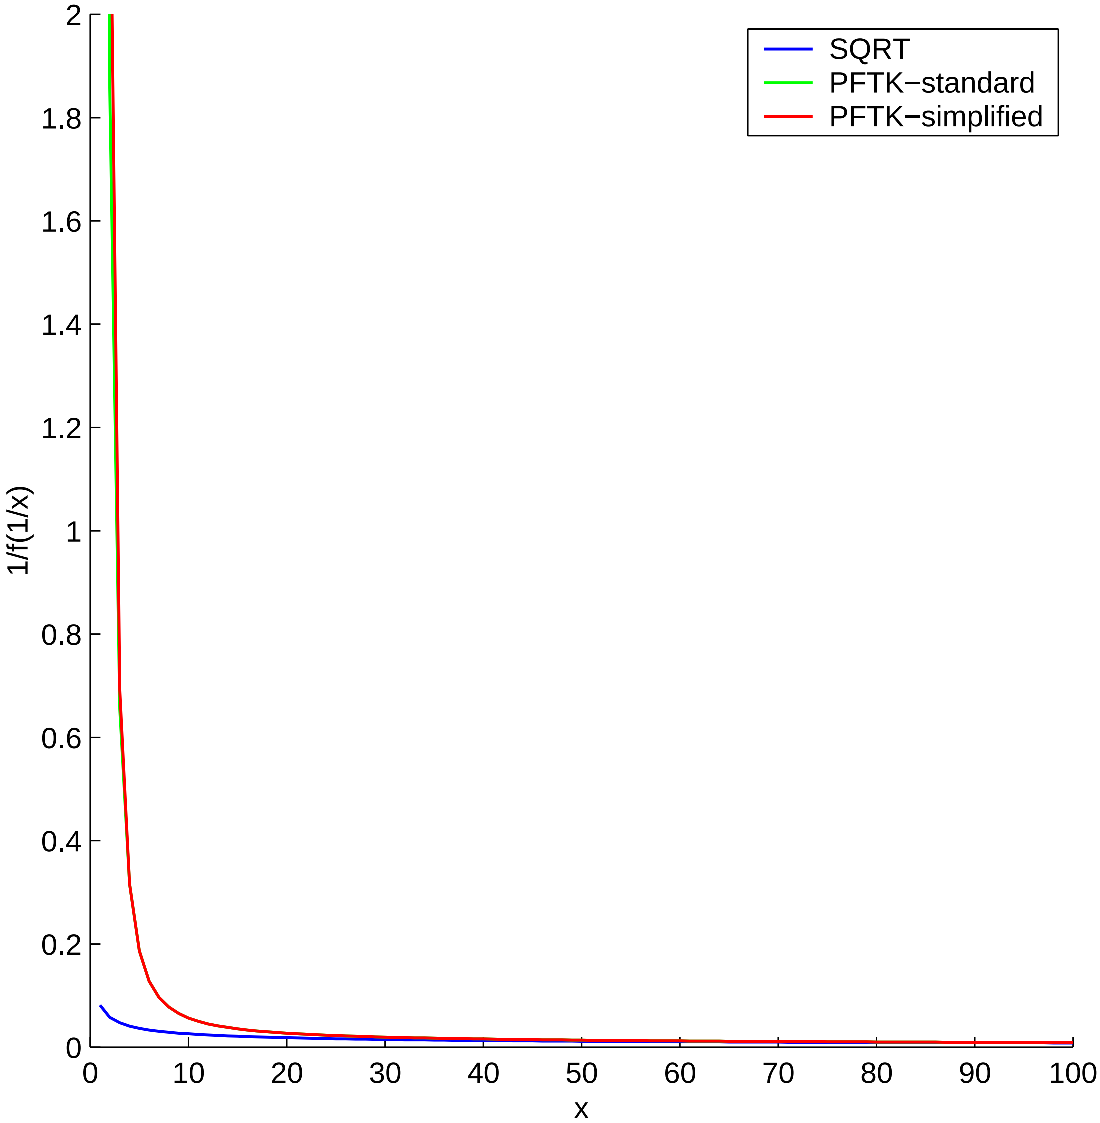

For PFTK-standard, $\frac{1}{f(1/x)}$ is Almost Convex

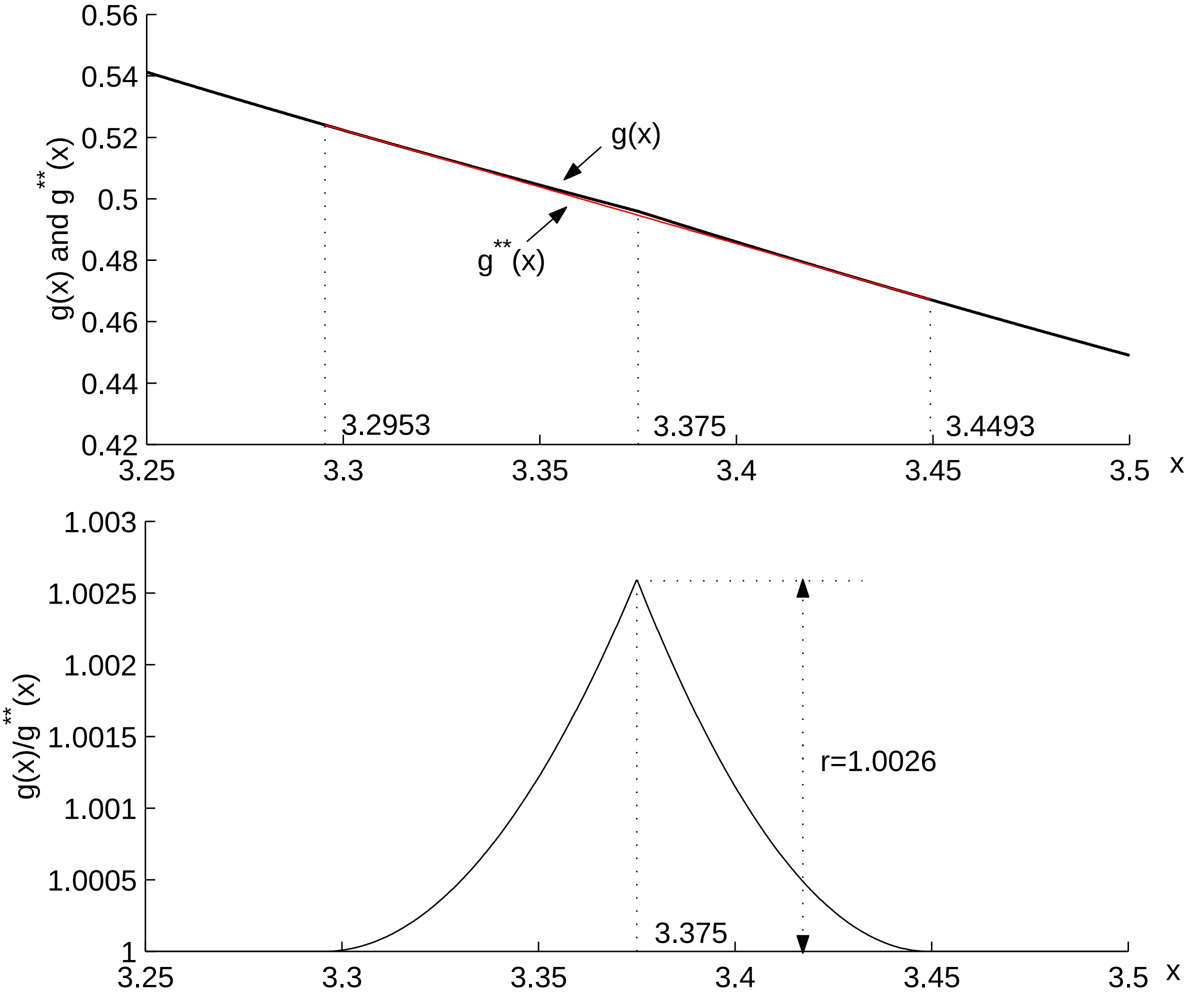

- Note: $g(x) \overset{def}{=} \frac{1}{f(1/x)}$, g^* is convex conjugate of g

First Suggestion

Assume θ_n and $\widehat{\theta}_n$ are negatively or lightly correlated. Consider the function f in the region where $\widehat{\theta}_n$ takes values.

