Non-asymptotic Delay Bounds for Networks with Heavy Tailed-Traffic

Jörg Liebeherr

Almut Burchard

University of Toronto

Florin Ciucu

Deutsche Telekom Labs/ TU Berlin

Infocom 2010

Overview

- Heavy-tailed and self-similar traffic envelope
  - Example: Pareto Traffic
- **Key result**: Sample path bound for htss traffic
- Delay Bounds:
  - Single node
  - End-to-end
- Scaling Properties

Heavy-Tailed Self-Similar Traffic

- A **heavy-tailed** process X satisfies

$$\boxed{Pr(X(t) > x) \sim Kx^{-\alpha}}$$

$$1 < \alpha < 2$$

- A **self-similar** process satisfies

$$\boxed{X(t) \sim_{dist} a^{-H} X(at)}$$

$$a > 0$$
$H \in (0, 1)$ Parameter

End-to-End Delays

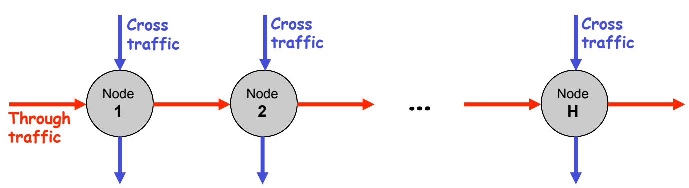

- Links have fixed capacity ***C***
- Through traffic and cross traffic are heavy-tailed and self-similar

Objective:
- Upper and lower bounds on end-to-end (E2E) delays
- Scaling properties of these delays

End-to-End Delays

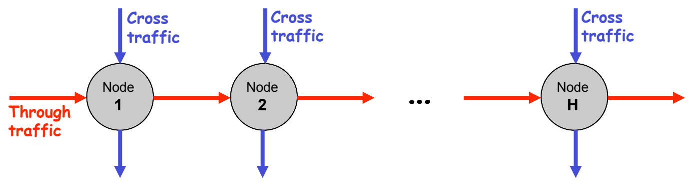

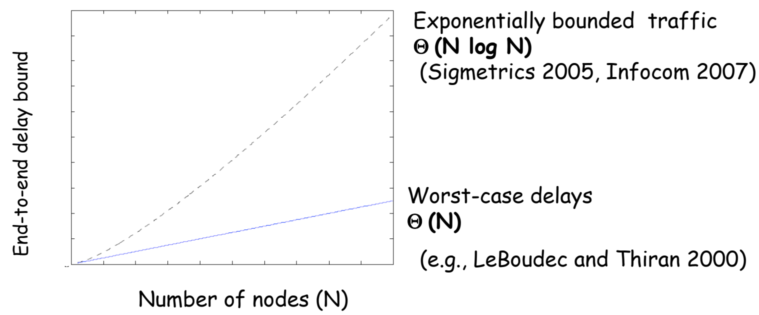

Stochastic Network Calculus

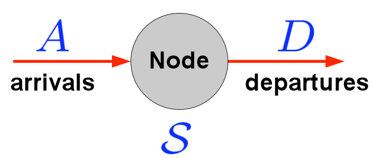

- Statistical arrival envelope:

$$Pr\{A(s,t) > \mathcal{G}(t-s) + \sigma\} \leq \varepsilon(\sigma)$$

- Service curve:

$$D(t) \geq A * \mathcal{S}(t)$$

- Apply min-plus algebra to obtain bounds on backlog, delay, and output burstiness

htts Traffic Envelope

- Heavy-tailed self-similar (**htss**) envelope:

$$Pr(A(s,t) > \underbrace{r(t-s) + \sigma(t-s)^H}_{\mathcal{G}(t-s;\sigma)}) \leq \underbrace{K\sigma^{-\alpha}}_{\varepsilon(\sigma)}$$