1. The more convex $\frac{1}{f(1/x)}$, the more conservative the control is.

2. The more variable $\widehat{\theta}_n$, the more conservative the control is.

First Suggestion: Numerical Example

$\{\theta_n\}$ i.i.d. with generalized exponential density

$f_{\theta_n}(x) = \lambda\,\mathrm{exp}(-\lambda(x - x_0))$, $x \geq x_0$, $\lambda, x_0 \geq 0$

- $\mathbb{E}[\theta_n] = x_0 + \lambda$
- Coeff. of variation: $\frac{\lambda}{\sqrt{\lambda + x_0}}$
- Skewness: 2
- Kurtosis: 6

First Suggestion: Numerical Example

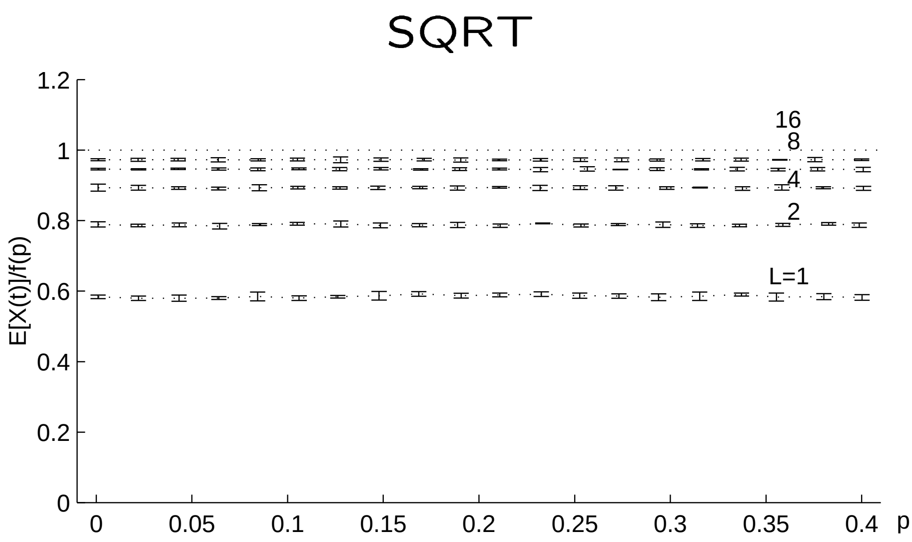

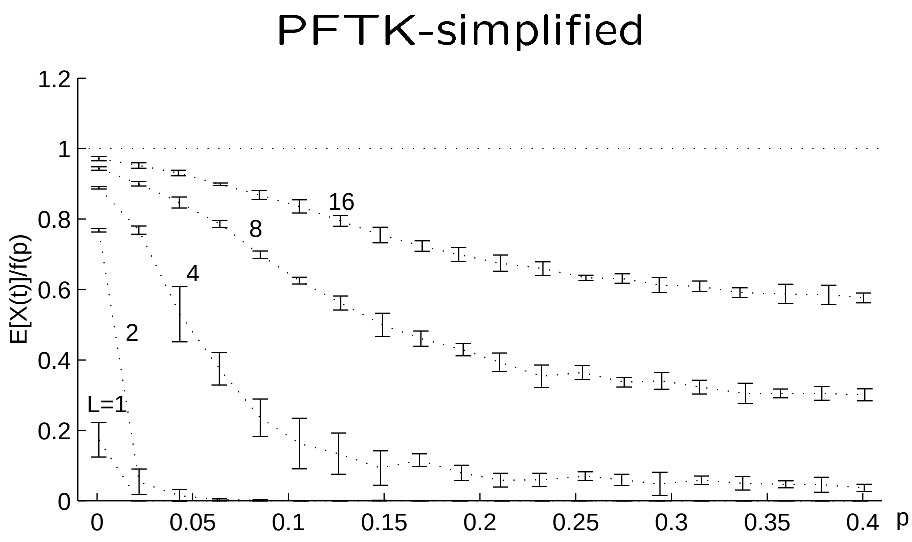

Ob. PFTK-simplified: the larger $\bar{p}$, the more convex $\frac{1}{f(1/x)}$, the more conservative the control.

First Suggestion: Numerical Example (PFTK-simplified, fixed $\bar{p}$)

$\bar{p} = 0.01$

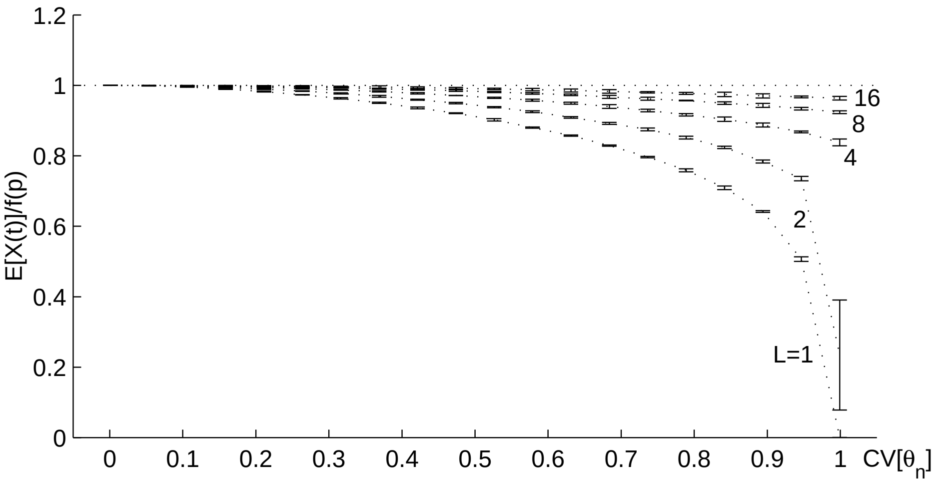

$\bar{p} = 0.1$

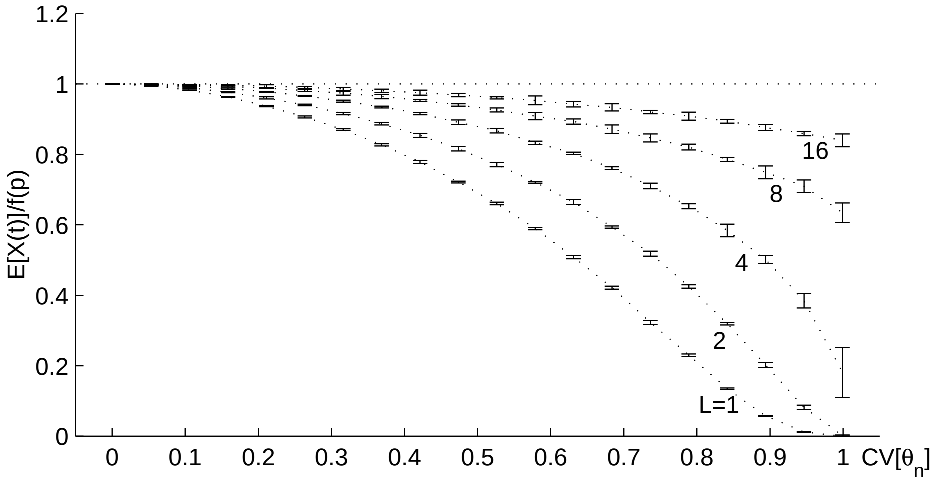

Ob. The more variable $\widehat{\theta}_n$, the more conservative the control.

First Suggestion: Numerical Example (Comprehensive Control)

SQRT

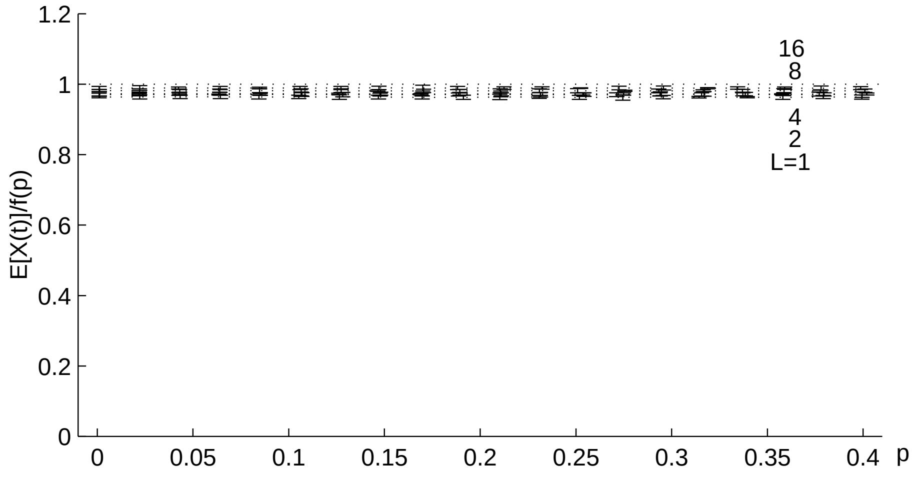

PFTK-simplified

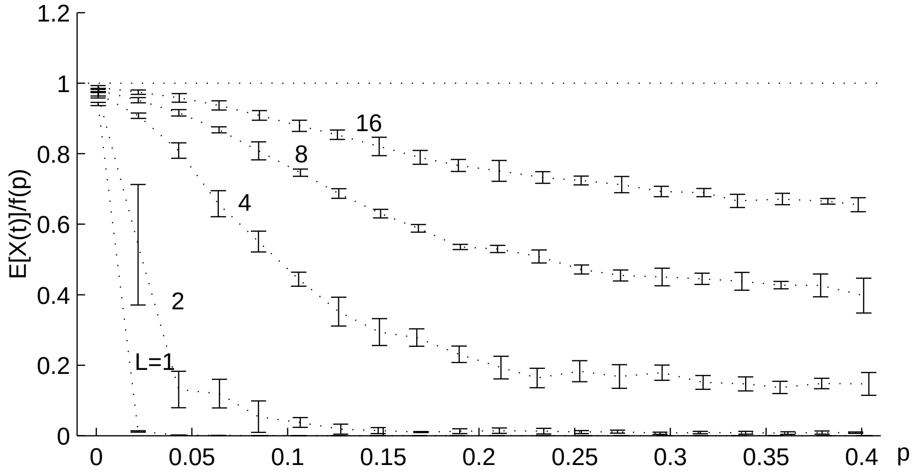

Ob. Qualitatively the same as for the basic control, but somewhat less pronounced.

First Suggestion: ns-2 Experiment (TFRC, PFTK-standard)

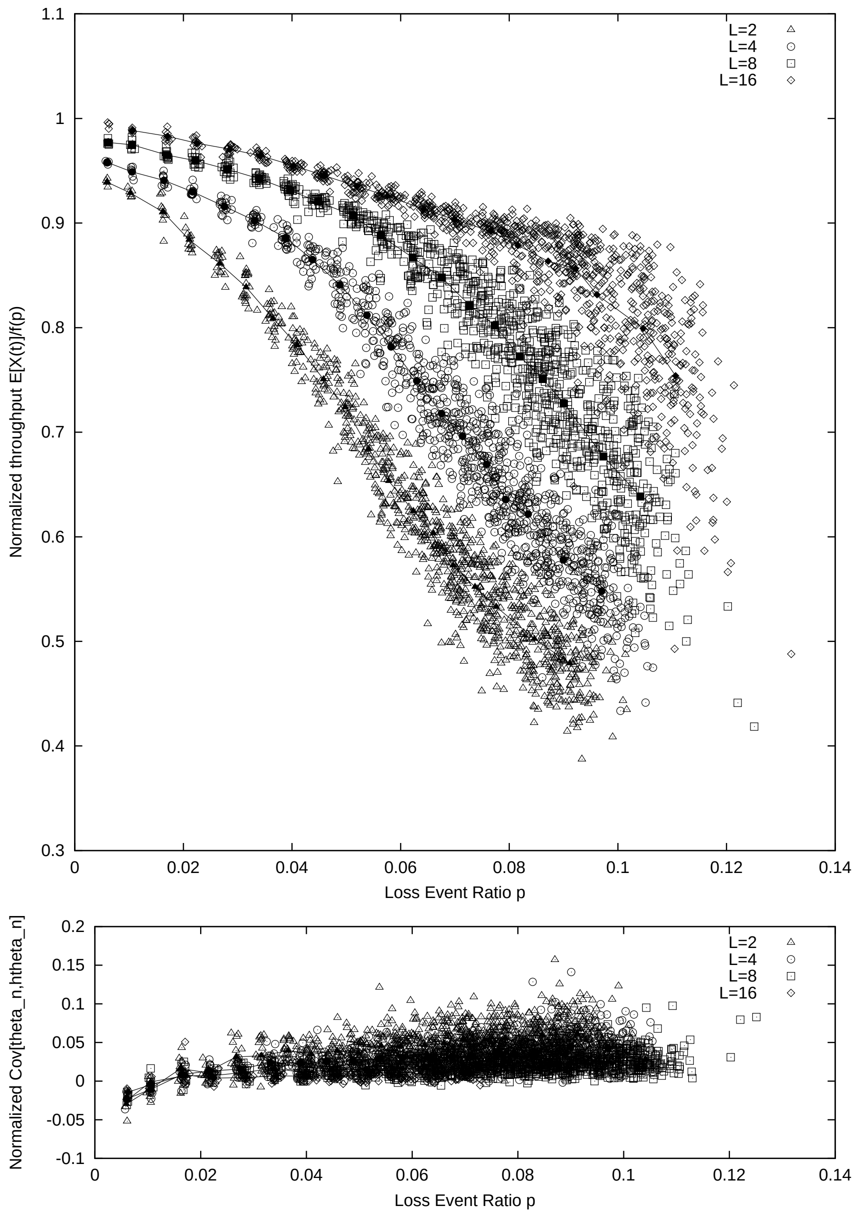

Setting: single RED link shared by TFRC and TCP-Sack1 connections (link capacity 15 Mb/s, round-trip time about 50 ms)

First Suggestion: ns-2 Experiment (TFRC, SQRT)