- **Main difficulty:** Backlog and delay bounds require sample path envelopes of the form

$$Pr(\sup_{s\leq t}\{A(s,t) - \overline{\mathcal{G}}(t-s;\sigma)\} > 0) \leq \varepsilon(\sigma)$$

- **Key contribution:**
  Derive sample path bound for htss traffic

Example: Pareto Traffic

- Size of i-th arrival: $Pr(X_i > x) = \left(\frac{x}{b}\right)^{-\alpha}$ $\quad x \geq b$, $1 < \alpha < 2$
- Arrivals are evenly spaced with gap λ: $A(t) = \sum_{i=1}^{N(t)} X_i$
- With Generalized Central Limit Theorem … $A(t) \approx \lambda t E[X] + c_\alpha(\lambda t)^{1/\alpha} S_\alpha$

  … and tail bound $Pr(S_\alpha > \sigma) \sim (c_\alpha \sigma)^{-\alpha}$

  α-stable distribution

- … we get htss envelope

$$\begin{aligned} \mathcal{G}(t;\sigma) &= \lambda E[X] t + \sigma t^{1/\alpha} \\ \varepsilon(\sigma) &= \lambda \sigma^{-\alpha} \end{aligned}$$

Example: Envelopes for Pareto Traffic

Parameters:

$$\alpha = 1.6$$
$$b = 150\ Byte$$
$$\lambda = 75\ Mbps$$

Comparison of envelopes:

- htss GCLT envelope
- Average rate
- Trace-based
  - deterministic envelope
  - htts trace envelope

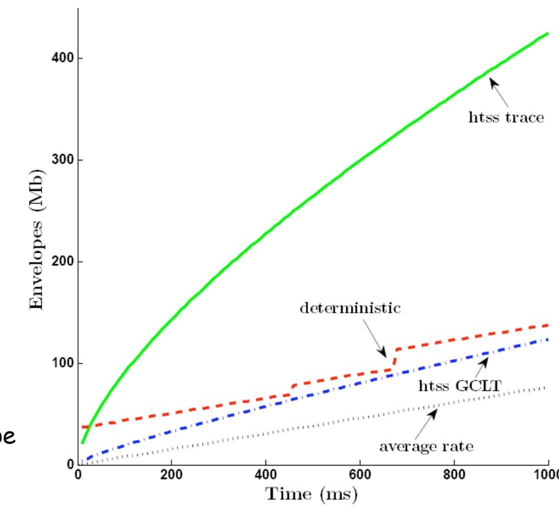

Backlog computation without sample path bound

- Backlog: $B(t) = \sup_{s \le t}\{A(s,t) - C(t-s)\}$

$$Pr(B(t) > \sigma) = Pr(\sup_{s \le t}\{A(s,t) - C(t-s)\} > \sigma)$$

- Without sample path bound:

$$Pr(B(t) > \sigma) \approx \sup_{s \le t} Pr(A(s,t) - C(t-s) > \sigma)$$

- Generally "≥" holds, leading to **lower bound**
- This is done in prior works on heavy-tailed analysis:
  - Gallardo, Makrakis, Orozco-Barbosa 2000, 2005
  - Jiang, Emstad 2005
  - Karasaridis, Hatzinakos, 2001

Construction of sample path bound

- Non-uniform sampling: $x_k = \tau\gamma^k$

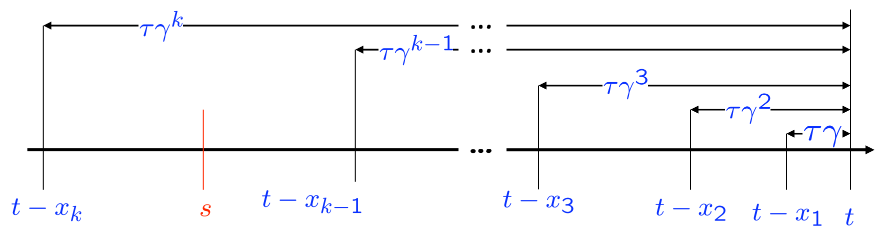

$$A(s,t) - C(t-s) \leq A(t - x_k, t) - Cx_{k-1}$$

Sample path bound

- Backlog bound with discretization:

$$\begin{aligned} Pr(B(t) > \sigma) &= Pr(\sup_{s \leq t}\{A(s,t) - C(t-s)\} > \sigma) \\ &\leq Pr(\sup_{k}\{A(t-x_k,t) - Cx_{k-1}\} > \sigma) \\ &\leq \sum_{k=-\infty}^{\infty} Pr(A(t-x_k,t) > \sigma + Cx_{k-1}) \end{aligned}$$

- Next steps:
  1. Bound individual terms
  2. Bound entire sum

Step 1

- Express $Pr(A(t-x_k,t) > \sigma + Cx_{k-1})$

  using $Pr(A(t-x_k,t) > rx_k + \sigma x_k^H) \leq K\sigma^{-\alpha}$

$$
\begin{aligned}
& Pr(A(t-x_k,t) > \sigma + Cx_{k-1}) \\
& \leq \left[\frac{\sigma}{x_k^H} + x_k^{1-H}(C/\gamma - r)\right]^{-\alpha} \\
& = \left[\sigma\gamma^{-kH} + \gamma^{k(1-H)}(C/\gamma - r)\right]^{-\alpha}
\end{aligned}
$$

Step 2

$$
\begin{aligned}
Pr(B(t) > \sigma) & \leq \sum_{k=-\infty}^{\infty} Pr(A(t-x_k,t) > \sigma + Cx_{k-1}) \\
& \leq \sum_{k=-\infty}^{\infty} \left[\sigma\gamma^{-kH} + \gamma^{k(1-H)}(C/\gamma - r)\right]^{-\alpha} \\
& \leq \underbrace{\sum_{k=-\infty}^{k_0} \sigma^{-\alpha}\gamma^{k\alpha H} + L_1 \sum_{k=k_0}^{\infty} \gamma^{-k\alpha(1-H)}}_{\text{Geometric series}}
\end{aligned}
$$

$$
\leq L_2\, \sigma^{-\alpha}\gamma^{k_0\alpha H} + L_3\, \gamma^{-k_0\alpha(1-H)}
$$

Pick: $\gamma^{k_0} = \sigma$ $\quad \leq L_4\, \sigma^{-\alpha(1-H)}$

ht service curves

- Heavy-tailed (**ht**) statistical service curve:

$$\mathcal{S}(t;\sigma) = [Rt-\sigma]_+\,, \quad \varepsilon(\sigma) = L\sigma^{-\beta}$$

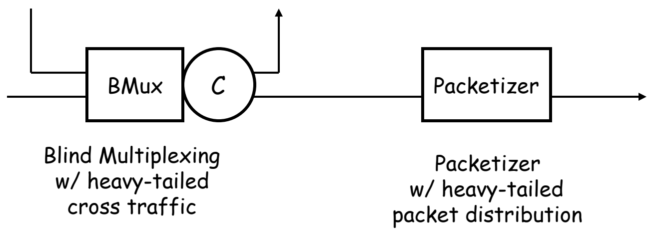

Packetizer

- Packetizer is a delay element that generates packet (burst) traffic from fluid-flow traffic
- Link rate is ***C***
- Utilization is ρ
- Heavy-tailed packet size distribution:
  $Pr(X > \sigma) \leq L\sigma^{-\alpha}$

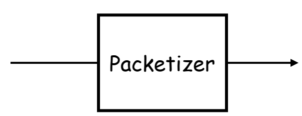

- ht service curve:

$$\begin{aligned} \mathcal{S}(t;\sigma) &= [Ct-\sigma]_+ \\ \varepsilon(\sigma) &= \frac{\rho L}{(\alpha-1)E[X]}\sigma^{-(\alpha-1)} \end{aligned}$$