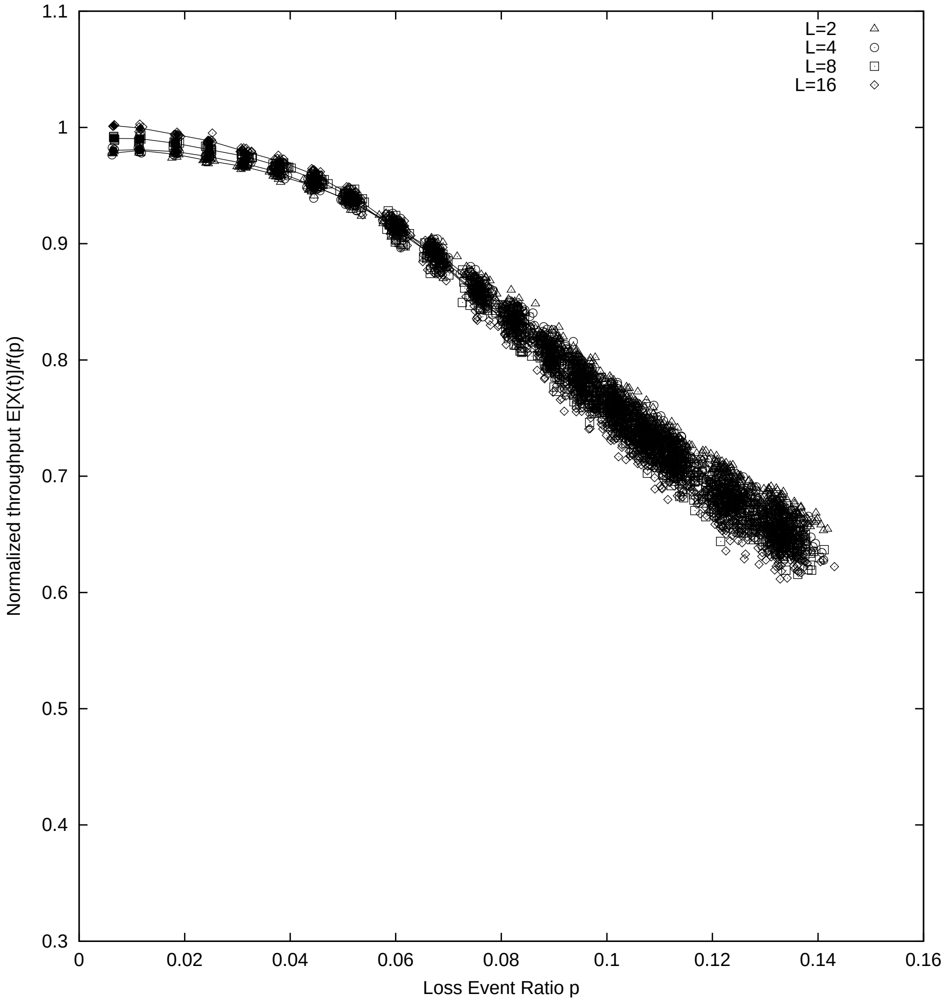

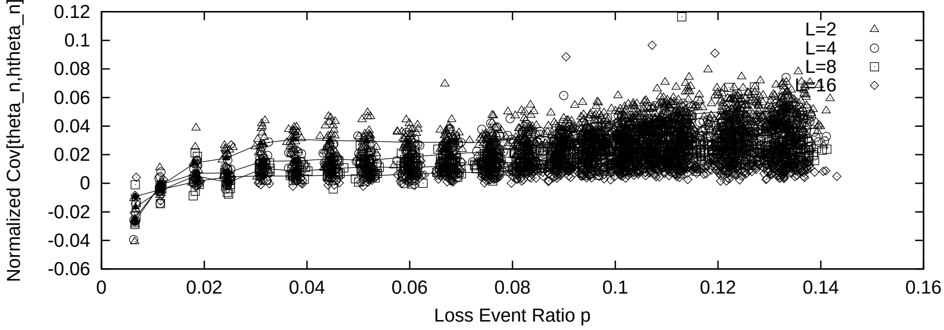

Second Set of Sufficient Conditions for Conservativeness

Assume

(F2) $f(1/x)$ is concave with x

(C2) Cov$[X_n, S_n] \leq 0$

Then, the basic control is conservative.

Conversely, if

(F2') $f(1/x)$ is convex with x

(C2') Cov$[X_n, S_n] \geq 0$

(V) $\widehat{\theta}_n$ has non-zero variance

Then, the basic control is non-conservative.

When $f(1/x)$ is concave or convex?

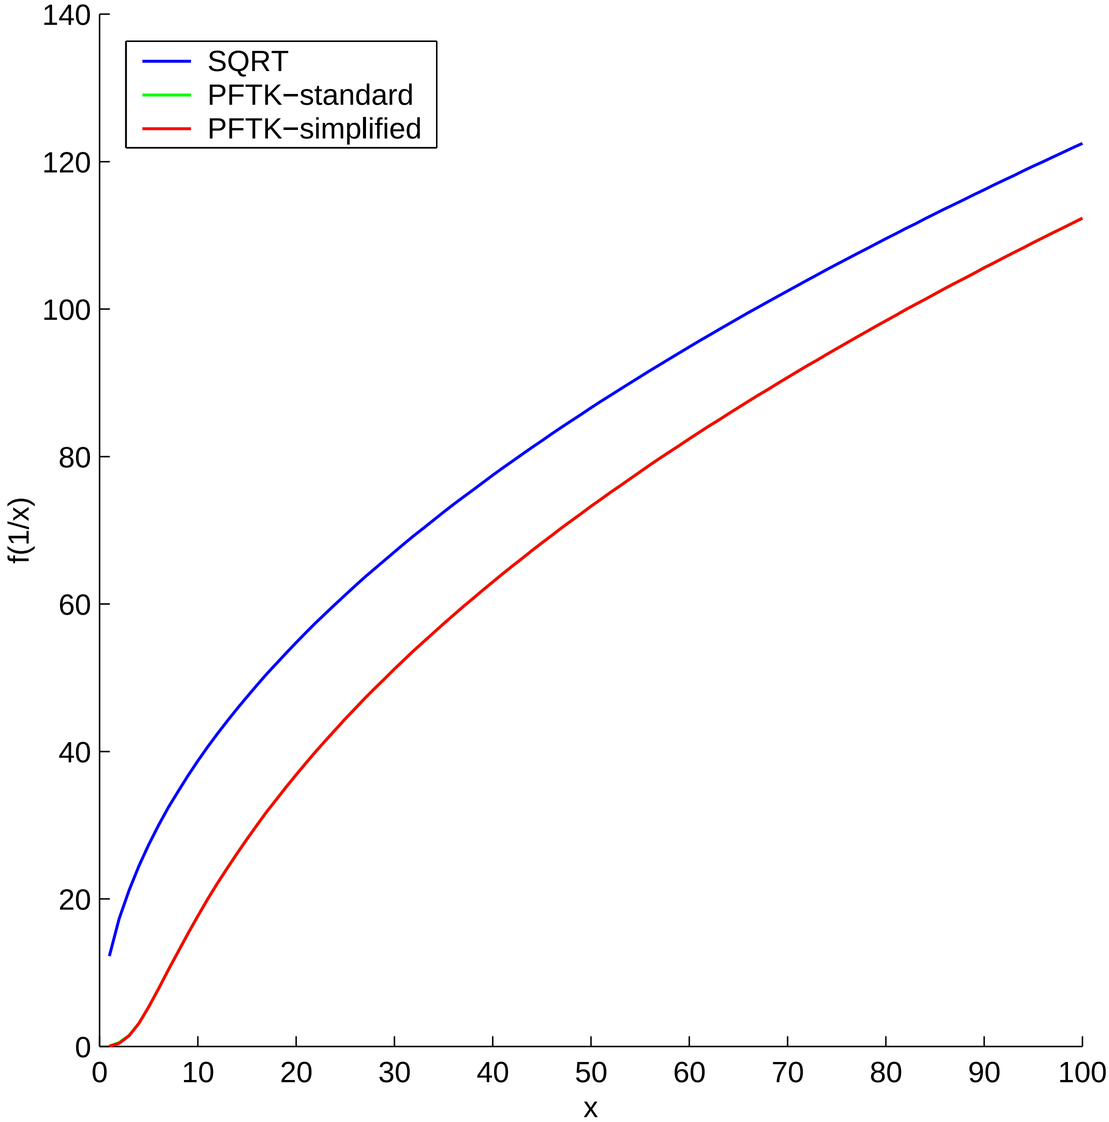

- SQRT: $f(1/x)$ concave
- PFTK: $f(1/x)$ concave for light loss, but convex for heavy loss

Importance of Feller's Paradox Type of Effects

Feller's Paradox: for a point process, average interval between two points seen by a random observer is larger than as seen at the interval boundaries.

Recall (C2): $\text{Cov}[X_0, S_0] \leq 0$

Thus, by Palm inversion formula:

$$\mathbb{E}[X(0)] = \mathbb{E}_T[X(0)] + \frac{\text{Cov}[X_0, S_0]}{E_T[S_0]}$$

$$\text{it follows (C2)} \Rightarrow \mathbb{E}[X(0)] \leq \mathbb{E}_T[X(0)]$$

Interpretation:

1. a random observer would more likely pick larger interval S_n
2. (C2) implies, on average, she would see smaller rate than as seen at $\{T_n\}$

When (C2) is True?

Ob. $\mathbb{E}[S_n|X_n = x]$ non-increasing with $x \Rightarrow$ Cov$[X_n, S_n] \leq 0$ (C2)

Ex. If θ_n is independent of data send rate X_n, then

$$\mathbb{E}[S_n|X_n = x] = \frac{1}{\bar{p}x}$$

Thus, (C2) holds.

Second Suggestion

- Assume S_n and X_n are negatively or non correlated.

  If $f(1/x)$ is concave in the region where $\widehat{\theta}_n$ takes its values, the control tends to be conservative.

- Conversely, assume S_n and X_n are positively or non correlated.

  If $f(1/x)$ is convex in the region where $\widehat{\theta}_n$ takes its values, the control is *non* conservative.

In both cases, more variable $\widehat{\theta}_n$ is, stronger the effect.

Second Suggestion: ns-2 Experiment

$L = 4$

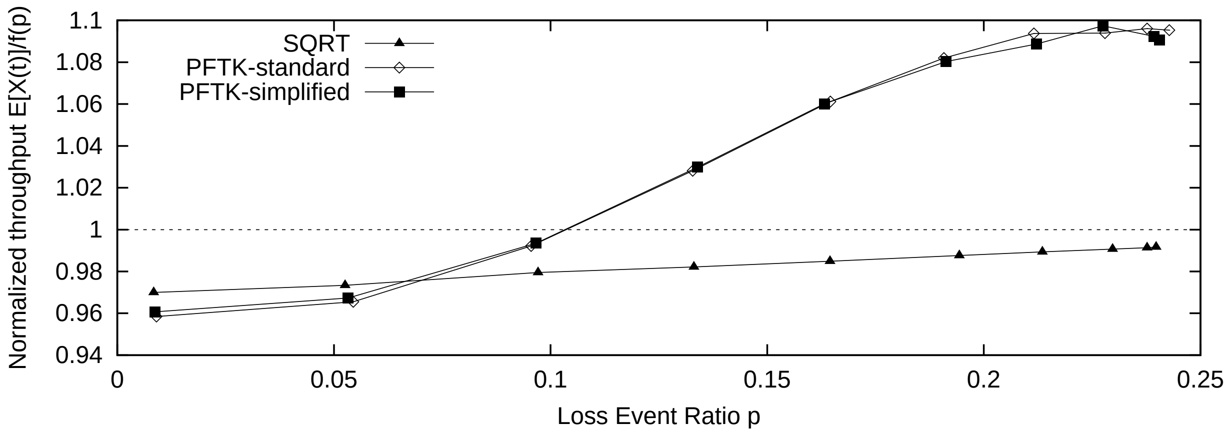

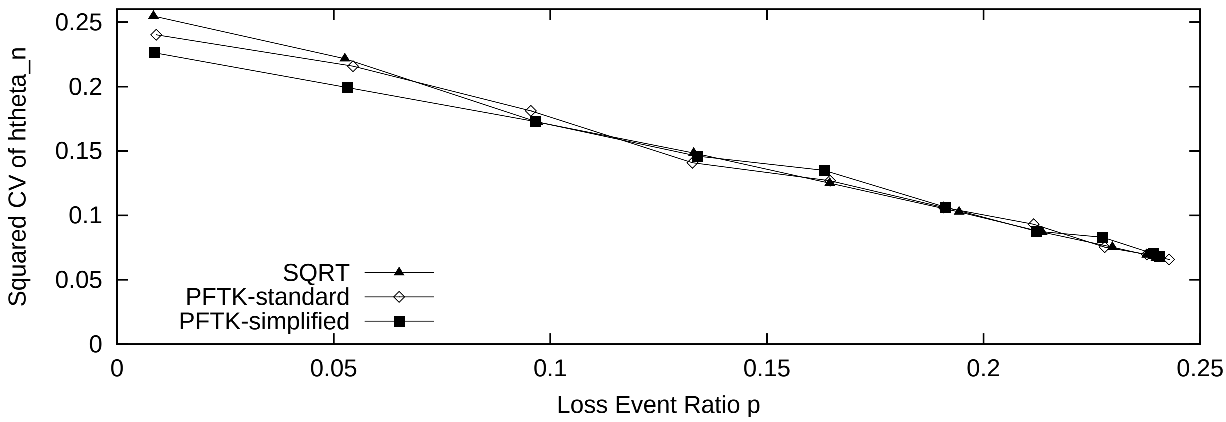

Setting: a rate controlled source with fixed packet send rate (each 20 ms) through a single loss link (fixed drop probability independent of the packet size)

Second Suggestion: ns-2 Experiment

$L = 8$

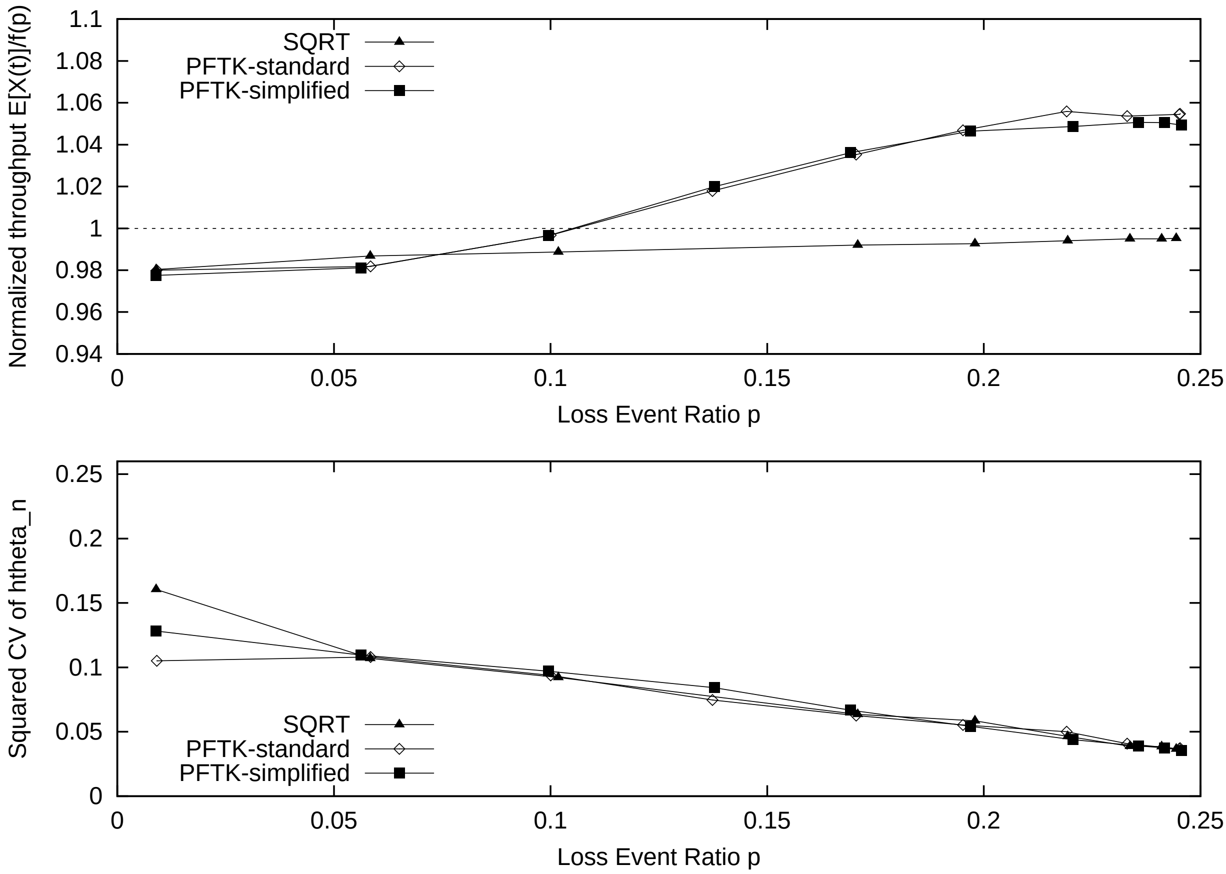

Ob. Qualitatively the same as for $L = 4$, but the effects less pronounced

Third Suggestion (Loss Event Ratios Seen by Sources)

The loss event ratios for TCP, our adaptive equation based rate controlled source (A), and a non-adaptive source (P) (Poisson) should be in the relation

$$(*) \qquad \bar{p}_{tcp} \leq \bar{p}_A \leq \bar{p}_P$$

The more responsive source A is, the closer $\bar{p}_A$ should be to $\bar{p}_{tcp}$.