Leftover Service Curve

- **Blind Multiplexing**: Through traffic has lowest priority and gets leftover capacity
- Link rate is ***C***
- Rate of cross traffic is r_c

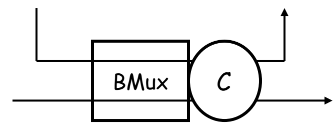

- ht service curve: $$\begin{aligned} \mathcal{S}(t;\sigma) &= [(C - r_c - \mu)t - \sigma]_+ \\ \varepsilon(\sigma) &= \tilde{K}\sigma^{-\alpha(1-H)} \end{aligned}$$

Single Node Delay Bound

- htss envelope: $$\begin{aligned} \mathcal{G}(t;\sigma) &= rt + \sigma t^H \\ \varepsilon(\sigma) &= K\sigma^{-\alpha} \end{aligned}$$
- ht service curve: $$\begin{aligned} \mathcal{S}(t;\sigma) &= [Rt - \sigma]_+ \\ \varepsilon(\sigma) &= L\sigma^{-\beta} \end{aligned}$$

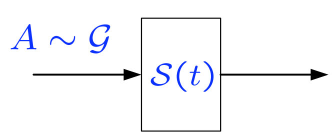

- Delay bound:

$$Pr(W(t) > w) \leq M(Rw)^{-\min\{\alpha(1-H),\beta\}}$$

Example: Node with Pareto Traffic

Traffic parameters:

$$\alpha = 1.6$$
$$b = 150 \; Byte$$
$$\lambda = 75 \; Mbps$$

Node:

- Capacity C=100 Mbps with packetizer
- No cross traffic

Compared with:

- Lower bound from Infocom 2007
- Simulations

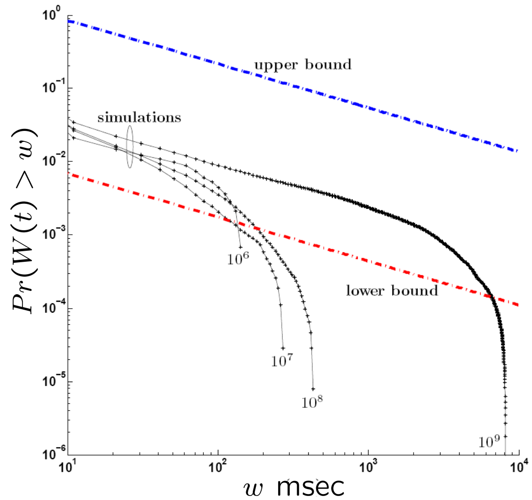

End-to-end Service Curve

Given a path of N nodes with ht service curves

$$\mathcal{S}_n(t;\sigma) = [Rt - \sigma]_+ \,, \quad \varepsilon(\sigma) = L\sigma^{-\beta}$$

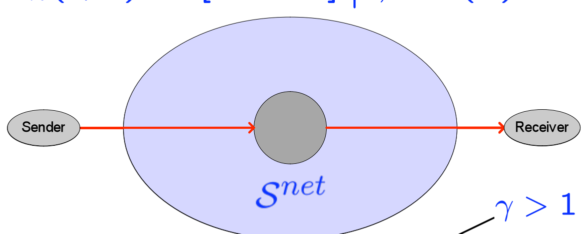

$\gamma > 1$ free choice

Example: Nodes with Pareto Traffic (End-to-end)

Parameters:

$$N = 1, 2, 4, 8$$

Compared with:

- Lower bound from Infocom 2007
- Simulation traces of 10^8 packets

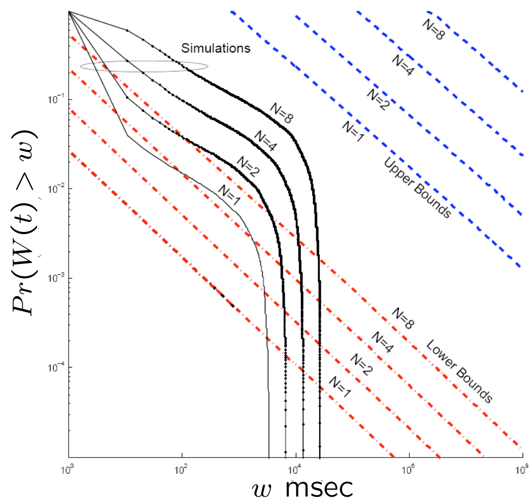

Scaling of delay bounds

- **Single node, large delays** $(w \to \infty)$ **:**

$$\begin{aligned} Pr(W(t) > w) &\leq M(Rw)^{-\alpha(1-H)} \\ &= O\left(w^{-\alpha(1-H)}\right) \end{aligned}$$

Delays scale with same power law as backlog bounds

Scaling of delay bounds

- **Multiple nodes, large delays $(w \to \infty)$:**

$$Pr(W_{net}(t) > w)$$
$$\leq (M_1 \log w + M_2 \log N + M_3) \cdot$$
$$\cdot N^{2+\alpha(1-H)} (Rw)^{-\alpha(1-H)}$$

$$= O\Big(w^{-\alpha(1-H)} \log w\Big)$$

Comparison with queueing theory result (Single Node, $w \to \infty$)

- Cohen 1973:

$$\lim_{t \to \infty} Pr(W(t) > w) \sim \frac{\rho}{1-\rho} c(\alpha) w^{-(\alpha-1)}$$

Same power law decay !

Scaling of delay bounds

- **Multiple nodes, long paths** $(N \to \infty)$:

Define the delay quantile as

$$w_{net}(\varepsilon) = \inf\{w > 0 \mid Pr(W_{net} > w) \le \varepsilon\}$$

We get

$$w_{net}(\varepsilon) = O\Big(N^{\frac{\alpha+1}{\alpha-1}} (\log N)^{\frac{1}{\alpha-1}}\Big)$$

$(H = 1/\alpha)$

Illustration of scaling bounds

Upper Bound: $\mathcal{O}(N^{\frac{\alpha+1}{\alpha-1}} (\log N)^{\frac{1}{\alpha-1}})$

Lower Bound: $\Theta(N^{\frac{\alpha}{\alpha-1}})$

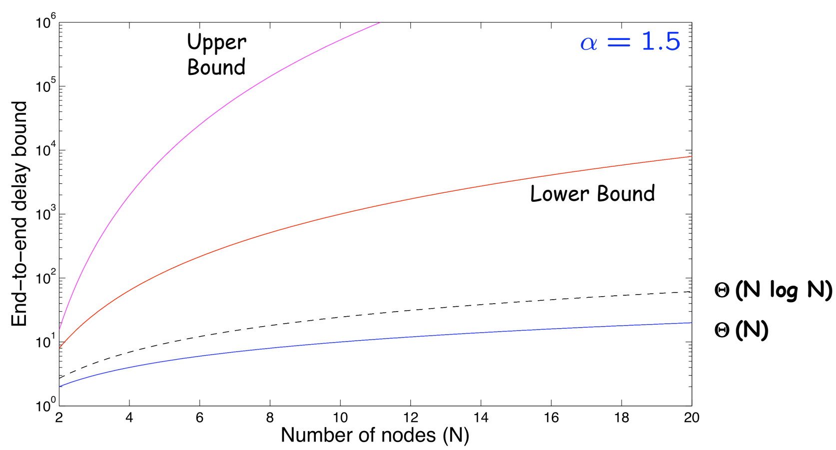

Conclusions

- First non-asymptotic end-to-end delay bounds for heavy-tailed traffic
  - Upper bounds
  - Lower bounds (from Infocom 2007)

- Rigorous extension of network calculus to heavy-tailed traffic

- Observation:
  Bounds are loose, but simulations are useless

Technical Report: **arXiv:0911.3856v1**