Ob. If $\mathbb{E}[X(0)] \leq f(\bar{p}_A)$, and $(*)$, then

$$\mathbb{E}[X(0)] \leq f(\bar{p}_{tcp})$$

(Conservativeness implies TCP-friendliness, for a TCP source that attains throughput $\geq f(\bar{p}_{tcp})$, with equality if f is accurate loss-throughput function of the given TCP)

Third Suggestion: ns-2 Experiment

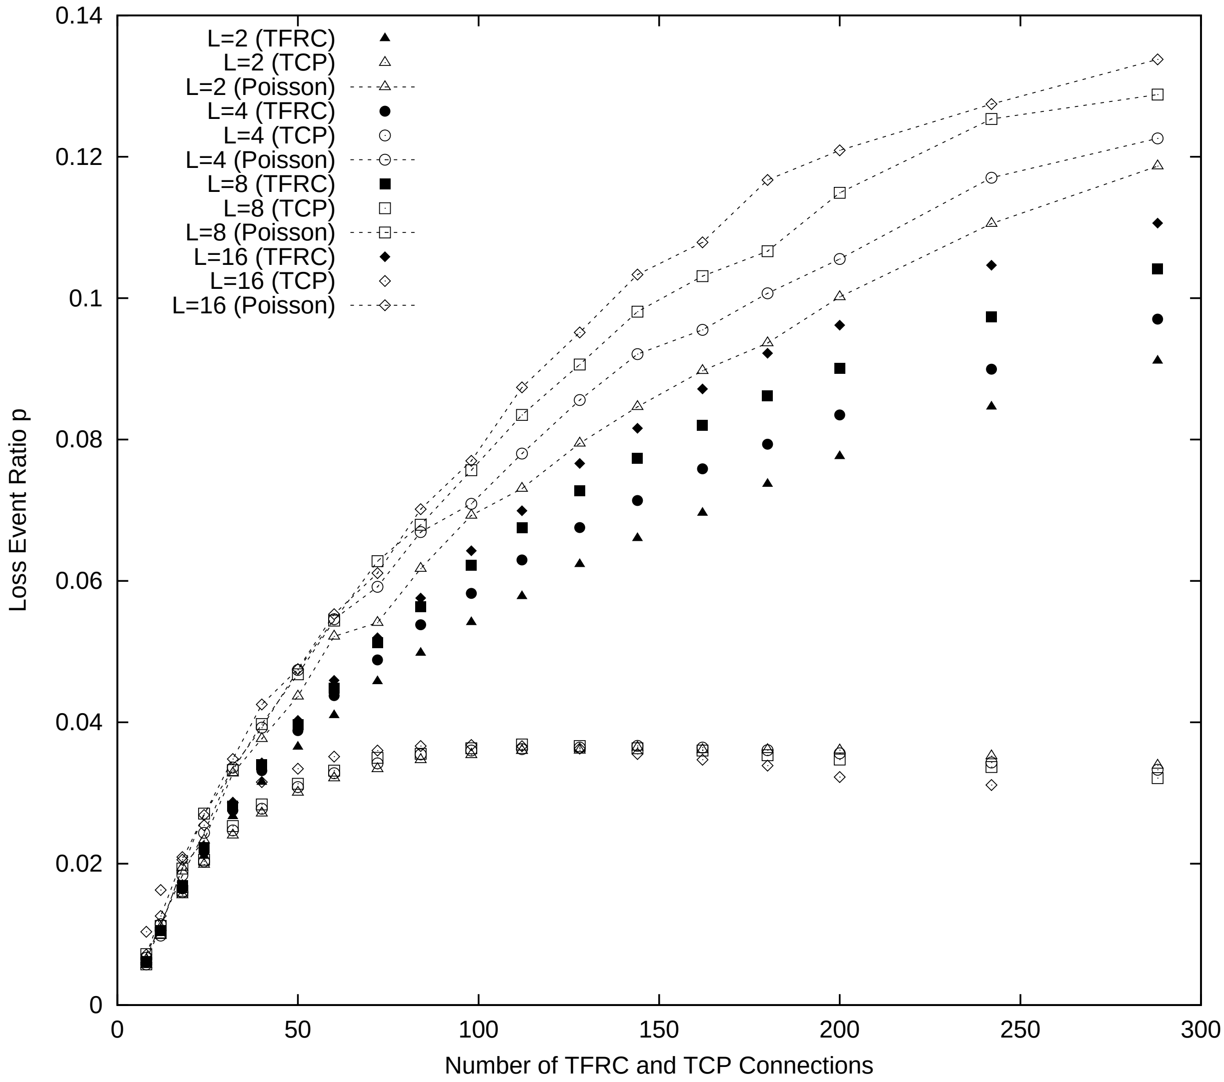

f may be Inaccurate Loss-throughput Function

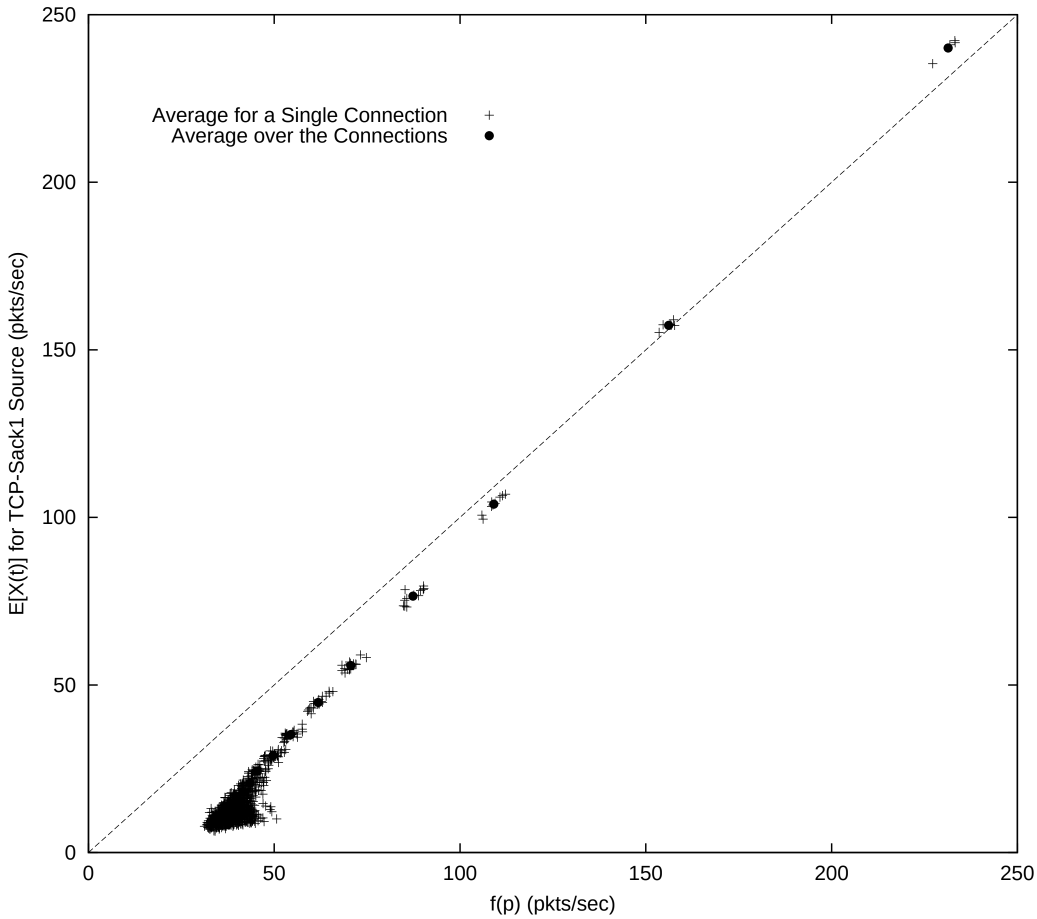

Ob. TCP-Sack1 does not verify PFTK loss-throughput function f

f may be Inaccurate Loss-throughput Function

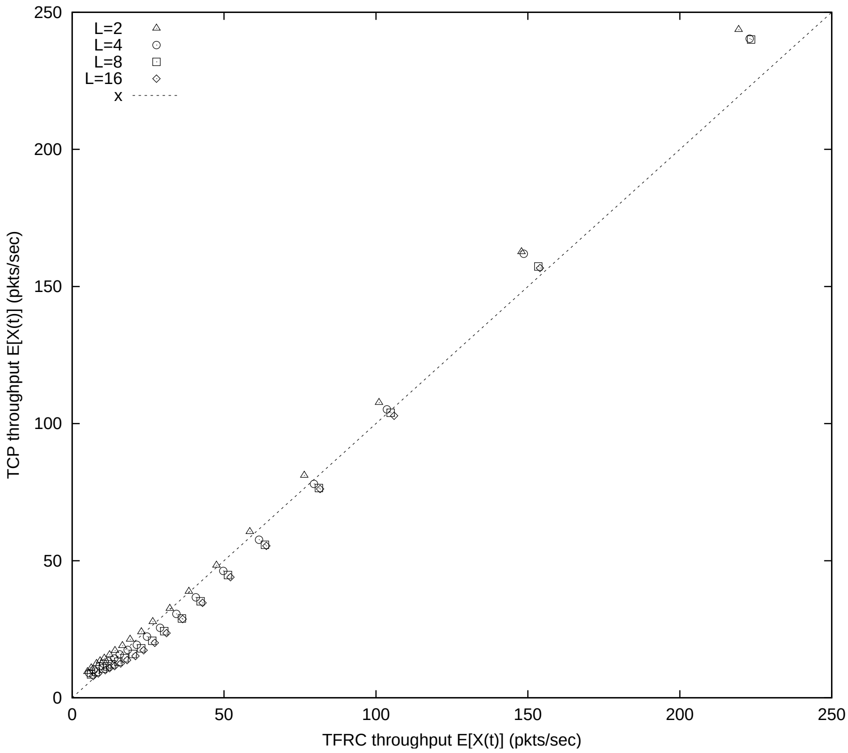

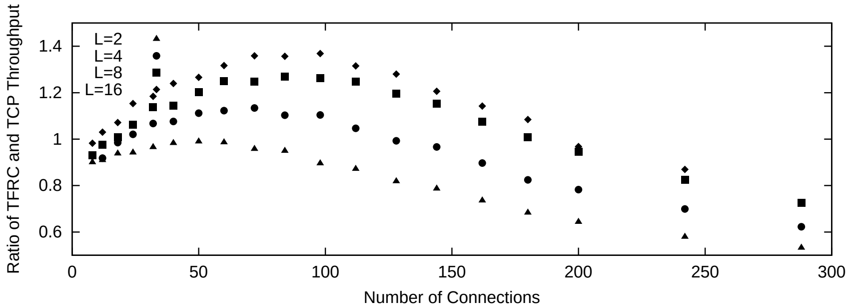

Ob. Control may be not friendly to the given TCP, even though being conservative and friendly to the function f

Ob. This is NOT problem of the control, but merely due to inaccuracy of f

Conclusion

Two causes of $\mathbb{E}[X(t)] \neq f(\bar{p})$:

- time versus event averages
- convexity properties of $f(1/x)$, $1/f(1/x)$

Important to separate:

- conservativeness
- this control loss event ratio versus TCP loss event ratio
- obedience of TCP to given function f

Pointers

1. V. and Le Boudec, "On the Long-Run Behavior of Equation-Based Rate Control" DSC Technical Report 02/006, February 2002.

2. V. and Le Boudec, "Some Observations on Equation-Based Rate Control," in Proc. of ITC-17, Salvador da Bahia, Brazil, September 24-28, 2001.

Available at: `http://lcawww.epfl.ch`

Appendix

Comparison of (Non-)Conservativeness Conditions

(F1)	$\frac{1}{f(1/x)}$ convex	**(F2)**	$f(1/x)$ concave
(F1')	$\frac{1}{f(1/x)}$ concave	**(F2')**	$f(1/x)$ convex

(F2) $\Rightarrow$ **(F1)**

Comparison of (Non-)Conservativeness Conditions (cont'd)

(C1)	Cov$[\theta_n, \widehat{\theta}_n] \leq 0$	**(C2)**	Cov$[X_n, S_n] \leq 0$
(C1')	Cov$[\theta_n, \widehat{\theta}_n] > 0$	**(C2')**	Cov$[X_n, S_n] > 0$

(C2) $\Leftrightarrow$

$$\text{Cov}[\theta_n, \frac{1}{f(1/\widehat{\theta}_n)}] \geq \mathbb{E}[\theta_n] \left(\frac{1}{\mathbb{E}[f(1/\widehat{\theta}_n)]} - \mathbb{E}[\frac{1}{f(1/\widehat{\theta}_n)}] \right)$$

(C2') $\Leftrightarrow$

$$\text{Cov}[\theta_n, \frac{1}{f(1/\widehat{\theta}_n)}] < \mathbb{E}[\theta_n] \left(\frac{1}{\mathbb{E}[f(1/\widehat{\theta}_n)]} - \mathbb{E}[\frac{1}{f(1/\widehat{\theta}_n)}] \right)$$

Ob. The RHS is negative.

Assume $g(x) = \frac{1}{f(1/x)}$ non-increasing convex.

(C2) $\Leftarrow$ Cov$[\theta_n, \widehat{\theta}_n] \leq \frac{\mathbb{E}[\theta_n]}{g'(\mathbb{E}[\theta_n])} \left(\frac{1}{\mathbb{E}[f(1/\widehat{\theta}_n)]} - \frac{1}{f(1/\mathbb{E}[\theta_n])} \right)$

(C2') $\Rightarrow$ Cov$[\theta_n, \widehat{\theta}_n] > \frac{\mathbb{E}[\theta_n]}{g'(\mathbb{E}[\theta_n])} \left(\frac{1}{\mathbb{E}[f(1/\widehat{\theta}_n)]} - \frac{1}{f(1/\mathbb{E}[\theta_n])} \right)$

Ob. For $f(1/x)$ convex with x, the RHS is positive. In this case, if **(C2')** holds, then necessarily **(C1')** holds.

Example of Non-Conservative Control

Assume:

- $\{\theta_n, Z_n\}$ is semi-Markov process
- $[p_{ij}]$ transition matrix of DTMC $\{Z_n\}$
- $\mathsf{P}(Z_{n+1} = j, \theta_n = m | Z_n = i) = p_{ij} g_i(m)$

Consider:

- $\{Z_n\}$ two-state DTMC with state space $\{g, b\}$
- periodic losses while in a given state; $\mathsf{P}(\theta_n = n_g | Z_n = g) = 1$ and $\mathsf{P}(\theta_n = n_b | Z_n = b) = 1$

Example of Non-Conservative Control (cont'd)

Slow MC limit, $p_{gb}, p_{bg} \to 0$, $\frac{p_{gb}}{p_{bg}} = u$,

$$\mathbb{E}[X(0)] \to \frac{p_{bg} n_g + p_{gb} n_b}{p_{bg} \frac{n_g}{f(1/n_g)} + p_{gb} \frac{n_b}{f(1/n_b)}}$$

Maximum attained for $\frac{p_{bg}}{p_{gb}} = \sqrt{\frac{n_b}{n_g}}$; maximum normalized throughput $(\mathbb{E}[X(0)]/f(\bar{p}))$,

$$x^* = \frac{1}{2} \sqrt{2 + \sqrt{\frac{n_g}{n_b}} + \frac{1}{\sqrt{\frac{n_g}{n_b}}}}$$

Example of Non-Conservative Control (cont'd)

$\mathbb{E}[X(0)]/f(\bar{p})$ vs. p_{gb} and p_{bg} (Basic Control)

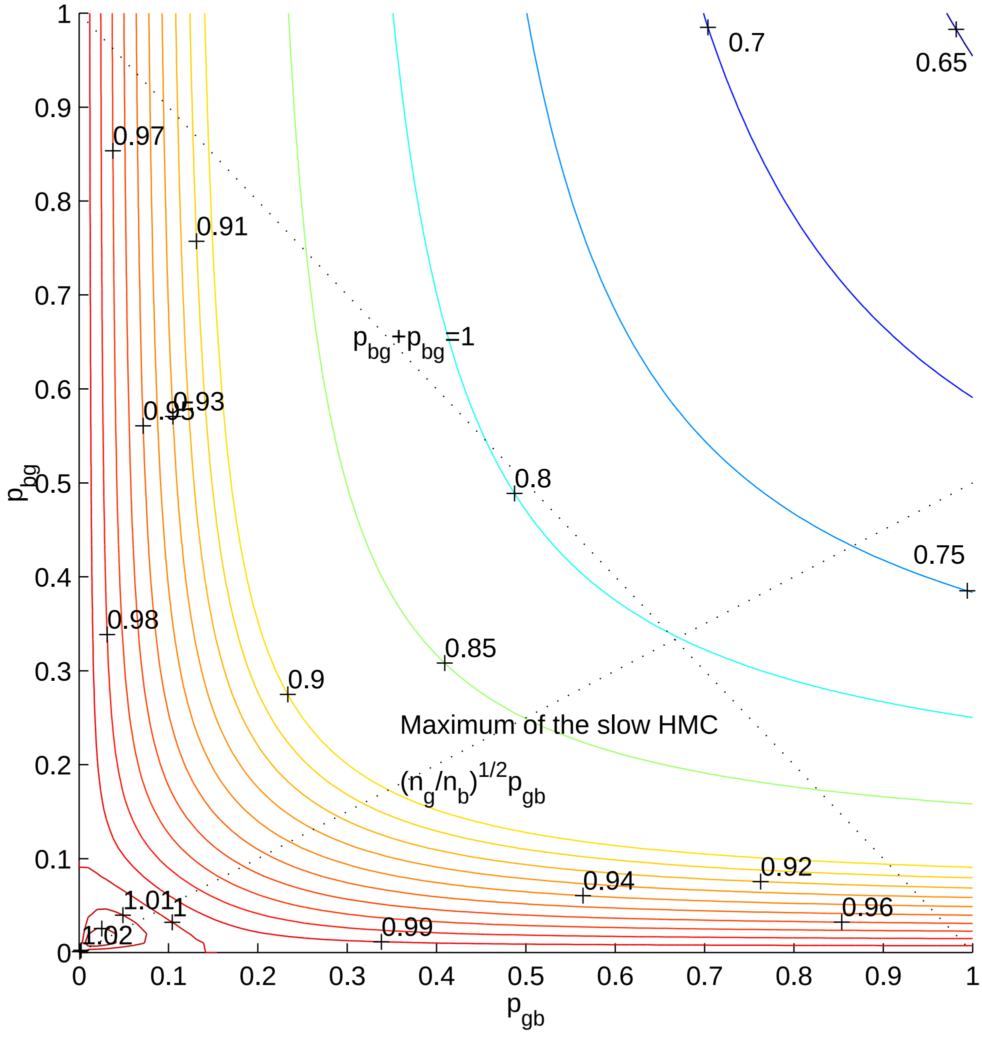

Example of Non-Conservative Control (cont'd)

$\mathbb{E}[X(0)]/f(\bar{p})$ vs. p_{gb} and p_{bg} (Comprehensive Control)

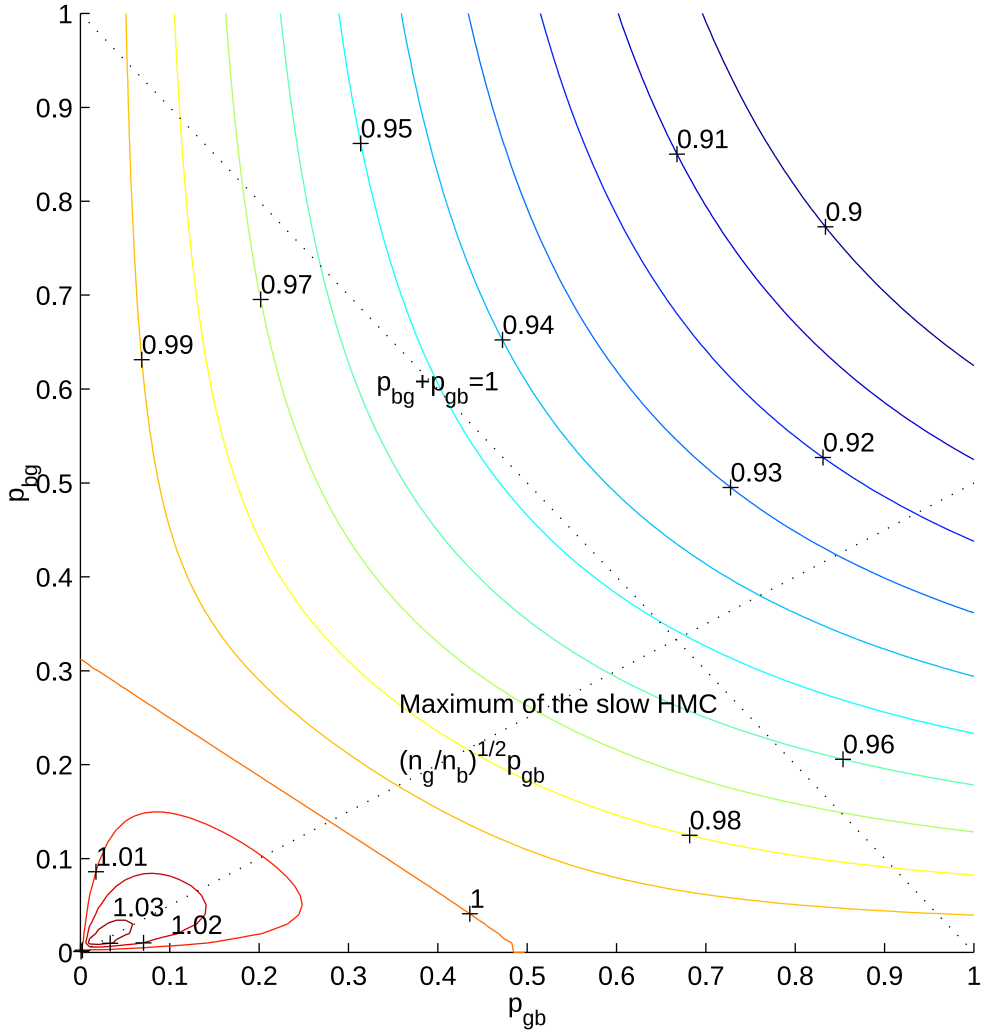

Example of Non-Conservative Control (cont'd)

Maximum $\mathbb{E}[X(0)]/f(\bar{p})$ for slow MC limit (x^*)